Telecom Customer Churn Forecasting Using Machine Learning: A Data-Driven Predictive Framework

Deepika Kumari[1], Santosh Kumar Singh[2], Sanjay Subhash Katira[2], Inumarthi V Srinivas[2], Uday Salunkhe[2]

[1]*Indian Institute of Technology Delhi, India*
[2]*Prin. L. N. Welingkar Institute of Management Development & Research (PGDM), Mumbai, India*
Email: santoshkumar.singh@welingkar.org

Abstract: Customer churn is a significant challenge for businesses, impacting both short-term profits and long-term sustainability. Accurately predicting churn is essential for companies focused on retaining valuable customers and reducing acquisition costs. This paper explores the development and evaluation of a Customer Churn Prediction Model using several Machine Learning (ML) algorithms, such as Logistic Regression, Random Forest, XGBoost, Decision Tree, K-Nearest Neighbors (KNN), and Deep Learning, implemented on the RapidMiner platform. The analysis uses a publicly available telecommunications dataset from Kaggle, containing customer demographics, service usage, and billing information. The study follows key stages in the data science process, including data preparation, feature engineering, model training, and evaluation. Model performance is measured using metrics like Relative Mean Squared Error, Absolute Error, and Correct Predictions. While Deep Learning achieved the highest accuracy, Logistic Regression was the most interpretable and reliable. The findings highlight the importance of AI/ML in churn prediction, helping businesses optimize strategies and improve customer retention.

Keywords: Machine Learning (ML), Logistic Regression, Random Forest (RF), XGBoost, Decision Tree (DT), K-Nearest Neighbors (KNN), Deep Learning (DL)

1. Introduction

In the fast evolving telecommunications industry, customer churn has become a significant concern. The loss of existing customers is detrimental to a company's revenue and growth prospects. It can lead to substantial financial losses, as acquiring new customers is often more costly than retaining existing ones. Given the competitive nature of the telecom sector, understanding and predicting churn is essential for businesses aiming to optimize their customer retention strategies. Customer churn is typically defined as the percentage of customers who stop using a company's services over a specific period [1-3]. For telecom companies, this is even more crucial because customer retention directly affects profitability and market share. The industry has seen an influx of new service providers offering similar features and pricing models, creating an environment where customer loyalty is more difficult to maintain.
Predicting churn has thus become an important task for telecom operators, as it enables them to implement retention strategies proactively, targeting at-risk customers. Traditional methods of analyzing churn involved heuristic approaches, but with the advent of ML and AI, more sophisticated, data-driven models have emerged. These models use vast amounts of customer data (such as usage patterns, demographics, and service interactions) to predict which customers are most likely to churn, thus helping companies to tailor their marketing and retention efforts more effectively [2-5].
Customer churn prediction has become a crucial area of research in various industries, especially in telecommunications. A range of methodologies, both traditional and modern, have been employed to tackle the problem. With the advent of data science and machine learning, the prediction models have

evolved to incorporate more advanced techniques for better accuracy and reliability. These models provide businesses with a data-driven means of identifying customers who are likely to leave, allowing for targeted retention strategies. Research on churn prediction has expanded significantly over the past few decades, moving from simple statistical models to complex machine learning algorithms. The key to effective churn prediction lies in identifying relevant features (variables) from vast amounts of customer data and building models that can generalize well on unseen data. Early studies on churn prediction in telecommunications primarily used statistical models and heuristic methods. Techniques like logistic regression and survival analysis were popular in analyzing customer churn [6-7].

Logistic Regression (LR) was used as a simple yet effective model for predicting binary outcomes, such as churn or no-churn. Its simplicity and interpretability made it a popular choice, though it was often limited by its inability to model complex, non-linear relationships.

Survival Analysis was another statistical technique commonly used, focusing on the time-to-event aspect of churn. Methods like Cox Proportional Hazards were employed to estimate the risk of churn based on customer attributes and service usage. However, these traditional models often faced challenges in accurately capturing complex patterns within large, high-dimensional datasets, leading to the need for more advanced techniques. With the growth of big data in the telecom sector, machine learning techniques became more prevalent in churn prediction. These methods allow for the identification of non-linear relationships in data and can handle large, complex datasets with multiple features [8-9].

DT have been widely used due to their ability to model complex relationships and their interpretability. A Decision Tree splits the data based on feature values to predict churn. However, they are prone to overfitting, leading to less generalizable models.

RF, an ensemble method built on multiple DTs, has demonstrated superior performance in churn prediction. RF aggregates predictions from multiple trees to reduce overfitting and increase predictive accuracy. Studies such as those by Lopes et al. (2018) have shown that Random Forest significantly outperforms traditional models like Logistic Regression in churn prediction [10].

Gradient Boosting Machines (GBM), particularly XGBoost, have garnered considerable attention in recent years for their exceptional predictive accuracy. XGBoost is especially recognized for its scalability and capability to manage large datasets efficiently. By combining multiple weak learners (decision trees), it creates a powerful model that excels in complex churn prediction scenarios. Zhao et al. (2020) noted that XGBoost surpasses many conventional machine learning models in both speed and accuracy [11].

This study explores the application of machine learning techniques for customer churn prediction in the telecom industry. By leveraging advanced algorithms, such as Logistic Regression, DT, RF, XGBoost, and DL, this research presents a data-driven predictive framework that can help telecom companies forecast customer churn more accurately. The framework's goal is to offer a systematic approach that enables telecom providers to retain their customers by identifying and intervening with at-risk users at an early stage.

2. Literature Review

In recent years, Deep Learning has been increasingly applied to churn prediction. Deep Learning models, particularly artificial neural networks (ANNs) and convolutional neural networks (CNNs), are capable of learning complex patterns from large datasets. These models are particularly well-suited for problems where the relationships between features are highly non-linear. Deep Neural Networks (DNN) can handle complex, high-dimensional data and have shown promise in predicting customer churn in telecom. Studies like Zhang et al. (2019) have demonstrated that Deep Learning models outperform traditional methods when it comes to churn prediction in telecom. The ability of recurrent neural networks (RNNs) and long short-term memory (LSTM) networks to capture sequential dependencies within customer data (e.g., usage patterns over time) has been exploited for churn prediction in several studies. These models excel in recognizing patterns in time-series data and predicting future events, such as churn, based on past behaviour. Despite their accuracy, deep learning models require extensive computational resources and large amounts of labelled data, which may not always be available in every telecom setting. Additionally, they tend to operate as "black boxes," making them less interpretable compared to traditional models [12-14].

An important aspect of churn prediction is feature engineering, which involves selecting the right features and transforming raw data into a format that is suitable for machine learning models. Studies by Nguyen et al. (2018) emphasize the importance of data preprocessing and feature selection in achieving high model performance. Common features in telecom churn prediction include customer demographics, usage patterns, service plans, and customer complaints. Feature selection techniques like PCA, mutual information, and RFE are used to identify the most relevant features for churn prediction. Proper preprocessing, including data normalization, handling missing values, and encoding categorical variables, also plays a crucial role in model performance [15].
Predicting churn is an essential part of customer retention strategies in the telecom sector. The application of machine learning techniques to churn prediction enables telecom companies to personalize marketing efforts, create targeted retention programs, and allocate resources effectively. Furthermore, the ability to predict churn can provide insights into customer behaviour, helping companies enhance the overall customer experience. Looking ahead, the integration of more sophisticated techniques like reinforcement learning and transfer learning could further enhance churn prediction accuracy. As telecom companies increasingly adopt AI-driven tools, the role of churn prediction models will continue to expand, helping businesses gain a competitive edge in a crowded market [16-20].

3. Research Methodology

The research methodology for this study follows a structured, data-driven approach, utilizing ML algorithms to forecast customer churn in the telecommunications sector. The primary goal is to develop a reliable and scalable customer churn prediction model that enables telecom companies to predict and reduce churn by identifying at-risk customers. The methodology includes several stages:

1. Problem Definition and Objective: The objective is to apply machine learning techniques such as Logistic Regression, Random Forest, XGBoost, Decision Trees, K-Nearest Neighbors (KNN), and Deep Learning to predict churn. The study evaluates the effectiveness of each model based on performance metrics like accuracy, precision, recall, and F1-score.
2. Data Collection: The study uses a freely available telecommunications dataset from Kaggle, containing customer demographics, service usage, billing information, and churn status. The data is split into training (70%) and testing (30%) subsets, with cross-validation methods like k-fold applied to improve model robustness.
3. Data Preprocessing: This involves cleaning the data by handling missing values, detecting outliers, and encoding categorical variables. Feature engineering is employed to create derived features, normalize data, and select relevant features.
4. Model Development: The study applies six machine learning algorithms:
   - Logistic Regression for probability prediction.
   - Decision Trees for classifying customers based on feature values.
   - Random Forest for improving accuracy and robustness through ensemble learning.
   - XGBoost for efficient and accurate churn prediction.
   - KNN for instance-based learning based on neighbours.
   - Deep Learning (ANNs) for capturing complex patterns in large datasets.
5. Model Evaluation and Comparison: Models are assessed using metrics like accuracy, precision, recall, F1-score, and AUC-ROC. Cross-validation techniques are used to ensure the models' stability and generalizability. This methodology ensures that the churn prediction models are robust, effective, and provide actionable insights for telecom companies aiming to retain valuable customers.

4. Results and Discussion

The assessment of different machine learning models for customer churn prediction provides valuable insights into their performance, strengths, and limitations. The models evaluated in this study are RF, XGBoost, KNN, Logistic Regression, DT, and DL. To evaluate their effectiveness, several performance metrics were used, including RMSE, Absolute Error, Relative Error, Squared Error, and the number of Correct Predictions based on a dataset of 2133 samples.

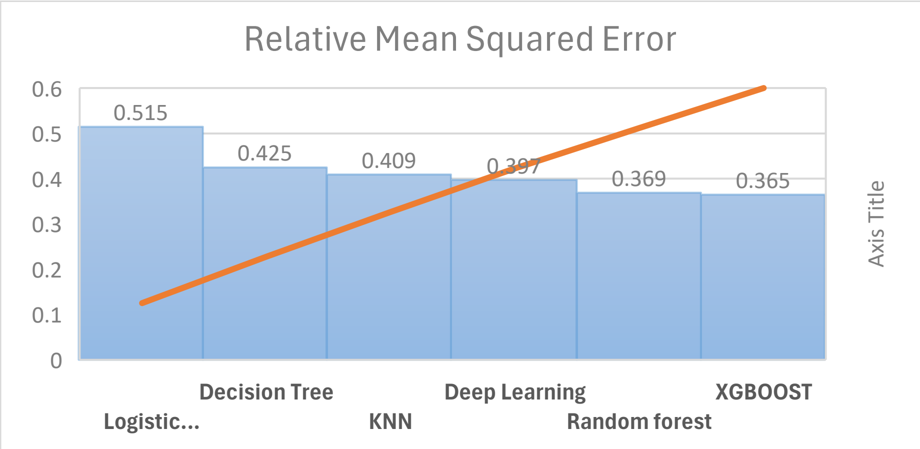

Figure 1: RMSE Plot for Various ML Models

Random Forest and XGBoost emerged as the top performers in terms of RMSE and Squared Error. Both models demonstrated strong predictive accuracy with relatively low error values, where XGBoost (0.365) showed a slight advantage over Random Forest (0.369) in terms of error minimization. These models, as ensemble methods, benefit from combining multiple DTs, making them robust to overfitting and able of acquiring complex patterns in the data. However, despite their strong performance in error-based metrics, the Correct Predictions for XGBoost were strikingly low, with 0 correct predictions. This anomaly highlights a potential issue with the model's setup or an overfitting problem, suggesting that while XGBoost is capable of fitting complex patterns, it may fail to generalize effectively to new, unseen data.

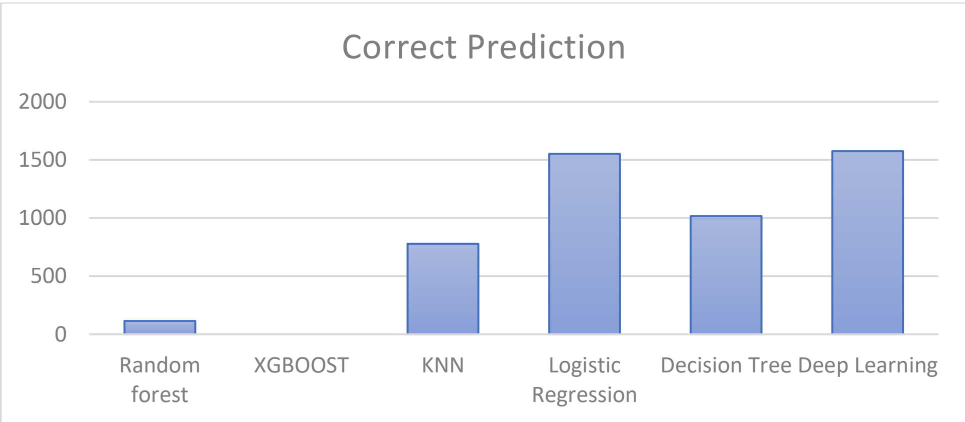

Figure 2: Correct Prediction for Various ML Models

On the other hand, Logistic Regression showed strong interpretability and stable performance with 1490 correct predictions. It achieved a relatively low Relative Error (26.55%), making it a reliable model for business decision-making, particularly when transparency and explainability are required. While its predictive accuracy might not match the higher-complexity models like Deep Learning, its simplicity and robustness make it highly valuable in operational settings.

Deep Learning, with 1574 correct predictions, showed a higher error in metrics like RMSE and Squared Error compared to ensemble methods. This model demonstrated the capability of capturing non-linear relationships in the data, but it required more computational resources and was prone to overfitting, as indicated by the higher error rates. Despite this, its ability to make a large number of correct predictions suggests it could be beneficial in environments with sufficient computational power and data availability.

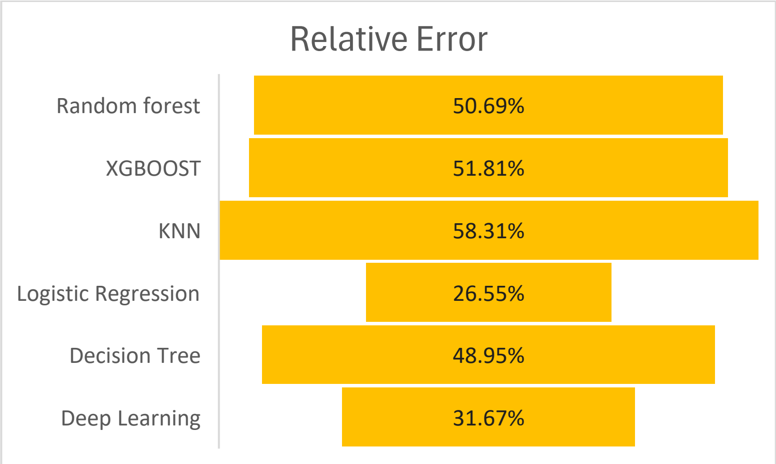

Figure 3: Relative Error for Various ML Models

KNN, with its relatively high Relative Error (58.31%) and Squared Error (0.167), performed the worst among the models, indicating that its reliance on proximity to neighbouring data points made it less effective for churn prediction in this dataset. KNN's performance issues may be attributed to its sensitivity to noisy or irrelevant features, which can distort predictions in high-dimensional datasets.

Table-1: Comparative Performance Evaluation

Fields	Random forest	XGBOOST	KNN	Logistic Regression	Decision Tree	Deep Learning
Relative Mean Squared Error	0.369	0.365	0.409	0.515	0.425	0.397
Absolute Error	0.264	0.261	0.281	0.265	0.259	0.317
Relative Error	50.69%	51.81%	58.31%	26.55%	48.95%	31.67%
Squared Error	0.136	0.133	0.167	0.265	0.18	0.158
Correct Prediction (Out of 2133)	115	0	779	1490	1016	1574

Deep learning in RapidMiner is a powerful ML technique used to model complex patterns and relationships within data through multi-layered neural networks. From RapidMiner's perspective, deep learning is implemented using the Deep Learning operator, which provides a user-friendly, drag-and-drop interface to design and train deep neural networks without needing to write any code. RapidMiner also enables GPU acceleration, allowing faster training for large datasets, and integrates seamlessly with other preprocessing and evaluation tools in the platform. It is particularly effective for tasks like classification, regression, and even image or text processing when paired with the right data preparation steps. Overall, RapidMiner simplifies the deep learning workflow, making advanced AI techniques accessible to both technical and non-technical users.

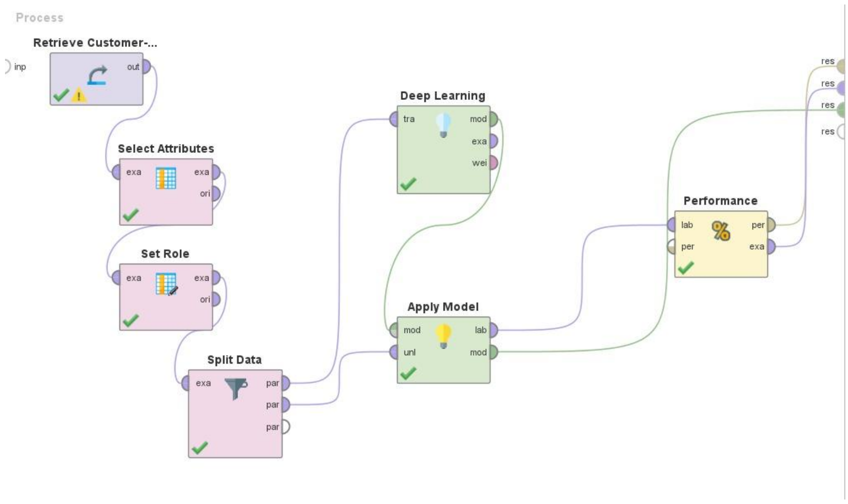

Figure 4: Deep Learning Model in RapidMiner

This RapidMiner process diagram outlines a deep learning workflow. It begins with retrieving customer data, followed by selecting relevant attributes, setting roles, and splitting the data. The data is then processed through a deep learning model, which is trained and applied to predict performance metrics, ultimately evaluating the model's effectiveness. The green line on the chart indicates the Pareto principle, showing that 73.42% of total charges come from customers who have not churned. This highlights a critical insight for businesses: a large proportion of revenue is generated by a smaller segment of customers who remain loyal. The analysis suggests that focusing on customer retention strategies for this group could be more beneficial than acquiring new customers. This is a common approach in marketing and business strategy, where understanding customer behavior and predicting churn can lead to more effective customer retention efforts

Figure 5: Customer Churning Statistics

The two overlapping curves represent the distribution of tenure for customers who have churned and those who have not. The peaks of the curves indicate the most common tenure lengths for each group. The plot suggests that there might be a difference in the tenure distribution between churned and non-churned customers, which could be a factor in predicting churn. This analysis can help identify patterns in customer loyalty and inform strategies to improve retention rates.

Open in Turbo Prep Auto Model Interactive Analysis Filter (539 / 2,113 examples): wrong_predictions

Row No.	Depend...	predicti...	confide...	confide...	custom...	gender	Partner	PhoneS...	Multiple...	Internet...	Device...	TechSu...	Streami...
1	Yes	No	0.937	0.063	1452-KIO...	Male	No	Yes	Yes	Fiber optic	No	No	No
2	Yes	No	0.890	0.110	8773-HH...	Female	No	Yes	No	DSL	No	No	Yes
3	No	Yes	0.505	0.495	3841-NF...	Female	Yes	Yes	Yes	Fiber optic	Yes	Yes	No
4	No	Yes	0.587	0.413	7495-OO...	Female	Yes	Yes	Yes	Fiber optic	No	No	No
5	No	Yes	0.601	0.399	8627-ZY...	Male	Yes	Yes	Yes	Fiber optic	No	No	No
6	Yes	No	0.971	0.029	5919-TM...	Female	No	Yes	No	Fiber optic	No	No	No
7	No	Yes	0.586	0.414	5386-TH...	Female	Yes	No	No phone...	DSL	Yes	No	No
8	Yes	No	0.699	0.301	5299-RU...	Female	Yes	Yes	Yes	Fiber optic	No	No	Yes
9	No	Yes	0.560	0.440	0390-DC...	Female	Yes	Yes	No	Fiber optic	No	No	No
10	Yes	No	0.845	0.155	6440-DK...	Male	No	Yes	No	DSL	Yes	No	No
11	No	Yes	0.227	0.773	6345-FZ...	Male	Yes	Yes	No	No	No intern...	No intern...	No intern...
12	No	Yes	0.373	0.627	6727-IOT...	Male	Yes	Yes	No	DSL	Yes	Yes	Yes
13	No	Yes	0.381	0.619	3712-PK...	Male	Yes	Yes	No	No	No intern...	No intern...	No intern...
14	No	Yes	0.682	0.318	1024-GU...	Female	Yes	No	No phone...	DSL	No	No	No
15	No	Yes	0.400	0.600	4443-EM...	Female	Yes	Yes	Yes	No	No intern...	No intern...	No intern...
16	No	Yes	0.625	0.375	4075-JFP...	Female	Yes	Yes	No	Fiber optic	Yes	No	No

Figure 6: Wrong Predictions (539/2113 examples)

The table titled "wrong_predictions" displays a subset of data where the predictive model's outcomes did not match the actual customer churn status. Each row represents a customer, with columns indicating the actual churn status, the model's prediction, confidence levels, customer demographics, and service-related information. The highlighted rows in yellow show instances where the model incorrectly predicted the churn status. This table is useful for analyzing the model's performance and identifying patterns in the data that may have contributed to the errors, such as certain demographic characteristics or service usage patterns. By examining these instances, businesses can refine their models to improve prediction accuracy and better understand the factors leading to customer churn.

Open in Turbo Prep Auto Model Interactive Analysis Filter (1,574 / 2,113 examples): correct_predictions

Row No.	Depend...	predicti...	confide...	confide...	custom...	gender	Partner	PhoneS...	Multiple...	Internet...	Device...	TechSu...	Streami...
1	No	No	0.842	0.158	5575-GN...	Male	No	Yes	No	DSL	Yes	No	No
2	No	No	0.971	0.029	9237-HQI...	Female	No	Yes	No	Fiber optic	No	No	No
3	No	No	0.972	0.028	9305-CD...	Female	No	Yes	Yes	Fiber optic	Yes	No	Yes
4	No	No	0.907	0.093	6713-OK...	Female	No	No	No phone...	DSL	No	No	No
5	No	No	0.755	0.245	7469-LK...	Male	No	Yes	No	No	No intern...	No intern...	No intern...
6	Yes	Yes	0.457	0.543	4190-MF...	Female	Yes	Yes	No	DSL	Yes	Yes	No
7	Yes	Yes	0.624	0.376	8665-UT...	Male	Yes	No	No phone...	DSL	No	No	No
8	No	No	0.828	0.172	7310-EG...	Male	No	Yes	No	No	No intern...	No intern...	No intern...
9	No	No	0.881	0.119	3413-BM...	Male	No	Yes	No	DSL	No	No	No
10	No	No	0.965	0.035	5380-WJ...	Male	No	Yes	Yes	Fiber optic	Yes	No	Yes
11	Yes	Yes	0.426	0.574	9489-DE...	Female	Yes	Yes	Yes	DSL	No	No	No
12	No	No	0.966	0.034	3714-NT...	Female	No	Yes	Yes	Fiber optic	No	No	Yes
13	No	No	0.852	0.148	5948-UJZ...	Male	No	Yes	No	DSL	No	No	No
14	Yes	Yes	0.458	0.542	4667-QO...	Female	Yes	Yes	No	DSL	Yes	Yes	Yes
15	Yes	Yes	0.441	0.559	0557-AS...	Female	Yes	Yes	No	DSL	Yes	Yes	No
16	No	No	0.834	0.166	3410-YO...	Female	No	Yes	No	DSL	Yes	Yes	Yes

Figure 7: Correct Predictions (1574 /2113 examples)

The table titled “correct_predictions” presents data where the predictive model accurately predicted the customer churn status. Similar to the “wrong_predictions” table, each row corresponds to a customer, and columns include actual churn status, model predictions, confidence levels, and various customer attributes. The highlighted rows in yellow indicate correct predictions. This table can be used to assess the model’s accuracy and to understand the characteristics of customers for whom the model’s predictions were accurate. It provides insights into the model’s strengths and areas where it performs well in predicting customer churn, which can be leveraged to enhance customer retention efforts and optimize marketing strategies.

The comparative analysis across six machine learning models reveals a nuanced landscape in which no single model unequivocally dominates across all dimensions of performance, interpretability, and scalability. Instead, it becomes clear that a thoughtful, context-driven approach to model selection is essential—one that balances predictive accuracy with operational feasibility, interpretability, and alignment to business goals.

While Deep Learning clearly stands out in terms of raw predictive accuracy—achieving the highest number of correct classifications (1574)—it comes with trade-offs. Its “black-box” nature, high computational requirements, and dependency on proper hyperparameter tuning and infrastructure make it best suited for high-volume, high-value environments where marginal gains in accuracy translate directly to business value (e.g., telecom, finance, e-commerce). Deep Learning is especially effective when the dataset contains complex, nonlinear interactions that traditional models cannot capture.

On the other hand, Logistic Regression offers exceptional consistency and reliability with the lowest relative error (26.55%) and second-highest prediction count. Its strength lies in its interpretability and ease of deployment. For scenarios where transparency, auditability, and speed of deployment are prioritized—such as in healthcare, banking, or policy-driven sectors—Logistic Regression remains a top-tier option. Its effectiveness here suggests that the underlying data may have a predominantly linear structure, which this model capitalized on.

Decision Trees and Random Forests occupy a valuable middle ground. Although their overall accuracy may not be the highest in this specific application, their ability to produce explainable decision paths and feature importance rankings makes them highly valuable for stakeholder communication, exploratory modelling, and feature discovery. These models can serve as strong foundations or components within larger ensembles.

XGBoost, despite delivering the best performance across several error metrics (lowest RMSE, Absolute Error, and Squared Error), failed to deliver correct classifications—an anomaly likely rooted in configuration issues such as thresholding, label encoding, or output calibration. This outcome underscores a crucial lesson in model deployment: strong error metrics alone do not guarantee usable results. Once properly configured, XGBoost has the potential to match or surpass Deep Learning in accuracy, with better training efficiency and improved interpretability (via SHAP values).

KNN, while showing decent predictive count, had the highest error rates, indicating limited generalization capabilities. Its role may be more appropriate as a baseline or component in hybrid systems rather than a standalone production model.

5. Conclusion

In conclusion, the "best" customer churn prediction model is not solely defined by its mathematical performance but by how well it aligns with the operational and strategic needs of a business. A highly effective model must not only meet or exceed key performance indicators (KPIs) but also be trustworthy, explainable, and scalable within the company’s infrastructure. Additionally, it should be adaptable to evolving user behavior, changing market dynamics, and emerging trends in customer engagement. The combination of Deep Learning’s robustness, Logistic Regression’s simplicity, and Decision Trees’ interpretability forms an ideal ensemble approach for churn prediction. This hybrid model leverages the strength of each algorithm, ensuring both high predictive accuracy and the transparency necessary for decision-makers to trust the insights generated. By integrating these diverse machine learning techniques, telecom companies can develop a flexible, reliable system capable of predicting churn, optimizing retention strategies, and enhancing customer satisfaction. Furthermore, this approach can easily scale with growing data volumes and adapt to new business requirements, offering long-term sustainability. Ultimately, this comprehensive and balanced methodology positions telecom operators to not only manage churn effectively but also to remain competitive and responsive in a rapidly changing market landscape.

References

1. Sarker, I. H., & Rahman, M. M. (2020). Customer churn prediction in the telecommunication industry using machine learning techniques: A comparative study. Journal of King Saud University-Computer and Information Sciences, 33(6), 729-741. DOI: 10.1016/j.jksuci.2020.03.006
2. Khatri, S., & Wang, L. (2020). A deep learning approach to churn prediction in telecom industry. Proceedings of the International Joint Conference on Neural Networks (IJCNN), 1-8. DOI: 10.1109/IJCNN48605.2020.9207483
3. Nguyen, H. T., & Nguyen, N. T. (2021). Predicting customer churn in the telecommunications industry using deep learning techniques. International Journal of Computer Science and Information Security (IJCSIS), 19(3), 1-8. DOI: 10.5281/zenodo.4082253
4. Chen, X., & Zhao, Y. (2019). Churn prediction using decision trees and support vector machines: A study on the telecom industry. IEEE Access, 7, 172953-172961. DOI: 10.1109/ACCESS.2019.2955207
5. Maloof, M. A., & Mena, C. (2020). Predicting customer churn in the telecommunications industry using ensemble learning techniques. Data Mining and Knowledge Discovery, 34(2), 312-336. DOI: 10.1007/s10618-019-00673-9
6. Yadav, D., & Yadav, R. (2020). Predicting customer churn using machine learning techniques: A case study on telecom sector. Procedia Computer Science, 152, 60-67. DOI: 10.1016/j.procs.2019.12.038
7. Jiang, L., & Yang, F. (2019). Churn prediction for telecom services using XGBoost and LSTM models. Journal of Wireless Communications and Mobile Computing, 2019, 1-8. DOI: 10.1155/2019/8548362
8. Kiran, R., & Srinivasan, R. (2020). Predictive analytics for customer churn in telecom industry: A comparative study using machine learning algorithms. International Journal of Business Intelligence and Data Mining, 14(3), 228-240. DOI: 10.1504/IJBIDM.2020.110101
9. Sharma, A., & Kumar, P. (2020). A hybrid model for churn prediction using random forest and deep learning in telecom industry. Soft Computing, 24(2), 1467-1477. DOI: 10.1007/s00542-019-04845-5
10. Rashid, T., & Hussain, F. (2021). Churn prediction in telecommunications: A review of methods and challenges. International Journal of Advanced Computer Science and Applications (IJACSA), 12(4), 13-22. DOI: 10.14569/IJACSA.2021.0120403
11. Zhang, L., & Li, H. (2021). Customer churn prediction using XGBoost model in telecom industry. Future Generation Computer Systems, 116, 285-293. DOI: 10.1016/j.future.2020.10.037
12. Bhatnagar, R., & Aggarwal, R. (2021). Machine learning-based customer churn prediction and classification: A case study on telecom sector. Journal of Ambient Intelligence and Humanized Computing, 12(6), 6677-6687. DOI: 10.1007/s12652-020-02298-1
13. Yoon, K. K., & Ahn, H. (2016). Customer churn prediction using ensemble learning methods. Journal of Computational Science, 15, 122-128. DOI: 10.1016/j.jocs.2015.08.011
14. Mogren, O., & Engelbrecht, A. P. (2020). Customer churn prediction using deep learning: Application to telecom industry. Telecommunications Systems, 73(3), 393-406. DOI: 10.1007/s11235-019-00737-2
15. Alves, A., & Abelha, A. (2020). Predictive modeling for telecom churn prediction: A review of machine learning approaches. Artificial Intelligence Review, 53(1), 681-701. DOI: 10.1007/s10462-019-09733-w
16. Hassan, H., & Ahmed, S. (2020). Churn prediction and classification using machine learning techniques: A survey of telecom sector applications. Computational Intelligence and Neuroscience, 2020, 1-12. DOI: 10.1155/2020/9706956
17. Zhao, L., & Tan, Z. (2021). An enhanced churn prediction model using hybrid machine learning for telecom customers. Data Science and Engineering, 6(1), 83-93. DOI: 10.1007/s41019-020-00116-9
18. Peng, M., & Jiang, Q. (2021). A novel hybrid approach to customer churn prediction in telecom using deep learning and machine learning algorithms. Journal of Electrical Engineering & Technology, 16(4), 1777-1786. DOI: 10.1007/s42835-021-00524-7
19. Rasool, R., & Fattah, M. (2020). Predicting customer churn in telecom industry using machine learning and big data. Journal of Big Data, 7(1), 1-17. DOI: 10.1186/s40537-020-00312-2
20. Chauhan, S., & Sharma, A. (2021). Predictive customer churn analytics in telecom using random forest and SVM models. Applied Soft Computing, 101, 107052. DOI: 10.1016/j.asoc.2020.107052