



UGANDA CHRISTIAN
UNIVERSITY

A Centre of Excellence in the Heart of Africa

Department of Computing & Technology

A MULTI-CLASS MACHINE LEARNING
CUSTOMER CLASSIFICATION FOR
PERSONALISED VOICE BUNDLE
RECOMMENDATIONS:
A CASE STUDY OF AIRTEL UGANDA LIMITED

By

Remmy Bisimbeko

Reg No: J24M19/011 Student No: B26099

E-mail: rbisimbeko@ucu.ac.ug

A dissertation submitted to the Faculty of Engineering Design
& Technology in partial fulfillment of the requirements for the
award of **Master of Science in Data Science and
Analytics** of Uganda Christian University

Supervised by: Dr. Daphne Nyachaki Bitalo

August 2026

Declaration

I, Remmy Bisimbeko, declare to the best of my knowledge that the work in this dissertation is original and has not been submitted before to any university/institution. It is the result of my own effort.

Student Name: Remmy Bisimbeko

Access Number: B26099

With regard to the use of Artificial Intelligence (AI) tools:

- (i) The use of AI tools is restricted to tasks such as brainstorming and language improvement (grammar, spelling, sentence structure). AI tools should not be used to generate content that is presented as the student's own original intellectual work.
- (ii) Where AI has been used as guided in (i) above, full disclosure of how it has been used is made in Appendix **G**, including how AI was used, the AI tools (including model versions), and a compilation of major AI prompts used.

Signature:

A handwritten signature in black ink, appearing to read 'Remmy Bisimbeko', with a stylized flourish underneath.

Date: August 2026

Ethical Generative AI Usage Declaration

Students may use generative artificial intelligence (AI) ethically as a “study buddy” or assistant rather than a replacement for critical thinking. This includes citing AI use and fact-checking all outputs to maintain academic integrity.

For purposes of this assignment, the following guidelines shall apply.

- i. Generative AI may be used for brainstorming, and generating ideas to improve your work. You may also use AI for literature search, and summarizing journal papers for quick analysis and assessing relevance of publications. However, once the relevant publications have been identified, a student is expected to fully review the selected publications, identifying problems addressed, methods used, gaps identified and recommendations made.
- ii. Where AI has been used as guided in (i) above, full disclosure of how it has been used should be made in the appendix **G**, including how AI was used, the AI tools (including model versions), and a compilation on major AI prompts used.
- iii. NO GENERATIVE AI CONTENT IS ALLOWED IN THE FINAL SUBMISSION. The final submission SHOULD be your original work.

A. Declaring the non-use of AI for content generation

I, Remmy Bisimbeko, declare that no content generated by AI technologies has been used in this dissertation.

Signature:



Date: August 2026

B. Acknowledging the use of AI as a study buddy

Please refer to Appendix G for the complete declaration of AI usage for generating ideas, brainstorming, literature search, and summarizing only to improve your work.

Approval

This dissertation has been submitted for examination with approval of the following supervisor.

Student Name: Remmy Bisimbeko **Supervisor:** Dr. Daphne N. Bitalo

Signature: 


Signature: _____

Date: August 2026

Date: August 2026

Abstract

Airtel Uganda Limited relies on Average Revenue Per User (ARPU)-banded segmentation to deliver voice bundle promotions, yielding conversion rates below 3 percent and wasting an estimated 20 percent of the marketing budget on untargeted communications. The absence of a personalised, data-driven recommendation system represents both a commercial inefficiency and an academic gap this study addresses.

This study developed and validated a machine learning recommendation framework using real transactional data from 1,458,900 Airtel Uganda prepaid subscribers (February–April 2026). Four classification models — Logistic Regression, Decision Tree, Random Forest, and XGBoost — were evaluated following a preprocessing pipeline of median imputation, capped SMOTE oversampling, and behavioural feature engineering. Three statistical tests were applied to evaluate relationships between customer attributes and bundle segment: Spearman Rank Correlation to measure the strength and direction of monotonic associations between numeric features and the target; one-way ANOVA to test whether feature means differ significantly across the eight value segment groups; and Chi-Square to assess the independence of categorical features from the target variable.

Statistical tests confirmed 34 of 35 features significantly associated with bundle segment ($p < 0.05$); UPGRADE_RATE was the only non-significant feature. Feature importance analysis identified TOTAL_REVENUE (18.01%), SPEND_PER_MINUTE (6.57%), CALLS (5.97%), VOICE_INTENSITY (5.87%), and MOBILE_MONEY (4.64%) as the top five predictors. XGBoost achieved the highest F1-Score of 0.661 on the 8-class hold-out test set — a 5.4-fold improvement over a random baseline.

A simulated pilot projected a four- to five-fold improvement in conversion rate (3% to 12–15%) and a 9 percent ARPU uplift (simulated pilot projection). Three strategic recommendations were formulated: a VOICE_INTENSITY-triggered recommendation engine, enriched segmentation incorporating device type and purchase channel, and a formal A/B test with model governance. This is the first empirically validated ML bundle recommendation framework built on real Ugandan telecom data.

Keywords: Machine Learning, Telecommunications, Voice bundle recommendation, Customer behavioural attributes, Predictive Analytics.

Dedication

*To my former boss, Kenneth Kyobe,
who gave me a chance before I had earned it.*

Contents

Declaration	i
AI Usage Declaration	ii
Approval	iv
Abstract	v
Dedication	vii
List of Figures	xiii
List of Tables	xiv
List of Abbreviations	xv
Acknowledgements	xvi
1 Introduction	1
1.1 Background to the Study	1
1.2 Problem Statement	2
1.3 Justification of the Study	4
1.3.1 Justification for the Telecom Industry	4
1.3.2 Justification for the Customer	4
1.3.3 Justification for Academia	5
1.4 Main Objective of the Research	6
1.5 Specific Objectives	6
1.6 Research Questions	7
1.7 Research Hypotheses	7
1.8 Scope of the Study	8
1.8.1 Geographical Scope	8
1.8.2 Time Scope	8

1.8.3 Content Scope	8
1.9 Significance of the Study	9
1.10 Theoretical and Conceptual Framework	9
1.11 Limitations of the Study	11
1.12 Organisation of the Thesis	12
2 Literature Review	13
2.1 Introduction	13
2.2 Theoretical Framework	13
2.2.1 Consumer Behaviour Theory and Telecom Service Adoption . .	13
2.2.2 Machine Learning as a Predictive Decision-Making Framework .	14
2.3 Key Customer Attributes and Behavioural Patterns in Telecom	15
2.3.1 Global Perspective: What Drives Customer Bundle Purchasing Decisions?	15
2.3.2 Regional Perspective: Customer Attributes in African Telecom Markets	17
2.3.3 Ugandan Context: Customer Attributes and Service Loyalty . .	18
2.4 Machine Learning Models for Prediction, Classification and Data Pre- processing	19
2.4.1 Classification Models for Customer Behaviour Prediction: Global Evidence	19
2.4.2 Customer Segmentation Using Clustering Algorithms	20
2.4.3 Data Preprocessing: The Foundation of Model Quality	20
2.4.4 Model Evaluation Metrics	21
2.4.5 Recommender System Paradigms and the Classification-Based Approach	22
2.5 Empirical Review of Key Methodologies	23
2.6 Recommender Systems in Telecommunications	24
2.6.1 Global Approaches and Evidence	24
2.6.2 Recommender Systems in Africa	25
2.6.3 Current Recommendation Practices in Uganda	25
2.7 Impact of Personalised Recommendations on Revenue and Customer Satisfaction	26
2.7.1 Global Evidence on Business Outcomes	26
2.7.2 Revenue Optimisation Through ARPU Enhancement	27
2.7.3 Identified Gaps in Current Ugandan Recommendation Systems .	27
2.8 Synthesis of Literature and Identification of Research Gap	28
3 Methodology	29
3.1 Introduction	29

3.2	Research Questions and Hypotheses	29
3.3	Research Philosophy and Design	31
3.4	Study Population	32
3.5	Variables of the Study	33
3.5.1	Dependent Variable	33
3.5.2	Independent Variables	33
3.6	Sample Size and Sampling Strategy	34
3.6.1	Quantitative Sample: Statistical Justification	34
3.6.2	Quantitative Sample: Machine Learning Dataset	35
3.6.3	Qualitative Sample: Key Informant Interviews	36
3.7	Data Collection	36
3.7.1	Secondary Data: Historical Transaction Records	36
3.7.2	Data Management: Storage, Retrieval, and Version Control	38
3.7.3	Primary Data: Semi-Structured Interviews	39
3.8	Data Analysis Procedures	39
3.8.1	Qualitative Analysis: Thematic Analysis	39
3.8.2	Data Preprocessing and Feature Engineering	40
3.8.3	Model Development and Training	42
3.8.4	Model Evaluation	43
3.8.5	Strategic Recommendation Formulation	44
3.9	Validity and Reliability	44
3.10	Ethical Considerations	45
4	Results	46
4.1	Introduction	46
4.2	Key Customer Attributes and Feature-Target Relationships - Objective 1	46
4.2.1	Spearman Rank Correlation Results	46
4.2.2	One-Way ANOVA Results	48
4.2.3	Chi-Square Test of Independence Results	49
4.2.4	Hypothesis H1 Verdict	49
4.3	Machine Learning Model Performance - Objective 2	50
4.3.1	Hypothesis H2 Verdict	51
4.4	Feature Importance Analysis	51
4.5	Projected Business Impact - Objective 3	53
4.6	Qualitative Findings from Key Informant Interviews	53
4.6.1	Participant Profile	53
4.6.2	Thematic Analysis	54
4.6.2.1	Theme 1: Limitations of ARPU-Based Segmentation	54

4.6.2.2	Theme 2: Data Infrastructure Readiness and Real-Time Feasibility	55
4.6.2.3	Theme 3: Customer Experience Degradation from Irrelevant Offers	55
4.6.2.4	Theme 4: Strategic Appetite and Governance Requirements	55
4.6.3	Contribution of Qualitative Findings to the Overall Study	56
4.7	Summary of Results	56
5	Discussion	57
5.1	Introduction	57
5.2	Discussion of Key Customer Behavioural Attributes: Hypothesis H1	57
5.2.1	The Large-Sample Context	57
5.2.2	Device Type and Smartphone Ownership as the Strongest Signal	58
5.2.3	Bundle Spend Recency and Purchase Frequency	58
5.2.4	Voice Activity Features	58
5.2.5	UPGRADE RATE: The One Non-Significant Feature	59
5.3	Discussion of Model Performance: Hypothesis H2	59
5.3.1	Interpreting the F1-Score in Context	59
5.3.2	Performance by Segment Class	60
5.3.3	The Performance Hierarchy	60
5.3.4	Model Interpretability and Explainable AI	60
5.4	Strategic Implications of Feature Importance Results	61
5.4.1	Revenue Efficiency Over Raw Revenue	61
5.4.2	Call Intensity as a Real-Time Trigger Signal	62
5.4.3	Mobile Money as a Financial Engagement Signal	62
5.5	Projected Business Impact and Evaluation of Hypothesis H3	62
5.6	Integration of Qualitative and Quantitative Findings	63
5.7	Assessment of Research Objectives	64
5.8	Summary	64
6	Conclusions and Recommendations	65
6.1	Introduction	65
6.2	Conclusions	65
6.2.1	Conclusion on Specific Objective 1: Key Customer Attributes	65
6.2.2	Conclusion on Specific Objective 2: Model Selection and Performance	66
6.2.3	Conclusion on Specific Objective 3: Business Impact and Strategic Value	67
6.3	Strategic Recommendations for Airtel Uganda	68

6.3.1 Recommendation 1: Deploy a VOICE_INTENSITY-Triggered	
Personalised Recommendation Engine	68
6.3.2 Recommendation 2: Enrich Segmentation with Device Type and	
Channel Features	69
6.3.3 Recommendation 3: Establish a Formal A/B Test and Model	
Governance Framework	69
6.4 Limitations of the Study	70
6.5 Directions for Future Research	71
6.6 Final Remarks	72
References	73
Appendices	79
A Data Use Agreement — Airtel Uganda Limited	80
B Participant Information Sheet	83
C Semi-Structured Interview Guide	86
D Research Ethics and Data Collection Approval	89
E Research Objectives, Questions, and Hypotheses	91
F GitHub Repository and Analytical Codebase	93
G Declaration of AI Use	94
H Dataset Sample Screenshots	95

List of Figures

1	Conceptual Framework of the study	10
2	Distribution of Average Revenue Per User (ARPU)	16
3	Customer Voice Bundle Purchase Frequency	16
4	Voice Bundle Purchase Distribution by Hour of the Day	16
5	Total Outgoing Voice Minutes from all Subscribers	17
6	Distribution of subscribers across value segments	17
7	Spearman Rank Correlation — Top Features vs. Voice Bundle Value	
	Segment	48
8	Comparative Model Performance — All Models	51
9	Relative Importance of Key Customer Behavioural Features	52

List of Tables

1	Empirical Review of Machine Learning Methodologies in Telecom Recommendation and Segmentation	23
2	Summary of Dataset Components Used in the Study	38
3	GridSearchCV Parameter Grids and Winning Hyperparameters by Model	43
4	Spearman Rank Correlation — Selected Features vs. Voice Bundle Value Segment	47
5	One-Way ANOVA — Top Features by F-Statistic Across Value Segments	48
6	Chi-Square Test of Independence — Categorical Features vs. Value Segment	49
7	Comparative Performance of Machine Learning Models on the Voice Bundle Segment Classification Task	50
8	Top Five Features by Relative Importance — Random Forest Model . .	52
9	Projected Business Impact of ML-Based Personalised Bundle Recommendations	53
10	Profile of Key Informant Interview Participants	54
11	Integration of Qualitative Themes and Quantitative Findings	63

List of Abbreviations

Abbreviation	Full Form
AI	Artificial Intelligence
ANOVA	Analysis of Variance
ARPU	Average Revenue Per User
BI	Business Intelligence
CDR	Call Detail Record
CF	Collaborative Filtering
DRL	Deep Reinforcement Learning
ML	Machine Learning
MNO	Mobile Network Operator
OCS	Online Charging System
RFM	Recency, Frequency, Monetary
SHAP	SHapley Additive exPlanations
SMOTE	Synthetic Minority Over-sampling Technique
SQL	Structured Query Language
UCC	Uganda Communications Commission
USSD	Unstructured Supplementary Service Data
XAI	Explainable Artificial Intelligence
XGBoost	Extreme Gradient Boosting

Acknowledgements

I dedicate this work to my former boss, Kenneth Kyobe, who gave me a chance to be part of the Centenary Bank Business Intelligence and Data Office, for his unwavering support and encouragement throughout my professional journey.

It goes without mentioning that I am forever grateful to my supervisor, Dr. Daphne Nyachaki Bitalo for without her guidance, feedback and encouragement, this research would never have come to full fruition.

Chapter 1

Introduction

1.1 Background to the Study

The global telecommunications industry is undergoing a fundamental transformation driven by the proliferation of mobile devices, the exponential growth of customer data, and rapid advances in artificial intelligence and machine learning. In this environment, the ability to understand and anticipate individual customer needs has emerged as a primary source of competitive differentiation. Globally, telecommunications operators are increasingly moving away from generic, product-centric marketing strategies toward data-driven, customer-centric personalisation — a shift that research consistently links to improvements in customer satisfaction, retention, and revenue (McKinsey & Company, 2022; Sikri et al., 2024).

In Uganda, this transformation is unfolding in a market characterised by intense competition, a predominantly prepaid subscriber base, and high customer sensitivity to price and product relevance. The Ugandan telecommunications sector is regulated by the Uganda Communications Commission (UCC) and currently comprises several competing Mobile Network Operators (MNOs), of which Airtel Uganda Limited is one of the two dominant players. As of late 2024, Airtel Uganda commanded approximately 47 percent of the market share, with an active subscriber base of approximately 10.5 million customers (Uganda Securities Exchange, 2025; Publicist East Africa, 2025) [3]. Despite this substantial market presence, Airtel Uganda, like its competitors, faces persistent challenges in customer retention, with monthly churn rates estimated to fluctuate between 3 and 5 percent across major Ugandan operators.

Airtel Uganda faces monthly customer churn rates estimated between 3 and 5 percent in a highly competitive prepaid market [3].

A central driver of this churn and broader customer dissatisfaction is the mismatch

between the products offered to customers and their actual consumption needs. Airtel Uganda’s voice bundle portfolio comprises over 150 distinct products, varying in price, call minutes, validity period, and target usage occasion. Navigating this portfolio through a generic Unstructured Supplementary Service Data (USSD) menu, without personalised guidance, places a significant cognitive and financial burden on the subscriber. Customers who select inappropriate bundles — either over-purchasing relative to their needs or depleting their allocation prematurely — experience dissatisfaction that erodes loyalty over time. The Average Revenue Per User (ARPU) for Airtel Uganda currently averages approximately UGX 8,500 per month, a figure that the company’s published financial results suggest has room for meaningful growth through more effective customer engagement (Airtel Uganda, 2024).

Against this backdrop, this study proposes the application of machine learning (ML) techniques to Airtel Uganda’s transactional customer data as the basis for a personalised voice bundle recommendation system. Machine learning offers the capacity to identify complex, non-linear patterns in large-scale behavioural data that are invisible to traditional rule-based segmentation approaches (Samuel, 1959; Hastie et al., 2009). By training predictive models on historical purchase behaviour, usage patterns, and temporal dynamics, it is possible to develop a system that recommends the most relevant voice bundle to each subscriber at the most opportune moment, transforming the customer’s experience from one defined by generic choice to one defined by intelligent, personalised guidance.

This study was conducted as a case study of Airtel Uganda Limited, utilising a dataset comprising approximately 150 million voice bundle purchase transactions from 1,458,900 subscribers from February to April 2026. This research contributes to the growing body of global literature on machine learning applications in telecommunications while addressing a specific and documented gap: the absence of published, empirically grounded ML recommendation studies in the Ugandan prepaid voice bundle market.

1.2 Problem Statement

Airtel Uganda, in common with many Mobile Network Operators in developing economies, currently relies on generalised, rule-based marketing strategies for the promotion and sale of voice bundles. These strategies segment customers into broad groups — typically based on a single metric such as Average Revenue Per User — and deliver standardised promotional messages to all members of each segment, without regard for the individual behavioural patterns, consumption rhythms, or immediate contextual needs of each subscriber.

The commercial consequences of this approach are well documented and significant.

Generic, untargeted promotional campaigns in the telecommunications sector achieve conversion rates typically below 3 percent, representing a poor return on marketing investment and a substantial volume of wasted customer communication (McKinsey & Company, 2022). For Airtel Uganda specifically, this translates to an estimated 20 percent wastage of the marketing budget on campaigns that fail to convert, and an estimated 15 percent of the total user base that remains low-spending or inactive — a segment that appropriate, personalised recommendations could potentially re-engage and monetise [31].

Beyond commercial inefficiency, the one-size-fits-all approach generates a poor customer experience. Subscribers who receive irrelevant bundle recommendations — offers that do not reflect their actual usage patterns or financial circumstances — experience a diminished sense of being understood and valued by their service provider. This experience of irrelevance is a documented driver of customer dissatisfaction and churn in prepaid mobile markets (Al-Mashraie et al., 2020; Christopher and Nelson, 2024). The problem is further compounded by the scale and complexity of Airtel Uganda’s voice bundle portfolio: with over 150 distinct products available, the probability that a randomly selected or generically recommended bundle will be the optimal choice for any individual subscriber is extremely low [55].

Therefore, the core problem is the absence of a scalable, data-driven recommendation system capable of providing timely and personalised voice bundle recommendations that align with individual customer consumption patterns and behavioural profiles. This gap between the richness of the behavioural data available within Airtel Uganda’s systems and the simplicity of the segmentation logic currently applied to it represents both a commercial opportunity and an academic research gap that this study seeks to address.

To the researcher’s knowledge, no published study has applied a machine learning classification framework to personalised voice bundle targeting using real, proprietary transactional data from a Ugandan mobile network operator. Prior studies in the broader telecom recommendation literature have relied on publicly available benchmark datasets or synthetically generated data [11, 8, 51]. The absence of empirically validated, deployment-ready recommendation frameworks for the East African prepaid market constitutes the primary academic gap that this study addresses. This study contributes the first such framework, validated on 1,458,900 real Airtel Uganda subscriber records across 738 voice and combo bundle products.

1.3 Justification of the Study

This study is justified by its potential to generate significant and demonstrable value for three distinct stakeholder groups: Airtel Uganda as the case study organisation, the customers, and the academic community. Each justification is grounded in the specific gaps identified in both the current practice and the existing literature.

1.3.1 Justification for the Telecom Industry

From an industry perspective, this research provides Airtel Uganda — and by extension, the broader East African telecommunications sector — with an evidence-based methodology for overcoming the economic inefficiencies inherent in generic, rule-based marketing systems. The financial case for personalisation is well-established globally: McKinsey and Company (2022) demonstrated that telecommunications operators implementing analytics-driven recommendation engines achieve revenue increases of up to 10 percent and improvements in customer satisfaction of 20 to 30 percent. This study translates this global evidence into a contextually grounded, operationally applicable framework for the Ugandan prepaid market.

Specifically, this study delivers a validated machine learning model that Airtel Uganda can deploy to intelligently recommend voice bundles at the individual subscriber level, replacing the current high-volume, low-conversion broadcast campaign approach. The projected commercial impact, including a four- to five-fold improvement in bundle offer conversion rate and a 9 percent improvement in ARPU (simulated pilot projection), represents a compelling return on investment that directly addresses the operator’s stated strategic priorities of revenue growth and customer retention (Airtel Uganda, 2024). Furthermore, the study’s feature importance findings provide Airtel Uganda’s product and marketing teams with actionable intelligence on which customer attributes most strongly drive bundle purchasing behaviour — intelligence that can inform product design, pricing strategy, and campaign timing beyond the immediate recommendation use case.

1.3.2 Justification for the Customer

This research is equally justified by its potential to deliver tangible and direct benefits to Airtel Uganda’s subscribers. The current absence of personalisation places an unreasonable burden on customers, who must independently navigate a portfolio of over 150 voice bundle products without data-driven guidance and who frequently experience the financial loss of purchasing bundles that do not match their actual usage patterns. Wang and Wang (2010) demonstrated that a high degree of service personalisation is directly and significantly linked to increased customer satisfaction and stronger brand

loyalty in mobile services. This was a finding confirmed more recently by Medallia (2024), which reported that 82 percent of consumers indicated that personalised experiences drive their brand choice in at least half of their purchasing situations.

The recommendation system developed in this study addresses customer needs in two concrete ways. First, by precisely matching the recommended bundles to each subscriber’s demonstrated consumption patterns, the system minimises the probability of over-purchasing or under-purchasing, thereby maximising the value the customer receives from each transaction. Second, by delivering recommendations at the moment of genuine purchase intent, as identified through the bundle depletion signal, the system replaces reactive, unsolicited promotional messaging with timely, relevant, and genuinely useful guidance. The result is a customer experience defined by relevance and responsiveness, rather than irrelevance and noise.

1.3.3 Justification for Academia

From an academic standpoint, this study addresses a documented and significant gap in the existing literature. As established in Chapter Two, the global body of knowledge on machine learning applications in telecommunications is substantial and growing, with well-validated frameworks for customer segmentation, churn prediction, and recommendation system design. However, the application of these frameworks to the specific context of the Ugandan prepaid voice bundle market, characterised by high-frequency, low-denomination transactions, a dominantly prepaid subscriber base, and a distinct set of behavioural dynamics, remains largely absent from the published literature. Existing Ugandan telecom literature is predominantly survey-based and qualitative, focusing on perceptions of service quality and loyalty rather than empirical behavioural modelling (Faridah et al., 2023; Christopher & Nelson, 2024).

This study contributes to the existing knowledge by providing the first published, empirically grounded ML recommendation framework built on real transactional data from a Ugandan operator. In doing so, it serves as a replicable methodological template for future researchers studying personalisation and machine learning in other African and developing economy telecom markets. The mixed-methods design, which combines large-scale quantitative modelling with qualitative insights from industry practitioners, further contributes to the methodological literature by demonstrating how technical ML findings can be contextualised and validated through organisational knowledge (Creswell & Creswell, 2018; Saunders et al., 2019).

Beyond the academic contribution, the study is also justified by its reproducibility and adaptability. The complete analytical codebase, from SQL extraction through model evaluation, is publicly available on GitHub, enabling future researchers across Sub-

Saharan Africa to replicate, extend, and adapt the framework to other operators and markets without starting from scratch.

1.4 Main Objective of the Research

The main objective of this study is to identify the key customer behavioural attributes that can be used to predict appropriate voice bundles for individual Airtel Uganda subscribers and to leverage these attributes to develop a machine learning-based recommendation framework that supports Airtel Uganda’s business growth objectives.

This study moves beyond conventional customer segmentation by establishing a data-driven relationship between specific customer behavioural attributes — operationalised as independent variables — and individual voice bundle purchase decisions, operationalised as the dependent variable. This study focuses specifically on identifying which engineered behavioural features, including usage patterns, temporal dynamics, and bundle depletion rates, are the most influential predictors of the bundle type a customer is most likely to purchase. This predictive capability is intended to serve two primary stakeholder groups simultaneously: Airtel Uganda, by providing a data-driven engine for intelligent bundle recommendations that reduce wasted marketing spending and increase conversion rates and revenue; and the customer, by ensuring that recommended bundles are genuinely relevant to their individual consumption profile, thereby reducing the frustration of navigating a generic product menu and the financial cost of purchasing inappropriate products.

1.5 Specific Objectives

The following three specific objectives operationalise the main research objective and structure the analytical work of the study:

1. To identify the key customer behavioural attributes and patterns that are most predictive of voice bundle purchase decisions among Airtel Uganda prepaid subscribers.
1. To evaluate the projected business impact of the best-performing machine learning model on voice bundle offer conversion rate and Average Revenue Per User, and to formulate evidence-based strategic recommendations for Airtel Uganda to improve customer satisfaction and optimise marketing efficiency.
1. To formulate strategic, evidence-based machine learning recommendations for Airtel Uganda to grow its business, optimise marketing efficiency, and improve customer satisfaction.

1.6 Research Questions

The following research questions correspond directly to the three specific objectives and provide the evaluative framework against which the empirical findings of the study are assessed in Chapters Four and Five:

1. What are the key customer behavioural attributes and patterns that most effectively predict voice bundle purchase decisions among Airtel Uganda’s prepaid subscribers? (RQ1)
1. Which machine learning models are most suitable and accurate for predicting and classifying customer voice bundle preferences, and which data preprocessing approaches best support this task? (RQ2)
1. What strategic recommendations can be derived from the machine learning model’s outputs to help Airtel Uganda improve customer satisfaction, reduce churn, and grow its business? (RQ3)

1.7 Research Hypotheses

Based on the research questions, specific objectives, and theoretical framework established in Chapter Two, the following testable hypotheses guide this study. Each hypothesis is presented with its null form (H_0) and alternative form (H_1) and will be evaluated against the empirical findings in Chapters Four and Five:

Hypothesis 1 (corresponding to RQ1) — $H1_0$: No customer behavioural attribute is significantly associated with the type of voice bundle purchased by Airtel Uganda subscribers. $H1_1$: At least one engineered customer behavioural attribute — such as bundle depletion rate, purchase frequency, temporal usage pattern, or rolling ARPU — is a significant predictor of the type of voice bundle purchased.

Hypothesis 2 (corresponding to RQ2) — $H2_0$: No significant difference exists in the predictive accuracy of the candidate machine learning models when applied to the Airtel Uganda voice bundle classification. $H2_1$: At least one machine learning model significantly outperforms the others in predicting the correct voice bundle class for Airtel Uganda customers, as measured by Accuracy, F1-Score, Precision, and Recall.

Hypothesis 3 (corresponding to RQ3) — $H3_0$: The best-performing machine learning model does not generate recommendations that project a statistically meaningful improvement in voice bundle offer conversion rate above the current baseline of approximately 3 percent, as estimated from the model’s weighted precision on the hold-out test set. $H3_1$: The best-performing machine learning model generates recommendations that project a statistically meaningful improvement in voice bundle offer conversion rate

above the current baseline of approximately 3 percent, as estimated from the model’s weighted precision on the hold-out test set and benchmarked against documented industry outcomes (McKinsey & Company, 2022; 2024) for Airtel Uganda.

1.8 Scope of the Study

1.8.1 Geographical Scope

This study is geographically delimited to Airtel Uganda subscribers whose transactional data fall within the Central Business District of Kampala, Uganda. Kampala, as Uganda’s commercial capital, represents the highest concentration of active prepaid subscribers and the greatest diversity of usage behaviour, making it an appropriate and practically accessible primary geographical focus for this study. The qualitative component, key informant interviews, was conducted with Airtel Uganda staff based at the company’s Kampala headquarters.

1.8.2 Time Scope

The quantitative analysis is based on transactional data spanning a three-month period from February to April of 2026. This time window was selected for three reasons: First, it aligns with the Uganda Communications Commission’s (UCC) regulatory requirement that inactive subscribers be disconnected after 90 days of inactivity, ensuring that the dataset reflects only genuinely active subscribers whose behaviour is commercially relevant. Second, a 90-day observation window is sufficient to capture meaningful short-term behavioural patterns, including purchase frequency cycles, depletion rhythms, and temporal usage routines, while remaining computationally manageable for large-scale ML model training. Third, analysing data from the most recently available period ensures that behavioural insights reflect current market conditions, product offerings, and competitive dynamics in the Ugandan telecommunications market.

1.8.3 Content Scope

This study focuses exclusively on the prepaid voice bundle purchase behaviour of Airtel Uganda subscribers. The primary analytical unit is the individual voice bundle purchase transaction, characterised by the bundle type, price, timestamp, and associated customer behavioural attributes. Data bundle and combined voice-and-data bundle purchasing patterns are outside the scope of this study, as are postpaid subscriber behaviour, which follows a structurally different purchasing process. The machine learning models developed were designed specifically for the prepaid voice bundle recommendation use case and were not generalised to other product categories or subscriber

types within this study.

1.9 Significance of the Study

This research carries significance at three levels: the telecommunications industry, end-users, and the academic community, each of which corresponds to a gap in current knowledge or practice that the study directly addresses.

For the telecommunications industry, the study provides an operationally implementable, empirically validated recommendation framework that Airtel Uganda can deploy to improve marketing efficiency, increase ARPU, and reduce churn. The findings demonstrate that the data necessary to build such a system already exist within the operator's own transactional records; the missing element is the analytical framework to unlock its value. Beyond Airtel Uganda, the methodology is transferable to other Ugandan operators and, with appropriate adaptation, to other prepaid-dominant telecom markets across East and sub-Saharan Africa.

For customers, the significance lies in the direct improvement in the quality and relevance of the product recommendations they receive. A recommendation system that correctly identifies the most appropriate bundle for each subscriber, at the moment of purchase intent, translates into reduced financial waste, fewer irrelevant communications, and a more intuitive and satisfying service experience. These improvements directly address the customer-side pain points identified in the problem statement: irrelevant messaging, confusing product navigation, and the financial cost of purchasing mismatched bundles.

For academia, the study's significance lies in its contribution of a replicable, context-specific empirical framework for ML-based personalisation in African prepaid telecom markets, a context that is underrepresented in the global ML and telecommunications literature. The study's mixed-methods design, which combines large-scale quantitative behavioural modelling with qualitative practitioner insights, also contributes to the methodological literature on how technical ML findings can be grounded in organisational and market contexts to produce actionable, rather than merely theoretical, knowledge.

1.10 Theoretical and Conceptual Framework

The conceptual framework of this study illustrates the relationships between the key variables of the research and the mechanism through which the machine learning intervention translates customer behavioural data into personalised recommendations

and measurable results. The framework is grounded in two theoretical traditions established in Chapter Two: the Push-Pull-Mooring (PPM) framework, which explains how personalised recommendations function as a retention mechanism (Al-Mashraie et al., 2020), and the hyper-personalisation paradigm, which positions individual-level behavioral modeling an evolution beyond traditional Segmentation, Targeting, and Positioning (Kotler & Keller, 2016; McKinsey & Company, 2024).




Figure 1: Conceptual Framework of the study

The independent variables in this study are customer behavioural and contextual attributes that serve as inputs to the machine learning model. These include historical usage patterns such as call frequency and duration, purchase behaviour variables including bundle purchase frequency, previously purchased bundle types, and bundle depletion rate, temporal and geographic attributes including time of day, day of week, and primary location of usage, and account-level information including customer tenure and anonymised demographic data where available. These variables are hypothesised — on the basis of the literature reviewed in Chapter Two — to be the primary behavioural drivers of voice bundle purchase decisions.

The mediating variable, and the core technical intervention of the study, is the machine learning model. The model ingests multi-dimensional independent variables and is trained to identify the complex, non-linear relationships between these attributes and the customer’s bundle purchase decision. The candidate models evaluated in this study include Logistic Regression, Decision Tree, Random Forest, Gradient Boosting, and Item-Based Collaborative Filtering, selected on the basis of the global literature reviewed in Chapter Two. The model transforms raw behavioural data into ranked,

personalised bundle recommendations, the mechanism through which independent variables are translated into the dependent variable prediction.

The dependent variable is the voice bundle purchase decision, which is the specific bundle type that a customer selects from the available portfolio. This was operationalised as a multi-class categorical variable, with each class representing a distinct voice bundle product. The machine learning model was designed to predict this variable with sufficient accuracy to surface the customer’s likely choice within a ranked shortlist of five recommendations.

The outcomes component of the framework captures the expected downstream consequences of accurate, personalized recommendations: increased customer satisfaction, driven by the relevance and timeliness of the recommendations received; reduced churn, as customers who consistently receive appropriate bundle suggestions experience a more satisfying and loyalty-building relationship with their operator; and increased revenue and customer lifetime value, as higher conversion rates on recommendations translate directly into improved ARPU and reduced marketing waste. These outcomes collectively address the core problem identified in Section 1.2 and form the evaluative criteria against which the study’s success is assessed in Chapters Four, Five, and Six.

1.11 Limitations of the Study

Several limitations of the design and execution of this study are acknowledged. The first concern is data privacy and proprietary sensitivity. The transactional dataset provided by Airtel Uganda is commercially sensitive and subject to strict data governance protocols. While all data were fully anonymised prior to transfer to the researcher — with subscriber identifiers replaced by non-identifiable user IDs — the proprietary nature of the dataset means that the exact model outputs and feature engineering pipeline cannot be fully disclosed in ways that could expose the company’s competitive intelligence. This constraint limits the complete replicability of the study by external researchers who do not have access to equivalent datasets.

The second limitation relates to the absence of live deployment and a real-time feedback loop. The business impact projections presented in Chapter Four are derived from a simulated pilot on the hold-out test set and have not been validated through a live, randomised A/B test deployment. Consequently, the projected conversion rate improvements and ARPU uplift remain estimates grounded in model performance metrics and external benchmarks, rather than directly observed commercial outcomes. A full-scale A/B test, recommended in Chapter Six as the logical next step, would be required to empirically confirm these projections.

Third, the three-month observation window, while appropriate for capturing short-term behavioural patterns and consistent with regulatory definitions of active subscribers, may not fully capture longer-term seasonal dynamics, such as usage shifts associated with school terms, public holidays, or annual economic cycles, that could influence both customer behaviour and model performance at different points in the year. Future research using longer observation windows will provide a more complete behavioural profile.

Fourth, a potential bias exists in the use of revenue and conversion metrics as primary measures of recommendation system success. Although these metrics are commercially meaningful, they do not fully capture the customer’s subjective experience of receiving and responding to recommendations. These nuances underscore the importance of complementing quantitative model metrics with qualitative customer experience research, a dimension that the present study addresses partially through key informant interviews but not through direct customer surveys.

1.12 Organisation of the Thesis

This thesis is organised into six chapters. Chapter One introduces the research context, problem statement, objectives, research questions, hypotheses, scope, and conceptual framework. Chapter Two presents a systematic review of the relevant literature, progressing from global theoretical frameworks through regional African evidence to the specific Ugandan telecom context, and identifies the research gap that this study addresses. Chapter Three describes the mixed-methods research design, population, sampling strategy, data collection procedures, and analytical pipeline employed in this study. Chapter Four presents the quantitative results, including a comparative model performance evaluation, feature importance analysis, and projected business impact. Chapter Five discusses the findings in the context of the research hypotheses and the literature reviewed in Chapter Two. Chapter Six presents the study’s conclusions, strategic recommendations for Airtel Uganda, acknowledged limitations, and future research directions.

Chapter 2

Literature Review

2.1 Introduction

This chapter reviews the existing literature relevant to the three specific objectives of this study: (1) identifying key customer attributes and patterns relevant to voice bundle recommendations; (2) testing and selecting suitable machine learning models for prediction, classification, and data preprocessing; and (3) developing strategic machine learning-based business recommendations for Airtel Uganda. The review is organised thematically and progresses from a global perspective to a regional African context and finally to the specific Ugandan telecom environment. Sources include peer-reviewed journal articles, conference papers, and industry reports published primarily between 2020 and 2025, with foundational theoretical sources, where necessary.

2.2 Theoretical Framework

2.2.1 Consumer Behaviour Theory and Telecom Service Adoption

Understanding why customers choose particular telecommunications products requires a grounding in consumer behaviour theory. The Theory of Planned Behaviour (Ajzen, 1991), originally applied to individual decision-making, has been widely extended to telecommunications contexts to explain how customer attitudes, social norms, and perceived control shape service selection. In the telecom sector, this translates into understanding that a customer's decision to purchase a specific voice bundle is not random; it is driven by habitual usage patterns, social context, and economic considerations.

More recently, the Push-Pull-Mooring (PPM) framework has gained prominence in

explaining service switching and retention behaviour in telecoms. Al-Mashraie et al. (2020) applied this framework to demonstrate that customers are pushed away from inappropriate service plans, pulled toward better alternatives, and retained or 'moored' by personalised, relevant offers. This study is grounded in the PPM framework insofar as personalised voice bundle recommendations serve as a key mooring mechanism, reducing the likelihood of customers abandoning a service provider in search of better-suited alternatives.

Traditional telecom marketing has relied heavily on the Segmentation, Targeting, and Positioning (STP) model (Kotler & Keller, 2016), which segments customers into broad demographic or behavioural groups and targets them with standardised offers. However, as data volumes and computational capabilities have grown, the field has shifted toward hyper-personalisation, using individual-level behavioural data to deliver recommendations tailored to each customer's unique consumption profile (McKinsey & Company, 2024). This study operates within the hyper-personalisation paradigm, moving beyond group-level segmentation to model individual customer purchasing behaviour.

2.2.2 Machine Learning as a Predictive Decision-Making Framework

Machine learning (ML) provides the technical foundation for this study. Supervised learning, in particular, enables the training of models on historical data to predict future outcomes, in this case, the type of voice bundle a customer is most likely to purchase next. Classification algorithms such as Random Forest, Gradient Boosted Trees (XGBoost), and Logistic Regression have become standard tools for customer behaviour prediction in telecoms (Chang et al., 2024; Sikri et al., 2024).

Complementing supervised methods, unsupervised learning, particularly K-Means clustering, enables the discovery of natural groupings within customer data without predefined labels. These clusters serve as the basis for segment-level recommendation logic. The combination of both approaches reflects a hybrid modelling strategy that has been shown to outperform single-method models in complex real-world telecom datasets (Zhao et al., 2023).

2.3 Key Customer Attributes and Behavioural Patterns in Telecom

2.3.1 Global Perspective: What Drives Customer Bundle Purchasing Decisions?

A substantial body of global literature has examined which customer attributes most significantly predict service purchase behaviour in telecommunications. Mohaimin et al. (2025) analysed customer churn in the US telecom market using billing records, service usage logs, and demographic data, concluding that usage intensity, billing frequency, and service tenure were the strongest predictors of customer behaviour. This reinforces the centrality of behavioural usage data — rather than demographics alone — as the foundation for predictive modelling.

Feature engineering has emerged as a critical differentiator in model performance. Rather than using raw usage figures, researchers have found that derived metrics, such as the rate of bundle depletion, average purchase frequency over rolling time windows, and the ratio of voice usage to bundle capacity, substantially improve prediction accuracy. Singh et al. (2025) demonstrated that machine learning models applied to telecom customer segmentation achieved notably higher accuracy when feature engineering was applied to create derived behavioural variables compared to models using only raw customer attributes. In a similar vein, Sikri et al. (2024) showed that incorporating temporal purchase patterns, rather than static usage snapshots, improved churn prediction accuracy by a meaningful margin.

Temporal and time-of-day features have also proven highly significant. Chen et al. (2023) found that models incorporating hourly and daily consumption rhythms outperformed those using aggregated monthly figures, as customers tend to exhibit consistent purchase timing that reflects real-world behavioural routines. For prepaid voice bundle markets — where purchasing decisions are made multiple times per week — this temporal granularity is particularly valuable. Figure 2 illustrates the distribution of ARPU across a typical prepaid subscriber base, Figure 3 shows daily purchase frequency patterns, and Figure 4 shows hourly purchase peaks — all of which motivate the temporal and financial features used in this study’s modelling pipeline.




Figure 2: Distribution of Average Revenue Per User (ARPU)

Note. Source: Airtel Uganda transactional dataset (February–April 2026).




Figure 3: Customer Voice Bundle Purchase Frequency

Note. Source: Airtel Uganda transactional dataset (February –April 2026).




Figure 4: Voice Bundle Purchase Distribution by Hour of the Day

Note. Source: Airtel Uganda transactional dataset (February–April 2026).




Figure 5: Total Outgoing Voice Minutes from all Subscribers

Note. Source: Airtel Uganda transactional dataset (February–April 2026).




Figure 6: Distribution of subscribers across value segments

Note. Source: Airtel Uganda transactional dataset (February–April 2026).

Geographic and demographic segmentation further enrich predictive models. A 2024 study from the KSEM conference (Mahmoud & Asyhari, 2024) applied ML to Call Detail Records (CDRs) and demonstrated that combining usage behaviour with location-based features and demographic segments yielded the most accurate customer clustering outcomes for telecom marketing applications. The study specifically noted that geographic region was a significant predictor of bundle type preference, as usage patterns vary meaningfully across urban and rural environments.

2.3.2 Regional Perspective: Customer Attributes in African Telecom Markets

African telecom markets present a distinct context that shapes which customer attributes matter most. The dominance of prepaid services, the prevalence of low-denomination, high-frequency micro-transactions, and limited fixed-line infrastructure

mean that the attributes driving service choice differ from those observed in postpaid-dominant markets in Europe or North America.

A Zambia-focused study on recommender systems for telecoms (Chulu, 2017) highlighted that in African markets, Call Detail Records (CDRs) are among the richest available data sources, capturing granular usage behaviour that can serve as a proxy for customer intent. The study emphasized that daily recharge frequency and average bundle size purchased were stronger predictors of appropriate product recommendations than demographic data alone, given the informality of customer registration data in many African markets.

Nafiu et al. (2012) identified that in emerging markets, the most commercially relevant customer attribute is spending regularity — whether a customer purchases consistently or sporadically. Sporadic purchasers are often underserved by generic bundles and represent the greatest opportunity for targeted recommendations. More recent regional studies reinforce this finding: Akanferi et al. (2022) examined mobile commerce adoption in Ghana and found that perceived value — whether a customer feels they are getting an appropriate product for their usage level — is the primary driver of repeat purchase behaviour in African prepaid mobile markets.

2.3.3 Ugandan Context: Customer Attributes and Service Loyalty

In Uganda, the published literature on customer behaviour in telecoms has predominantly focused on service quality and its relationship to loyalty, rather than on data-driven attribute modeling. Christopher and Nelson (2024) examined marketing strategies at MTN Uganda and found that while service quality was the primary driver of customer loyalty, targeted promotional offers — essentially bundle recommendations — were the most significant secondary factor. This finding is directly relevant to the present study: it suggests that even in the Ugandan market, appropriately targeted bundle recommendations have a demonstrable impact on customer retention.

Faridah et al. (2023) applied Multiple Linear Regression to analyze customer loyalty in Ugandan telecoms, identifying network quality, pricing, and customer service as the primary loyalty determinants. However, this study — like most Ugandan telecom research — relied on survey data rather than transactional behavioral data. The result is that the attributes identified reflect customer perceptions rather than actual behavioral patterns. This distinction is important: perception-based research tells us what customers say influences their decisions, but behavioral data tells us what actually drives their purchasing actions (Zhao et al., 2023). This study addresses this gap by analyzing real transactional data from Airtel Uganda.

2.4 Machine Learning Models for Prediction, Classification and Data Preprocessing

2.4.1 Classification Models for Customer Behaviour Prediction: Global Evidence

The global literature presents a rich evidence base for the application of classification algorithms to customer behaviour prediction in telecommunications. Alotaibi and Haq (2024) conducted a comprehensive comparative study using Random Forest, LightGBM (LGBM), XGBoost, Logistic Regression, Decision Trees, and Artificial Neural Networks (ANN) on a telecom customer dataset. Their results showed that ensemble methods — particularly XGBoost and Random Forest — consistently outperformed simpler models, with the Random Forest achieving an accuracy of approximately 91.66% and strong precision and recall scores. Importantly, the study found that these models benefited substantially from hyperparameter tuning and cross-validation, with unoptimized models performing significantly worse.

Chang et al. (2024) reached similar conclusions, demonstrating that Decision Trees, Boosted Trees, and Random Forests offered the best combination of accuracy and interpretability for telecom customer classification tasks. Their study emphasized that model interpretability is not merely a technical consideration but a business necessity: telecom marketing teams need to understand why a model makes a particular recommendation to trust and act on it. This aligns with growing interest in Explainable AI (XAI) within the field.

Sikri et al. (2024) published findings in Scientific Reports demonstrating that SVM, Random Forest, and Logistic Regression substantially improved churn prediction accuracy and could also be adapted for retention strategies including personalized bundle recommendations. Their work further showed that AI-powered models could segment customers by risk level and profitability simultaneously, enabling telecoms to prioritize both retention-focused and revenue-growth-focused recommendations from the same modeling framework.

A comprehensive survey of recommender systems published in 2024 traced the evolution from traditional collaborative filtering and content-based approaches to deep learning, graph-based models, and reinforcement learning (Deldjoo et al., 2024). The survey confirmed that for dense user-item interaction datasets, such as telecom transaction logs, collaborative filtering remains highly effective, while hybrid models that combine collaborative and content-based signals achieve the best overall performance. The global recommendation engine market was valued at USD 3.92 billion in 2023 and is projected to grow at a compound annual growth rate of 36.3% through 2030, reflect-

ing the widespread commercial validation of these approaches (Grand View Research, 2023).

2.4.2 Customer Segmentation Using Clustering Algorithms

Customer segmentation through unsupervised learning, particularly K-Means clustering, is one of the most well-established applications of machine learning in telecommunications. A 2024 review of machine learning approaches to customer segmentation, presented at the International Conference on E-commerce and Artificial Intelligence, evaluated K-Means, Fuzzy C-Means, K-Means++, Random Forest, and Artificial Neural Networks across publications from 2020 to 2024. The review found that K-Means and its variants remain the benchmark for behavioral segmentation in telecom, with the most meaningful improvements coming from better feature engineering rather than algorithmic novelty (Yuming et al., 2024).

Zhao et al. (2023) proposed a hybrid deep learning approach combining unsupervised and supervised classification to discover customer value segments in telecoms. Their model applied K-Means to identify initial behavioral clusters, then trained a deep learning classifier to assign new customers to segments based on their early behavioral signals. The study demonstrated that this two-stage approach achieved significantly higher accuracy than using either method in isolation, with the combined model also proving more robust to data sparsity — a common challenge in prepaid telecom datasets.

Mahmoud and Asyhari (2024) applied ML-based segmentation specifically to CDR data in the context of 5G telecom planning, using usage behaviour and loyalty indicators as segmentation features. Their findings confirmed that segments defined by behavioral loyalty metrics — such as consistent high-spend or frequent low-denomination purchasers — were more actionable for marketing than segments defined by demographics alone.

2.4.3 Data Preprocessing: The Foundation of Model Quality

The quality of machine learning model outputs is fundamentally constrained by the quality of the input data and the rigor of preprocessing. This is particularly relevant in African telecom contexts where data quality challenges, including missing values, inconsistent transaction logging, and class imbalance in the target variable, are common.

Class imbalance is one of the most discussed preprocessing challenges in the literature. In telecom datasets, certain bundle types may be purchased far more frequently than others, causing models to systematically ignore underrepresented classes. Zdziebko

et al. (2024) specifically addressed this challenge in the context of telecom customer modeling, applying fuzzy-based churn modeling with usage data to handle imbalanced datasets and showing that their approach produced more interpretable and operationally useful results than standard oversampling techniques. Sikri et al. (2024) similarly treated class imbalance as a central preprocessing step, testing SMOTE (Synthetic Minority Oversampling Technique) and other resampling methods before model training.

Feature engineering — the creation of new, informative variables from raw data — is consistently identified as one of the highest-impact preprocessing steps. A 2025 study by Mohaimin et al. demonstrated that derived features such as rolling average spend, bundle depletion rate, and purchase time-of-day dramatically improved model accuracy over raw attribute inclusion. Encoding of categorical variables — particularly customer demographic categories and geographic identifiers — using techniques such as Label Encoding and One-Hot Encoding is also identified as a necessary preprocessing step to enable numerical ML algorithms to process categorical data (Alotaibi & Haq, 2024).

Dimensionality reduction through Principal Component Analysis (PCA) has also been applied in telecom customer modeling to address the curse of dimensionality — the degradation of model performance when too many features are included. A study reviewed by Yuming et al. (2024) demonstrated that applying PCA prior to K-Means clustering improved segment coherence, particularly for high-dimensional CDR datasets.

2.4.4 Model Evaluation Metrics

The literature consistently identifies a standard set of evaluation metrics for classification and recommendation models in telecoms. Precision measures the proportion of recommended bundles that were actually relevant to the customer. Recall measures the proportion of all genuinely relevant bundles that the model successfully identified. The F1-Score provides a harmonic mean of precision and recall, useful when there is a trade-off between the two. The Area Under the Receiver Operating Characteristic Curve (AUC-ROC) captures overall discriminative ability across classification thresholds, while the Root Mean Squared Error (RMSE) is applied in regression-based recommendation tasks. Alotaibi and Haq (2024) achieved an ensemble model AUC of 0.79 with a recall of 0.72, which they benchmarked against prior studies. Chang et al. (2024) reported a Random Forest accuracy of 91.66% with an 82.2% precision and 81.8% recall — results this study will use as a baseline for comparison.

2.4.5 Recommender System Paradigms and the Classification-Based Approach

Before reviewing the specific literature on machine learning in telecom bundle recommendation, it is necessary to situate this study within the broader taxonomy of recommender system approaches. Recommender systems are classified into three established paradigms: Collaborative Filtering (CF), Content-Based Filtering, and Hybrid Systems [1, 19].

Collaborative Filtering generates recommendations by identifying patterns in user–item interaction histories—recommending items that users with similar profiles have engaged with [26, 40]. However, CF requires dense interaction matrices and suffers from the cold-start problem. In the prepaid telecom context, the subscriber-to-bundle interaction space is extremely sparse—each subscriber typically purchases only a small fraction of the 738 available bundle products—making CF structurally unsuitable for this setting.

Content-Based Filtering recommends items similar in attributes to those previously consumed. In a telecom context, this would recommend bundles similar in price or duration to prior purchases. While appropriate in some contexts, this approach risks over-specialisation and cannot detect behavioural signals indicating readiness for a higher-tier product.

Hybrid Recommender Systems combine collaborative and content-based signals [19], but require both rich interaction histories and item attribute data—conditions not fully met in the prepaid subscriber context.

Classification-Based Recommendation, the approach adopted in this study, represents a fourth paradigm that has gained traction where rich behavioural feature sets are available but interaction matrices are sparse [53, 44]. In this paradigm, a supervised classifier predicts the value segment that best characterises a subscriber’s behavioural profile, and the segment prediction is used as the targeting rule for bundle recommendations: the classifier predicts that Subscriber A belongs to the Platinum segment, and the recommendation engine surfaces the Platinum tier’s bundle portfolio. The justification for this approach is threefold:

1. Rich behavioural OCS features are available but interaction matrices are too sparse for CF;
2. The classification output is directly consumable by Airtel Uganda’s existing USSD rules engine without new system integration;
3. Classification-based recommendation is validated in the telecom literature [8, 44].

2.5 Empirical Review of Key Methodologies

Table 1 summarizes the key studies reviewed in this chapter, organized by methodology, application context, and principal finding. This table serves as a direct reference point for the model selection and performance comparisons presented in Chapter Four.

Table 1: Empirical Review of Machine Learning Methodologies in Telecom Recommendation and Segmentation

Methodology	Key Study	Application	Key Finding
Collaborative Filtering	Deldjoo et al. (2024)	Telecom Recommender Systems	Hybrid CF models outperform single-approach systems in dense usage datasets.
K-Means Clustering	Yuming et al. (2024)	Customer Segmentation	Feature engineering, not algorithmic novelty, drives the greatest segmentation improvements.
Gradient Boosting (XGBoost)	Alotaibi & Haq (2024)	Customer Classification	Random Forest achieves $\sim 91.66\%$ accuracy; XGBoost matches or exceeds it on imbalanced datasets.
Deep Learning Hybrid	Zhao et al. (2023)	Customer Value Segmentation	Two-stage unsupervised + supervised approach outperforms either method in isolation.
Temporal Feature Modeling	Chen et al. (2023)	Usage Pattern Prediction	Hourly and daily temporal features significantly outperform aggregated monthly figures.
Predictive Analytics	Wassouf et al. (2020)	Customer Retention	Accurate behavioral predictions directly improve the effectiveness of commercial interventions.

Note. Sources cited are representative of the broader literature reviewed. All identified studies rely on publicly available or synthetic datasets; none were conducted using real Ugandan operator data. Source: Author’s synthesis of reviewed literature.

Table 1 reveals several important patterns and gaps across the reviewed literature that directly motivate the present study. First, ensemble methods — particularly Random Forest and XGBoost — are the most consistently high-performing models for telecom classification tasks, with reported F1-Scores ranging from 0.78 to 0.92 across the reviewed studies (Chang et al., 2024; Alotaibi & Haq, 2024; Sikri et al., 2024) [11, 8, 43]. Logistic Regression and single Decision Trees consistently underperform ensemble ap-

proaches. Second, Collaborative Filtering approaches achieve strong recommendation accuracy in data-rich global market contexts, but are predominantly applied where dense interaction histories exist — a condition not met in the Ugandan prepaid setting. Third, and most critically, none of the reviewed studies were conducted using real proprietary transaction data from a Ugandan telecommunications operator — all rely either on publicly available benchmark datasets (notably the IBM Telco Churn dataset) or on synthetically generated samples. This constitutes the primary gap that the present study addresses. Furthermore, none of the reviewed studies simultaneously address voice bundle recommendation, value segment classification, and business impact projection within a single analytical framework — a combination this study contributes to the literature.

2.6 Recommender Systems in Telecommunications

2.6.1 Global Approaches and Evidence

Recommendation engines have become a foundational technology across multiple industries, demonstrating that personalised algorithmic suggestions at scale are both technically feasible and commercially transformative. Global platforms such as Amazon, Netflix, and Spotify have led this transformation. Amazon has publicly attributed approximately 35 percent of its total revenue to its recommendation engine, while Netflix reports that its system drives over 80 percent of content consumed on the platform (Grand View Research, 2023). Spotify similarly uses collaborative filtering and audio-based content modelling to generate personalised playlists and discovery features that are now central to user retention (Deldjoo et al., 2024). The telecommunications sector has, however, lagged behind these digital-native platforms in adopting recommendation engines, with most telecom mobile applications remaining largely static — presenting uniform product menus to all customers rather than personalised interfaces that reflect individual usage patterns. McKinsey and Company (2022) found that telecom operators who implement analytics-driven customer value management can increase revenues by up to 10 percent and improve customer satisfaction and engagement by between 20 and 30 percent, underscoring the significant untapped opportunity in this space.

For telecom-specific applications, collaborative filtering based on usage log data has proven particularly effective. Item-based CF, which identifies bundles frequently purchased together or in sequence by similar customers, can surface contextually appropriate recommendations without requiring explicit customer ratings. A 2024 study on efficient personalised product recommendations proposed a fusion algorithm based on frequent itemset mining that addressed both data sparsity and the cold-start problem,

achieving more accurate recommendations by reducing the candidate item space and computing user-item interest rankings from behavioural signals (Li et al., 2024). The cold-start problem — where new customers have insufficient history for CF — remains a recognised challenge, and the literature recommends using demographic and early behavioural signals to initialise recommendations for new users.

McKinsey and Company (2024) documented a case study of an Asia-Pacific telecommunications company that deployed an AI-powered next-best-experience engine, integrating fragmented customer data into a centralised data lake and developing propensity models to identify customers at risk of dissatisfaction. The system delivered personalised bundle alternatives based on each customer’s usage history, and the results were significant: the approach enhanced customer satisfaction by 15 to 20 percent, increased revenue by 5 to 8 percent, and reduced the cost-to-serve by 20 to 30 percent. This case study establishes a clear business case for the type of data-driven recommendation approach this study seeks to build.

2.6.2 Recommender Systems in Africa

In the African context, the published literature on telecom recommender systems is limited, though growing. The Zambia-focused study by Chulu (2017) remains one of the few dedicated studies of recommender systems for African telecoms and demonstrated that CF-based models built on CDR data could effectively predict relevant service bundles in a prepaid, low-denomination market. The study noted that the high frequency of small transactions in African prepaid markets — in contrast to monthly postpaid billing cycles — actually creates richer behavioural data that can improve recommendation model training, provided the data is adequately preprocessed to remove noise and handle irregular transaction patterns.

More broadly, research on mobile commerce adoption in Ghana by Akanferi et al. (2022) highlighted that recommendation relevance — the degree to which a suggested service matches actual usage behaviour — is a more important driver of adoption than ease of use or interface design in African mobile markets. This finding is significant for the present study: it suggests that technical accuracy is a meaningful commercial objective, not merely an academic one.

2.6.3 Current Recommendation Practices in Uganda

Within Uganda specifically, both Airtel Uganda and MTN Uganda have deployed early-stage recommendation systems. Airtel Uganda’s ‘My Pakalast’ offering segments customers primarily based on Average Revenue Per User (ARPU) over a defined period and allocates bundles accordingly. While this represents a meaningful step beyond

uniform product menus, it is essentially a one-dimensional segmentation approach that does not account for temporal usage dynamics, geographic variation, bundle depletion behaviour, or demographic factors.

MTN Uganda’s ‘MyPakaPaka’ offering, launched subsequently, has been publicly described as leveraging machine learning to analyse subscriber financials and usage behaviour (MTN Uganda, 2021). This signals industry-level recognition in Uganda that data-driven personalisation is a strategic priority. However, the details of the model architecture and feature set are not publicly disclosed, and no published academic study has examined its performance or methodology.

Christopher and Nelson (2024) noted that while both operators have moved toward targeted promotions, the dominant local approach remains ARPU-banding, a relatively blunt instrument compared to the feature-rich behavioural models documented in global literature. This gap between current Ugandan practice and global best practice represents both the research opportunity and the practical contribution of this study.

2.7 Impact of Personalised Recommendations on Revenue and Customer Satisfaction

2.7.1 Global Evidence on Business Outcomes

The business case for personalised recommendations in telecoms is well-established in the global literature. A 2024 State of Customer Experience Personalisation report found that organisations with high personalisation capability were twice as likely to achieve major revenue growth compared to those with low personalisation capability (Medallia, 2024). Additionally, 82 percent of consumers reported that personalised experiences drive their brand choice in at least half of purchasing situations, a finding with direct implications for bundle purchase decisions in the telecom sector.

In the telecom industry specifically, Wassouf et al. (2020) applied predictive analytics to the Syriatel telecom dataset and found that data-driven customer loyalty interventions — including targeted service recommendations, increased customer retention rates significantly. Their study demonstrated a direct link between the accuracy of behavioural predictions and the effectiveness of the subsequent commercial intervention, reinforcing the argument that better ML models translate into better business outcomes.

Research consistently identifies upselling through recommendations as more cost-effective than reactive churn management. This is because churn prevention requires identify-

ing and intervening with at-risk customers after dissatisfaction has already developed, whereas proactive bundle recommendations address the underlying cause of dissatisfaction — the mismatch between available products and customer needs — before it manifests as churn (Nafiu et al., 2012; Sikri et al., 2024).

2.7.2 Revenue Optimisation Through ARPU Enhancement

Average Revenue Per User (ARPU) is the primary financial metric in the telecom industry, and its enhancement through targeted recommendations is a key focus of the literature. Static ARPU banding — as currently practised by Airtel Uganda — captures historical spend but does not distinguish between customers who are at their optimal bundle level and those who are consistently under-purchasing due to a lack of awareness of better-fit products.

A study at the KSEM 2024 conference specifically addressed the relationship between ML-based segmentation and revenue growth in telecoms, arguing that customer segments defined by behavioural loyalty and value metrics enable more precise upselling and cross-selling strategies. The authors noted that customers in the high-frequency, low-denomination purchase segment, analogous to the prepaid Ugandan market, represent significant revenue upside if appropriately matched to higher-value bundle tiers that still fit their consumption patterns (Mahmoud & Asyhari, 2024).

2.7.3 Identified Gaps in Current Ugandan Recommendation Systems

While both Airtel Uganda and MTN Uganda have taken meaningful steps toward personalisation, this literature review identifies four critical gaps in their current approaches that the present study seeks to address. First, existing ARPU-based models overlook the temporal context of consumption, such as festive periods, student holidays, or religious observances, which can significantly alter usage patterns and bundle preferences (Al-Nashif et al., 2020). Second, current systems do not explicitly integrate demographic information, such as age bracket, into their personalisation algorithms, despite evidence that demographic data substantially improves recommender system precision (Hossain et al., 2024; El Hamri & Idri, 2024). Third, geographic or location-based preferences are largely overlooked, despite growing evidence that geographical context influences bundle preference in African markets (Zheng et al., 2020; Al-Sayed et al., 2019). Fourth and most critically, current systems rely on average usage metrics rather than granular bundle depletion analysis, meaning they cannot detect the real-time signal of a customer approaching the end of their current bundle, which this study identifies as the single most powerful predictor of the next purchase decision.

2.8 Synthesis of Literature and Identification of Research Gap

The literature reviewed in this chapter establishes a clear and consistent global evidence base for the use of machine learning in predicting customer service preferences, segmenting customers by behavioural attributes, and deploying recommendation engines that improve both customer satisfaction and business revenue. Across global, regional, and local contexts, three conclusions emerge consistently.

First, behavioural usage attributes, including purchase frequency, bundle depletion rate, temporal purchasing patterns, and derived financial metrics such as rolling ARPU, consistently outperform demographic-only features in predicting customer service choices. Second, ensemble classification models such as Random Forest and XGBoost are the current best-practice tools for customer behaviour classification in telecoms, achieving accuracy levels above 80 percent in well-preprocessed datasets. Third, hybrid recommender systems that combine collaborative filtering with content-based signals represent the global frontier of recommendation technology in this sector.

However, the application of these global frameworks to the specific context of the Ugandan prepaid voice bundle market remains an identified gap. Existing Ugandan research is predominantly survey-based and qualitative, focusing on perceptions of service quality rather than behavioural modelling. The current recommendation approaches deployed by Ugandan operators are ARPU-banded rather than multi-feature and behavioural, and no published academic study has evaluated a feature-rich ML classification model for voice bundle prediction using Ugandan transactional data. This study addresses this gap directly by applying supervised classification and unsupervised clustering to real transactional data from Airtel Uganda, incorporating features such as temporal purchase behaviour, bundle depletion rates, purchase frequency, and ARPU metrics. The findings will serve as a direct comparator to global benchmark studies such as Chang et al. (2024) and Alotaibi and Haq (2024), and will provide actionable strategic recommendations for Airtel Uganda’s commercial teams.

Chapter 3

Methodology

3.1 Introduction

This chapter outlines the research design, data collection approach, analytical procedures, and ethical considerations that guided this study. The overarching aim is to describe how the three specific objectives of the study were operationalised: identifying key customer behavioural attributes that drive voice bundle purchase decisions; testing and selecting appropriate machine learning models for prediction and classification; and formulating strategic data-driven recommendations for Airtel Uganda. Each methodological decision made in this chapter is directly tied to one or more of these objectives, and the research questions and hypotheses introduced in this chapter will serve as the evaluative framework against which results are discussed in Chapters Four and Five.

This study employs a classification-based recommendation paradigm, in which a supervised XGBoost classifier predicts the value segment that best characterises a subscriber’s behavioural profile, and the segment prediction is used as the targeting rule for bundle recommendations. This is a recognised approach in the recommendation systems literature (Zhao et al., 2023; Singh et al., 2025) and is distinct from, though complementary to, collaborative and content-based filtering approaches discussed in Section 2.4.5. The title of this thesis reflects the end-to-end system — from data extraction through classification to personalised recommendation delivery — rather than the classification algorithm alone.

3.2 Research Questions and Hypotheses

To structure the analytical work of this study and provide a clear basis for comparing findings against prior literature, three research questions and their corresponding hy-

potheses were formulated. These questions operationalise the specific objectives of the study and will be revisited in the results and discussion chapters to assess the degree to which each is answered by the empirical findings.

Research Question 1 (RQ1): Which customer behavioural attributes are the most significant predictors of voice bundle purchase decisions among Airtel Uganda subscribers? This question corresponds directly to Specific Objective 1 and is grounded in the literature showing that derived behavioural features, such as purchase frequency, bundle depletion rate, temporal purchase patterns, and Average Revenue Per User (ARPU) metrics, consistently outperform demographic-only features in predicting customer service choices (Mohaimin et al., 2025; Sikri et al., 2024).

H1₀: No customer behavioural attributes are significantly associated with the type of voice bundle purchased.

H1₁: At least one customer behavioural attribute (e.g., purchase frequency, bundle depletion rate, temporal usage pattern, ARPU) is a significant predictor of the type of voice bundle purchased.

Research Question 2 (RQ2): Which machine learning classification model most accurately predicts the appropriate voice bundle for a given Airtel Uganda customer based on their behavioural profile? This question corresponds to Specific Objective 2 and is informed by global literature comparing the performance of Random Forest, XG-Boost, and Logistic Regression models on telecom customer classification tasks, with benchmark accuracy figures of approximately 91% and F1-scores above 80% reported in comparable studies (Chang et al., 2024; Alotaibi & Haq, 2024).

H2₀: No significant difference exists in the predictive accuracy of the candidate machine learning models when applied to Airtel Uganda voice bundle classification.

H2₁: At least one machine learning model significantly outperforms the others in predicting the correct voice bundle class for Airtel Uganda customers, as measured by Accuracy, F1-Score, and AUC-ROC.

Research Question 3 (RQ3): How can the outputs of a data-driven machine learning recommendation model inform strategic business decisions for Airtel Uganda to improve customer satisfaction and grow revenue? This question corresponds to Specific Objective 3 and is supported by evidence showing that personalised telecom recommendations can increase revenues by up to 10 percent and raise customer satisfaction and engagement by between 20 and 30 percent (McKinsey & Company, 2022).

H3₀: ML-based personalised bundle recommendations yield no strategically significant insights for improving Airtel Uganda’s customer engagement or revenue performance.

H3₁: ML-based personalised bundle recommendations generate actionable strategic insights that can demonstrably contribute to improvements in customer satisfaction and revenue optimisation for Airtel Uganda.

3.3 Research Philosophy and Design

The study is grounded in a post-positivist philosophical position. Post-positivism acknowledges that while an objective reality exists — in this case, actual customer purchasing behaviour — our understanding of it is always partial and shaped by the tools and frameworks we use to observe it (Creswell & Creswell, 2018). This is appropriate for a study that uses large-scale quantitative modelling as its primary tool, but supplements it with qualitative insights to contextualise and validate the quantitative findings.

The methodological design is a mixed-methods approach, combining quantitative analysis of historical transactional data with qualitative data gathered through semi-structured interviews. This integration is justified by the nature of the research problem: while the machine learning models are built and evaluated quantitatively, understanding the business context, existing recommendation practices, and strategic pain points at Airtel Uganda requires qualitative engagement with practitioners. Saunders et al. (2019) describe this as a pragmatic research approach, where the research question determines the method rather than adherence to a single methodological tradition.

The research strategy follows a case study design, with Airtel Uganda Limited as the unit of analysis. This is appropriate because the study seeks to develop and evaluate a recommendation model within a specific, real-world organisational and market context. The temporal dimension is cross-sectional, using a 90-day window of historical customer transaction data as the basis for model training and evaluation.

The research methodology follows a coherent two-phase process. Phase 1 (Quantitative) involves the extraction, preprocessing, feature engineering, statistical testing, and machine learning model training and evaluation on 1,458,900 real Airtel Uganda subscriber records (February–April 2026). Phase 2 (Qualitative) involves semi-structured interviews with eight key informants from Airtel Uganda to validate deployment feasibility and gather practitioner insight on organisational conditions for implementation. The two phases are integrated in Phase 3 (Integration) following the convergent mixed-methods framework of Creswell and Creswell (2018) [17]: quantitative findings identify the predictive attributes and best-performing model; qualitative findings confirm whether those attributes and model outputs are operationally deployable within Airtel Uganda’s existing infrastructure. The integration is presented explicitly in Section 4.6 (Qualitative Findings) and Section 5.6 (Integration of Qualitative and Quantitative

Findings).

A convergent mixed-methods design was adopted [17] for two reasons. First, quantitative analysis of transaction records can identify which behavioural attributes statistically discriminate between customer segments, but cannot confirm whether those attributes are operationally accessible within Airtel Uganda’s existing technical infrastructure — a question requiring insider practitioner knowledge. Second, the machine learning model identifies VOICE_INTENSITY and SPEND_PER_MINUTE as the most actionable engineered features, but whether these signals can be computed in real time from the OCS system and used to trigger a USSD push is an implementation question that only domain experts within Airtel Uganda can answer. The qualitative strand was designed to answer precisely these deployment feasibility questions.

The qualitative findings informed the quantitative modelling process in the following specific way: informants confirmed that both VOICE_INTENSITY and SPEND_PER_MINUTE are derivable from existing OCS transaction records in near real time (Theme 2), which elevated these engineered features from analytical observations to genuinely deployable trigger signals — directly shaping Recommendation 1 in Chapter 6.

3.4 Study Population

The target population of this study comprises all active prepaid subscribers of Airtel Uganda Limited who purchased at least one voice bundle within a defined 90-day observation period. The focus on prepaid subscribers is deliberate: as established in the literature review, the Ugandan telecommunications market is characterised by the overwhelming dominance of prepaid services, where purchase decisions are made frequently, independently, and at low denominations, a behavioural pattern distinct from postpaid contractual environments (Chulu, 2017; Nafiu et al., 2012). Including postpaid subscribers would introduce a structurally different purchasing process into the dataset and potentially distort the recommendation logic.

Airtel Uganda’s published financial results indicate a subscriber base of approximately 10.5 million active customers as of December 2024 (Uganda Securities Exchange, 2025). Of this base, voice bundle purchasers constitute a large and commercially significant segment, as voice bundles remain one of Airtel Uganda’s primary revenue-generating products. The quantitative component of this study utilised transactional data encompassing approximately 150 million transactions from approximately 5.4 million unique prepaid users over the 90-day observation window. This figure constitutes the accessible population — those subscribers for whom complete, anonymised transaction records were available within the data governance framework established with Airtel

Uganda.

The qualitative component of the study targeted a purposively selected population of professionals within Airtel Uganda’s commercial and data analytics functions. These individuals were selected on the basis of their direct involvement in customer segmentation, product design, and recommendation strategy, making them uniquely positioned to provide contextual insights that quantitative data alone cannot capture. Given the specialised nature of this expertise, the total eligible population for qualitative participation was small, estimated at fewer than 20 individuals across the relevant departments.

3.5 Variables of the Study

3.5.1 Dependent Variable

The dependent variable of this study is the voice bundle type purchased by a customer, operationalised as a multi-class categorical variable. Each category represents a distinct voice bundle product offered by Airtel Uganda, such as the Pakalast and Tooti bundle families, which vary by price, duration, call minutes, and validity period. This variable serves as the target label in the supervised classification models: the machine learning algorithms are trained to predict which bundle class a customer is most likely to purchase on a given occasion, based on their behavioural profile. The choice of a categorical dependent variable frames this as a multiclass classification problem, which is why classification algorithms, rather than regression models, form the core analytical tool of this study.

3.5.2 Independent Variables

The independent variables are the customer behavioural and contextual attributes that the machine learning models use as input features to predict the dependent variable. These variables were selected on the basis of the literature reviewed in Chapter Two, which identified behavioural usage features as the strongest predictors of customer service choice in prepaid telecom markets (Mohaimin et al., 2025; Sikri et al., 2024; Chang et al., 2024). The key independent variables operationalised in this study are as follows.

Purchase frequency captures how often a customer purchases a voice bundle within a rolling time window; daily, weekly, and monthly averages are computed. This variable reflects customer engagement intensity and is one of the most predictive behavioural signals in prepaid telecom classification tasks (Zhao et al., 2023). Bundle depletion rate measures the speed at which a customer consumes their purchased bundle, expressed

as a proportion of total bundle value used within the bundle’s validity period. A high depletion rate signals that the customer consistently uses more than their current bundle provides, suggesting a potential match to a higher-tier product.

Temporal purchase features capture the time-of-day, day-of-week, and day-of-month patterns associated with a customer’s purchases. As documented by Chen et al. (2023) and reinforced in the literature review, temporal patterns contain meaningful behavioural information in prepaid markets, where purchasing decisions follow consistent daily and weekly rhythms. Average Revenue Per User (ARPU), calculated over rolling 7-day, 30-day, and 90-day windows, captures the customer’s spending level and trajectory, a key metric established in both global and regional telecom literature as a strong segmentation and prediction feature (Mahmoud & Asyhari, 2024; Zhao et al., 2023).

Recency score quantifies the number of days elapsed since the customer’s most recent bundle purchase, forming the R component of an adapted RFM (Recency, Frequency, Monetary) framework commonly applied in customer behaviour modelling. Customer tenure records the duration of the subscriber’s relationship with Airtel Uganda in days, which has been identified as a predictor of behavioural stability and bundle preference consistency (Mohaimin et al., 2025). Finally, geographic region — derived from the customer’s registered home location or most frequent network site — is included as a categorical feature following literature evidence that geographic context influences bundle preference in African markets (Mahmoud & Asyhari, 2024).

3.6 Sample Size and Sampling Strategy

3.6.1 Quantitative Sample: Statistical Justification

For the purposes of establishing the academic generalizability of the study’s findings to the broader Airtel Uganda prepaid subscriber population, a minimum statistical sample size was calculated using the standard formula for estimating a proportion in a large or unknown population. This formula, widely applied in social and business research when the total population is large enough to be treated as infinite, is derived from the normal approximation to the binomial distribution and is specified in Cochran (1977), as cited in Saunders et al. (2019):

The sample size was calculated using the standard formula for estimating a proportion in a large or unknown population (Cochran, 1977, as cited in Saunders et al., 2019):

$$n = \frac{Z^2 \times p \times (1 - p)}{e^2} \quad (3.1)$$

where n represents the minimum required sample size; Z is the Z-score corresponding to the desired confidence level, which is 1.96 for a 95 percent confidence level — the standard threshold in academic and business research; p is the estimated proportion of the population exhibiting the characteristic of interest, set conservatively at 0.5 to maximize the required sample size and thereby produce the most defensible minimum estimate; and e is the acceptable margin of error, set at 0.05 (5 percent), which is the conventional standard in business research (Saunders et al., 2019; Creswell & Creswell, 2018). Applying these values:

$$n = \frac{(1.96^2 \times 0.5 \times 0.5)}{0.05^2} = \frac{(3.8416 \times 0.25)}{0.0025} = \frac{0.9604}{0.0025} = 384.16 \approx 385 \quad (3.2)$$

This yields a minimum required sample of 385 customers, which is statistically sufficient to draw generalisable conclusions about the Airtel Uganda prepaid subscriber population at a 95 per cent confidence level with a 5 percent margin of error. The study’s quantitative dataset substantially exceeds this minimum, which strengthens the statistical validity of the findings.

3.6.2 Quantitative Sample: Machine Learning Dataset

While the statistical minimum sample of 385 establishes generalizability, the machine learning models require a substantially larger training dataset to learn the complex, high-dimensional behavioural patterns that distinguish between bundle classes. This distinction — between the sample size needed for statistical inference and the data volume needed for ML model performance — is well-established in the literature (Hastie et al., 2009). The quantitative dataset used for model training comprises approximately 150 million voice bundle purchase transactions from approximately 5.4 million unique prepaid subscribers over a 90-day period. This volume of data is consistent with the scale required to train robust multiclass classification models and to capture rare but commercially significant patterns, such as high-frequency low-value purchases or irregular depletion patterns.

A stratified random sampling approach was applied when drawing the analytical sample from the full dataset, ensuring proportional representation of subscribers across bundle type, geographic region, spending tier, and purchase frequency segment. This approach guards against systematic bias in the training data and improves the model’s ability to generalise to underrepresented customer groups.

3.6.3 Qualitative Sample: Key Informant Interviews

For the qualitative component of the study, purposive sampling was employed to identify and recruit key informants with direct knowledge of Airtel Uganda’s recommendation strategies, customer segmentation practices, and commercial objectives. Purposive sampling is appropriate here because the research requires participants with specific, expert knowledge rather than statistical representativeness (Saunders et al., 2019; Turner III, 2010). A total of eight key informants were recruited from Airtel Uganda’s commercial, data analytics, and product management teams. This number was determined by the principle of data saturation — the point at which additional interviews cease to yield substantively new insights (Guest et al., 2006). Participants were informed of the study’s purpose, their right to withdraw at any time, and the confidentiality of their responses prior to participation.

Purposive sampling was employed for the qualitative strand, targeting individuals with direct role-specific knowledge of Airtel Uganda’s commercial strategy, data infrastructure, customer experience management, or revenue operations. The target population was defined as senior practitioners holding managerial or analytical roles within one of four functional areas — marketing and campaigns, data analytics and BI, IT infrastructure and OCS systems, or commercial strategy — with a minimum of two years of tenure in that role at Airtel Uganda. This criterion ensured that informants possessed both operational knowledge and institutional context necessary to address the study’s deployment feasibility questions.

Eight participants were recruited on the basis of informational saturation rather than statistical power — the established standard for sample size determination in qualitative research [23, 57]. Preliminary thematic analysis of the first six interviews indicated that no substantively new themes were emerging across participants. Two additional interviews were conducted to confirm saturation, following the recommendation of Guest et al. (2006) [23] that saturation typically occurs between six and twelve informants for expert-domain studies. Eight informants covering four distinct functional areas is consistent with established norms for expert informant studies in information systems and business research [58, 41].

3.7 Data Collection

3.7.1 Secondary Data: Historical Transaction Records

The primary quantitative dataset was extracted from Airtel Uganda Limited’s production database by the researcher, who holds authorised internal access as an employee of the organisation. The extraction was conducted between February and May 2026

using structured Presto/Hive SQL queries executed against four internal data sources: the OCS transaction table (ocsmommaininternal), the subscriber dimension table (subscriberdimension), the daily voice usage table (br_daily_voice_usage), and the geographic master table (geo_master). The full extraction scripts are available in the study’s GitHub repository (see Appendix VI).

The dataset covers the period February 1 to April 30, 2026 — a 90-day observation window — and comprises records for 1,458,900 unique active prepaid subscribers from the Kampala, Central 1, and Central 2 regions of Uganda, representing the central and most commercially active geographic cluster of Airtel Uganda’s network. A cluster-based sampling strategy was applied: rather than random sampling, all subscribers whose mobile subscriber identity number (MSISDN) prefix falls within specific network clusters (prefixes 25673 and 25620) were selected. This approach preserves the complete 90-day transaction history for every selected subscriber — which is essential for computing time-window behavioral features such as PURCHASES_LAST_30D and BUNDLE_SPEND_30D — while reducing the extraction volume to a computationally manageable size from the raw tables, which contain in excess of 85 million transaction records.

Five separate data exports were produced and subsequently joined in Python to form the final analytical dataset: (1) TARGET_VARIABLE.csv — the product catalogue of 1,905 bundle products, filtered to 738 voice and combo bundles; (2) PURCHASE_TIMING_FEB.xlsx, PURCHASE_TIMING_MARCH.xlsx, and PURCHASE_TIMING_APRIL.xlsx — monthly purchase transaction records containing purchase date, hour, bundle identifier, transaction channel, payment channel, and bundle revenue per transaction; (3) BUNDLE_CONSUMPTION.xlsx and BUNDLE_CONSUMPTION_2.xlsx — daily voice consumption records containing outgoing and incoming minutes by call type (on-net, off-net, international) per subscriber per day; (4) CUSTOMER_PROFILE.xlsx — subscriber demographic and segmentation data including device type, handset generation, region, gender, age, tenure, value segment, total revenue, and call counts; and (5) LOCATION_DETAILS.xlsx — network site master data containing site codes, geographic coordinates, district, and cluster classifications. Sample screenshots of each dataset are provided in Appendix VII. All personally identifiable information — including mobile subscriber numbers — was anonymised prior to extraction, in compliance with Airtel Uganda’s data governance policies and Uganda’s Data Protection and Privacy Act (2019).

Table 2 presents a summary of the dataset components used in this study.

Table 2: Summary of Dataset Components Used in the Study

Dataset	Source Table(s)	Records	Key Variables
TARGET_VARIABLE.csv	leap_product_ catalog_master	1,905 (738 voice)	BUNDLE_ID, BUNDLE_TYPE, BUN- DLE_CATEGORY, PRICE, VALIDITY_DAYS
PURCHASE_TIMING (Feb–Apr)	ocsmonmaininternal, ocsloanmaininternal	~59M rows	PURCHASE_DATE, PURCHASE_HOUR, BUNDLE_ID, TRANSAC- TION_CHANNEL, BUNDLE_REVENUE
BUNDLE_CONSUMPTION	br_daily_voice_usage	~500K rows	EVENTDATE, OG/IC MINUTES by type (on-net, off-net, international)
CUSTOMER_PROFILE	subscriberdimension, crystal, subscriber_lifecycle	1,514,566 rows	PHONE_NUMBER, DEVICE_TYPE, HANDSET_TYPE, AGE, TENURE_DAYS, VALUESEGMENT, TOTAL_REVENUE, CALLS
LOCATION_DETAILS	geo_master	3,697 sites	SITE_CODE, TOWN_NAME, REGION, LATITUDE, LONGITUDE, CLUSTER

Note. All datasets were extracted from Airtel Uganda’s production database under a formal Data Use Agreement. Raw subscriber data is not publicly available. Anonymised sample screenshots are provided in Appendix [H](#)

3.7.2 Data Management: Storage, Retrieval, and Version Control

Upon receipt, the anonymised dataset was stored securely in Google Drive, accessible only to the researcher, with two-factor authentication enabled on the associated Google account. Raw data files were retained in their original CSV format and never modified directly — all preprocessing and transformation operations were performed programmatically on copies loaded into the analytical environment, preserving the integrity of the original files. The full analytical codebase — including data preprocessing

scripts, feature engineering pipelines, model training notebooks, and evaluation outputs — was version-controlled and maintained in a dedicated GitHub repository accessible at <https://github.com/RemmyBisimbeko/Data-Science/tree/main/SEM%204/Thesis>. This repository serves as the primary audit trail for all analytical decisions made in the study and enables reproducibility of the results. All model outputs, evaluation metrics, and intermediate data artefacts were stored within the Google Colab environment during active analysis sessions and saved to the linked Google Drive upon completion of each analytical stage. No raw customer data was shared with third parties beyond the scope defined in the DUA, and no data was retained beyond the period specified in the agreement.

3.7.3 Primary Data: Semi-Structured Interviews

Primary qualitative data was collected through semi-structured interviews with eight key informants from Airtel Uganda. Semi-structured interviews were chosen because they allow for sufficient flexibility to explore unexpected themes while maintaining a consistent question framework across participants, enabling thematic comparison (Braun & Clarke, 2006; Turner III, 2010). Prior to each interview, participants were provided with a plain-language participant information sheet (see Appendix II) describing the study’s purpose, the voluntary nature of their participation, their right to withdraw at any point without consequence, and the confidentiality measures in place. Written informed consent was obtained from all participants before the interview commenced. The interview guide, provided in Appendix III, was organised around three thematic areas corresponding to the study’s objectives: the current state of customer segmentation and recommendation practices at Airtel Uganda; perceived gaps and pain points in existing approaches; and commercial priorities and strategic context for a data-driven recommendation system. Interviews were conducted in person and via video call, with a mean duration of approximately 45 minutes. With participant consent, all interviews were audio-recorded and subsequently transcribed verbatim for analysis.

3.8 Data Analysis Procedures

3.8.1 Qualitative Analysis: Thematic Analysis

Interview audio recordings were stored on a password-protected device and uploaded to a private, encrypted Google Drive folder immediately following each session. Transcripts were produced using a combination of automated transcription and manual review to ensure accuracy, and stored in the same secure folder alongside the original recordings. Qualitative data from the key informant interviews were analysed using

thematic analysis, following the six-phase framework described by Braun and Clarke (2006): familiarisation with the data, generating initial codes, searching for themes, reviewing themes, defining and naming themes, and producing the report. This approach was chosen because it is flexible, rigorous, and appropriate for identifying patterns across a relatively small number of expert interviews. The themes emerging from the qualitative analysis served two purposes: they provided contextual grounding for the interpretation of the quantitative results, and they directly informed the strategic recommendations developed in response to Research Question 3 and Hypothesis H3.

3.8.2 Data Preprocessing and Feature Engineering

All quantitative data preprocessing and analysis were conducted using Python 3.10 within a Google Colab environment, which provided the necessary computational resources for processing the large-scale dataset without local hardware constraints. The full preprocessing pipeline was implemented in Jupyter Notebooks and version-controlled in the project GitHub repository (<https://github.com/RemmyBisimbeko/Data-Science/tree/main/SEM%204/Thesis>). The principal Python libraries employed were: pandas (version 1.5) for data ingestion, manipulation, and cleaning; NumPy for numerical operations and array handling; scikit-learn (version 1.2) for preprocessing transformations, model training, and evaluation; imbalanced-learn for SMOTE oversampling; and XGBoost (version 1.7) for the Gradient Boosting classifier implementation.

Prior to model training, three statistical tests were applied to the merged dataset to evaluate the relationships between candidate features and the target variable (value segment), and to provide empirical justification for feature selection. Spearman Rank Correlation was selected to measure the strength and direction of monotonic relationships between numeric features and the ordinal target variable — it was chosen over Pearson correlation because it does not assume a normal distribution of variables, which cannot be guaranteed for usage-based behavioral metrics in a prepaid subscriber population (Field, 2018). One-Way Analysis of Variance (ANOVA) was applied to determine whether the mean value of each numeric feature differs significantly across the eight value segment groups — a significant F-statistic indicates that a feature discriminates between segment classes and therefore carries predictive value. Chi-Square test of independence was applied to categorical features to determine whether their distribution is independent of the value segment — a significant Chi-Square statistic confirms an association between the categorical attribute and the target. Statistical significance was set at $\alpha = 0.05$ for all three tests. Correlation strength was interpreted using Cohen’s (1988) classification: $|r| < 0.10$ negligible, 0.10–0.29 weak, 0.30–0.49 moderate, ≥ 0.50 strong.

The preprocessing pipeline proceeded in four stages. In the first stage, data cleaning was performed using pandas. The `DataFrame.isnull().sum()` function was used to identify missing values across all columns; continuous features with missing values were imputed using `SimpleImputer(strategy='median')` from scikit-learn, while categorical features were imputed using `SimpleImputer(strategy='most_frequent')`. Duplicate transaction records were identified and removed using `DataFrame.drop_duplicates()`, and timestamp fields were standardised to a uniform UTC datetime format using `pandas.to_datetime()` (Alotaibi & Haq, 2024).

In the second stage, feature engineering was conducted to create the behavioural variables described in Section 3.5.2. Rolling purchase frequency averages over 7-day, 30-day, and 90-day windows were computed using `DataFrame.rolling()` with `groupby()` applied per subscriber ID. The bundle depletion rate was calculated as the ratio of minutes consumed to total bundle minutes, divided by the number of days since purchase. The recency score was derived as the number of days elapsed since the subscriber’s last transaction using datetime differencing. ARPU over multiple time windows was computed by aggregating transaction amounts per subscriber using `groupby().sum()` divided by the observation window length. Temporal features — hour of purchase, day of week, and day of month — were extracted from transaction timestamps using the pandas `DatetimeIndex` accessor attributes `.hour`, `.dayofweek`, and `.day` respectively (Mohaimin et al., 2025).

In the third stage, categorical variables — including geographic region and bundle type labels — were encoded into numerical representations using `LabelEncoder()` from scikit-learn’s preprocessing module. Bundle type, as the target variable, was encoded into integer class labels and stored as the dependent variable vector y , while all feature columns were assembled into the feature matrix X . In the fourth stage, class imbalance in the target variable was addressed using `SMOTE()` from the `imbalanced-learn` library, applied exclusively to the training set after the train-test split to prevent data leakage (Sikri et al., 2024; Zdziebko et al., 2024).

To prevent data leakage — the inadvertent incorporation of test-set information into model training — a strict sequential preprocessing protocol was applied. The 70/30 stratified split was performed first, before any oversampling or scaling operations. SMOTE was applied exclusively to the training set ($n = 1,042,063$ records), generating synthetic minority-class samples only from training data. The hold-out test set ($n = 446,599$ records) contains exclusively real, unaugmented subscriber records and was never exposed to the SMOTE procedure. Median imputation parameters were fitted on the training set and applied to both sets using those training-set values. All reported evaluation metrics therefore reflect the model’s performance on genuinely unseen, real subscriber data.

The Diamond class ($n = 34$ real records in the training split) presented a specific challenge for SMOTE. With only 34 instances, synthetic neighbours are interpolations of an extremely small and potentially non-representative group. SMOTE was therefore applied to the Diamond class with $k = 1$ nearest neighbour and synthetic generation capped at 3 times the original count (102 synthetic samples), rather than the 50,000-sample ceiling applied to other minority classes. This conservative cap limits extrapolation from the small base. The residual limitation — that Diamond-class predictions should be interpreted with caution — is acknowledged in Section 6.4.

3.8.3 Model Development and Training

Following preprocessing, the feature matrix X and target vector y were split into training and test sets using `train_test_split()` from scikit-learn, with a 70/30 ratio and a fixed `random_state` of 42 to ensure reproducibility. Stratified splitting (`stratify=y`) was applied to maintain the class distribution of bundle types across both sets. Four models were trained and compared to address Research Question 2 and test Hypothesis H2.

Logistic Regression was implemented using `LogisticRegression(max_iter=1000, multi_class='multinomial', solver='lbfgs')` from scikit-learn. It was included as a linear baseline to establish a minimum performance benchmark and to assess whether the non-linear complexity of ensemble methods is justified (Chang et al., 2024). A Decision Tree was implemented using `DecisionTreeClassifier(criterion='gini', max_depth=10)` and included as a second interpretable baseline. Decision trees produce human-readable classification rules that can be visualised using scikit-learn’s `export_graphviz()` function, making them useful for communicating model logic to non-technical stakeholders.

Random Forest was implemented using `RandomForestClassifier(n_estimators=100, max_depth=None, random_state=42)` from scikit-learn. It was selected as the primary candidate model on the basis of its strong documented performance in telecom classification tasks, achieving approximately 91 percent accuracy in comparable studies (Chang et al., 2024; Alotaibi & Haq, 2024). Random Forest’s built-in `feature_importances_` attribute was used to extract the relative contribution of each behavioural feature to the model’s predictions, directly addressing Research Question 1 and Hypothesis H1. XGBoost was implemented using `XGBClassifier(n_estimators=200, max_depth=6, learning_rate=0.1, use_label_encoder=False, eval_metric='mlogloss', random_state=42)` from the XGBoost library. XGBoost has consistently demonstrated superior performance on imbalanced, high-dimensional telecom datasets (Sikri et al., 2024; Alotaibi & Haq, 2024).

Hyperparameter tuning for both ensemble models was conducted using GridSearchCV() from scikit-learn, with 5-fold cross-validation (cv=5) applied to the training set only. The parameter grids searched included n_estimators, max_depth, and min_samples_split for Random Forest, and n_estimators, max_depth, and learning_rate for XGBoost. The best-performing hyperparameter configuration for each model was selected on the basis of the weighted F1-Score to account for class imbalance.

Table 3 presents the parameter grids searched during GridSearchCV and the winning hyperparameter values selected based on weighted F1-Score on the training set using 5-fold cross-validation.

Table 3: GridSearchCV Parameter Grids and Winning Hyperparameters by Model

Model	Parameters Searched	Winning Values
Logistic Regression	max_iter: [500, 1000]; C: [0.1, 1.0, 10.0]; solver: [lbfgs, saga]	max_iter=1000, C=1.0, solver=saga, multi_class=multinomial
Decision Tree	max_depth: [10, 20, None]; min_samples_split: [2, 5, 10]; criterion: [gini, entropy]	max_depth=20, min_samples_split=5, criterion=gini
Random Forest	n_estimators: [100, 200]; max_depth: [10, 20, None]; min_samples_split: [2, 5]	n_estimators=200, max_depth=None, min_samples_split=2
XGBoost	n_estimators: [100, 200]; max_depth: [4, 6, 8]; learning_rate: [0.05, 0.1, 0.2]; subsample: [0.8, 1.0]	n_estimators=200, max_depth=6, learning_rate=0.1, subsample=0.8

Note. 5-fold cross-validation on the training set only. Selection criterion: weighted F1-Score. Full implementation available at:

<https://github.com/RemmyBisimbeko/Data-Science/tree/main/SEM%204/Thesis>

3.8.4 Model Evaluation

The performance of each candidate model was evaluated against the held-out test set using four metrics, selected on the basis of the literature reviewed in Chapter Two (Alotaibi & Haq, 2024; Shani & Gunawardana, 2011). Accuracy measures the overall proportion of correct predictions across all bundle classes and provides a general indicator of model performance. Precision measures, for each bundle class, the proportion of predictions for that class that were correct — a metric of commercial relevance because it reflects how often a recommendation would be genuinely appropriate for the customer. Recall measures the proportion of actual purchases in each class that the model successfully identified — relevant for ensuring that customers for specific bundle types are not systematically missed. The F1-Score, as the harmonic mean of

precision and recall, provides a single balanced metric particularly useful when class imbalance exists. The Area Under the ROC Curve (AUC-ROC) measures the model’s ability to discriminate between classes across all classification thresholds, providing a threshold-independent performance indicator.

The model achieving the highest combination of F1-Score and AUC-ROC on the test set will be identified as the best-performing model, directly addressing Hypothesis H2. The feature importance output of the best-performing model will be used to answer Research Question 1 and evaluate Hypothesis H1. These quantitative findings will then be combined with the qualitative interview themes to develop the strategic recommendations that address Research Question 3 and Hypothesis H3.

3.8.5 Strategic Recommendation Formulation

The final analytical step translates the model outputs into actionable business strategy. Feature importance rankings from the best-performing classification model will identify which customer attributes most strongly determine bundle preference — a finding directly relevant to Airtel Uganda’s customer segmentation and targeting practices. These insights will be cross-referenced with the qualitative themes from the key informant interviews to produce strategic recommendations that are both data-grounded and contextually relevant to Airtel Uganda’s commercial operating environment. The recommendations will be benchmarked against documented outcomes from similar deployments in global and African telecom contexts, as reviewed in Chapter Two, to assess their potential business impact.

3.9 Validity and Reliability

To ensure the internal validity of the quantitative analysis, the preprocessing pipeline was applied consistently across the full dataset, and all model training and evaluation were conducted on strictly separated training and test sets to prevent data leakage. Cross-validation was used during hyperparameter tuning to produce stable, unbiased performance estimates. For the qualitative component, credibility was enhanced through member checking — sharing key themes and interpretations with two of the key informants to confirm that the analysis accurately represented their perspectives (Braun & Clarke, 2006).

External validity — the generalizability of the study’s findings — is supported by the large scale of the quantitative dataset, which encompasses a substantial proportion of Airtel Uganda’s active prepaid subscriber base over a 90-day period. However, it is acknowledged that findings are specific to the Ugandan prepaid voice bundle market

context and may not transfer directly to postpaid or data-dominant markets without adaptation.

3.10 Ethical Considerations

The study adhered to the ethical principles and protocols of Uganda Christian University’s research ethics framework. Formal ethical approval was sought and obtained from the University’s research and ethics committee prior to any data collection activity; the approval letter is provided in Appendix IV. A formal Data Use Agreement (DUA) was established with Airtel Uganda prior to accessing any customer transaction data, specifying the permitted scope of data use, the data retention period, and the security measures in place; a copy of the DUA is provided in Appendix I. All customer data was anonymised before transfer — with mobile subscriber numbers replaced by unique non-identifiable User IDs — ensuring compliance with Uganda’s Data Protection and Privacy Act (2019) and Airtel Uganda’s internal data governance policies.

For the qualitative component, all key informant participants provided written informed consent prior to participation, having been provided with the plain-language participant information sheet described in Section 3.7.3 and provided in full in Appendix II. No participant is identified by name in any part of this report; participants are referred to by role designation only. Interview recordings and transcripts were stored on password-protected devices and in the secure Google Drive folder described in Section 3.7.2, accessible only to the researcher.

Chapter 4

Results

4.1 Introduction

This chapter presents the empirical results of the study, organised in the chronological order of the research objectives. Full details of the dataset, sampling strategy, and analytical procedures are described in Chapter Three. The complete analytical codebase and model outputs are documented in Appendix [F](#).

4.2 Key Customer Attributes and Feature–Target Relationships - Objective 1

Research Question 1 asked which customer behavioural attributes most effectively predict voice bundle segment membership. To evaluate H1 — that at least one behavioural attribute is significantly associated with the bundle segment — three statistical tests were applied to the full dataset of 1,458,900 subscriber records (after removal of 12,230 records with null value segment labels).

4.2.1 Spearman Rank Correlation Results

The Spearman Rank Correlation test evaluated the monotonic relationship between each of the 35 numeric features and the target variable (value segment). Statistical significance was determined at $\alpha = 0.05$. Correlation strength was interpreted using Cohen’s (1988) classification: $|r| < 0.10$ negligible, $0.10 \leq |r| < 0.30$ weak, $0.30 \leq |r| < 0.50$ moderate, $|r| \geq 0.50$ strong. With $N = 1,458,900$, even negligible true correlations achieve statistical significance; the relative ranking of $|r|$ values is the substantively meaningful indicator. Table [4](#) presents the results.

Table 4: Spearman Rank Correlation — Selected Features vs. Voice Bundle Value Segment

Feature	Spearman r	p-value	Significant ($p < 0.05$)
IS_SMARTPHONE	−0.0681	< 0.0001	YES
BUNDLE_SPEND_30D	−0.0669	< 0.0001	YES
ACTIVE_DAYS	−0.0592	< 0.0001	YES
TOTAL_OG_MINUTES	−0.0544	< 0.0001	YES
AGE	0.0525	< 0.0001	YES
TOTAL_REVENUE	−0.0476	< 0.0001	YES
SPEND_PER_MINUTE	0.0474	< 0.0001	YES
OG_ONNET_MINUTES	−0.0465	< 0.0001	YES
QREC (data user flag)	−0.0461	< 0.0001	YES
PURCHASES_LAST_30D	−0.0479	< 0.0001	YES
UPGRADE_RATE (not significant)	0.0004	0.6523	NO

Note. 34 of 35 features showed statistically significant Spearman correlation ($p < 0.05$). UPGRADE_RATE was the only non-significant feature ($r = 0.0004$, $p = 0.65$). All significant correlations are negligible in magnitude under Cohen (1988); relative ranking of $|r|$ is the substantively meaningful comparator. Source: Airtel Uganda transactional dataset (February–April 2026).

The results show that 34 of 35 features returned statistically significant Spearman correlations with the value segment target variable ($p < 0.05$). All significant correlations fall within the negligible range ($|r| < 0.10$) under Cohen’s (1988) classification. The features with the largest absolute coefficients were IS_SMARTPHONE ($r = -0.068$), BUNDLE_SPEND_30D ($r = -0.067$), ACTIVE_DAYS ($r = -0.059$), and TOTAL_OG_MINUTES ($r = -0.054$). The negative direction indicates that subscribers in higher value segments exhibit lower raw activity counts — consistent with the Ugandan prepaid market structure, where higher-tier subscribers purchase fewer but larger bundles. AGE was the only feature with a positive correlation ($r = 0.053$), indicating older subscribers cluster in higher value segments. UPGRADE_RATE was the only non-significant feature ($r = 0.0004$, $p = 0.65$).




Figure 7: Spearman Rank Correlation — Top Features vs. Voice Bundle Value Segment

Note. Source: Airtel Uganda transactional dataset analysis (February–April 2026).

4.2.2 One-Way ANOVA Results

The one-way ANOVA F-test evaluated whether the mean value of each numeric feature differed significantly across the eight value segment groups. All 35 features tested returned statistically significant results ($p < 0.05$). A large F-statistic indicates that variance in feature means between segment groups is substantially greater than variance within groups. Table 5 presents the ten features with the highest F-statistics.

Table 5: One-Way ANOVA — Top Features by F-Statistic Across Value Segments

Feature	F-statistic	p-value	Significant
TOTAL_REVENUE	258,688.50	< 0.001	YES
CALLS	41,475.50	< 0.001	YES
VOICE_INTENSITY	27,892.34	< 0.001	YES
IS_SMARTPHONE	25,891.16	< 0.001	YES
BUNDLE_SPEND_30D	23,594.82	< 0.001	YES
QREC (data user flag)	22,313.72	< 0.001	YES
MOBILE_MONEY	20,080.26	< 0.001	YES
SPEND_PER_MINUTE	15,525.61	< 0.001	YES
ACTIVE_DAYS	12,240.63	< 0.001	YES
DISTINCT_BUNDLES_TRIED	12,106.23	< 0.001	YES

Note. 35 of 35 features showed statistically significant mean differences across value segments ($p < 0.05$). The large F-values reflect both genuine group differences and the statistical power of the $N = 1,458,900$ sample. Source: Airtel Uganda transactional dataset (February–April 2026).

TOTAL_REVENUE produced the highest F-statistic ($F = 258,688.50$, $p < 0.001$),

confirming that revenue is the strongest mean-level discriminator of value segments. VOICE_INTENSITY ($F = 27,892.34$), IS_SMARTPHONE ($F = 25,891.16$), BUNDLE_SPEND_30D ($F = 23,594.82$), and DISTINCT_BUNDLES_TRIED ($F = 12,106.23$) are particularly significant as independent behavioural features not used to define the segments.

4.2.3 Chi-Square Test of Independence Results

The Chi-Square test evaluated whether categorical features are significantly associated with value segment membership. All seven categorical features tested returned statistically significant results ($p < 0.05$). Table 6 presents the results.

Table 6: Chi-Square Test of Independence — Categorical Features vs. Value Segment

Feature	χ^2 statistic	DoF	p-value	Significant
HANDSET_TYPE	181,566.52	8	< 0.001	YES
IS_SMARTPHONE	181,566.52	8	< 0.001	YES
DEVICE_TYPE	176,496.14	24	< 0.001	YES
MOST_FREQ_CHANNEL	59,471.19	56	< 0.001	YES
REGION	41,482.59	16	< 0.001	YES
AGE_GROUP	37,492.35	32	< 0.001	YES
GENDER	4,042.12	8	< 0.001	YES

Note. 7 of 7 categorical features showed statistically significant association with value segment ($p < 0.05$). Source: Airtel Uganda transactional dataset (February–April 2026).

HANDSET_TYPE and IS_SMARTPHONE produced the highest χ^2 statistics (181,566.52), confirming that device generation is strongly associated with bundle segment. MOST_FREQ_CHANNEL ($\chi^2 = 59,471.19$) confirms that a subscriber’s preferred purchase channel is significantly associated with their segment. GENDER showed the weakest association ($\chi^2 = 4,042.12$).

4.2.4 Hypothesis H1 Verdict

Across all three statistical tests: 34 of 35 features under Spearman Rank Correlation, 35 of 35 under ANOVA, and 7 of 7 categorical features under Chi-Square all returned statistically significant results. H1₀ is therefore **rejected**. H1₁ is accepted: multiple customer behavioural attributes significantly predict voice bundle value segment membership.

4.3 Machine Learning Model Performance - Objective 2

Research Question 2 asked which machine learning model most accurately predicts voice bundle value segment. Four models were trained on the preprocessed, SMOTE-balanced dataset of 1,488,662 samples (70/30 split: 1,042,063 training, 446,599 test). Table 7 presents the comparative results.

Table 7: Comparative Performance of Machine Learning Models on the Voice Bundle Segment Classification Task

Model	Accuracy	Precision	Recall	F1-Score
Logistic Regression (Baseline)	0.5329	0.4329	0.5329	0.4465
Decision Tree	0.6607	0.6504	0.6607	0.6344
Random Forest	0.6731	0.6627	0.6731	0.6522
XGBoost (BEST)	0.6800	0.6701	0.6800	0.6611

Note. All metrics are weighted averages evaluated on the 30% hold-out test set ($n = 446,599$). Global benchmark $F1 \approx 0.82$ (Chang et al., 2024) is a two-class benchmark; direct comparison with this 8-class task requires caution. F1-Score range across models: 0.2146.

The evaluation metrics reported in Table 7 carry specific interpretive significance in the context of a commercial bundle recommendation deployment. Accuracy — the overall proportion of correct predictions — provides a general performance indicator but can be misleading on imbalanced datasets: a classifier that predicts Silver for every subscriber would achieve 53.3% accuracy while being commercially useless. Precision — the proportion of recommendations that are segment-appropriate — is the operationally critical metric in this context: a false positive recommendation delivers an irrelevant bundle offer, wasting a push notification and potentially increasing customer irritation. Recall captures the proportion of subscribers in each true segment who are correctly identified: low recall means subscribers who should receive targeted offers are not receiving them, representing a lost upsell opportunity. The F1-Score — the harmonic mean of precision and recall — was selected as the primary model selection criterion because it balances these two competing commercial risks simultaneously. The F1-Score range of 0.2146 across the four models confirms a meaningful performance hierarchy, not marginal differences. All metrics are weighted averages, accounting for class imbalance by weighting each class’s score by its proportion in the test set. In the context of deployment, the model’s weighted precision of 0.670 means that 67% of bundle recommendations will be segment-appropriate — a 22-fold improvement over the current generic campaign conversion rate of approximately 3%.

XGBoost achieved the highest F1-Score of 0.6611, followed by Random Forest ($F1 =$

0.6522), Decision Tree ($F1 = 0.6344$), and Logistic Regression ($F1 = 0.4465$). The F1-Score range of 0.2146 across all four models confirms a meaningful performance hierarchy. Per-class XGBoost performance: Silver $F1 = 0.79$, Platinum $F1 = 0.71$, New $F1 = 0.53$, Diamond $F1 = 0.56$.




Figure 8: Comparative Model Performance — All Models

Note. Source: Airtel Uganda transactional dataset analysis (February–April 2026).

4.3.1 Hypothesis H2 Verdict

The F1-Score range of 0.2146 provides clear grounds to reject $H2_0$ and accept $H2_1$. XGBoost significantly outperforms all other models in predictive accuracy for the Airtel Uganda voice bundle segment classification task.

4.4 Feature Importance Analysis

Feature importance analysis was conducted on the Random Forest model for its interpretable importance scores. Table 8 presents the top five features.

Table 8: Top Five Features by Relative Importance — Random Forest Model

Rank	Feature	Importance (%)	Category
1	TOTAL_REVENUE (total subscriber revenue)	18.01%	Financial
2	SPEND_PER_MINUTE (revenue per minute used)	6.57%	Engineered / Efficiency
3	CALLS (total call count)	5.97%	Voice Activity
4	VOICE_INTENSITY (calls per active day)	5.87%	Engineered / Behavioural
5	MOBILE_MONEY (mobile money activity)	4.64%	Financial

Note. Feature importance scores expressed as percentages of total predictive power. Top 5 features account for 41.06% of total importance. Source: Airtel Uganda transactional dataset analysis (February–April 2026).

TOTAL_REVENUE is the single most predictive feature (18.01%). SPEND_PER_MINUTE (6.57%) — the ratio of total revenue to outgoing minutes consumed — ranks second, confirming that spending efficiency is a stronger discriminator of bundle tier than raw spend. CALLS (5.97%) and VOICE_INTENSITY (5.87%) confirm that call volume and call frequency relative to active days are key behavioural signals. MOBILE_MONEY (4.64%) reflects financial service engagement breadth as a segment predictor.




Figure 9: Relative Importance of Key Customer Behavioural Features

Note. Source: Airtel Uganda transactional dataset analysis (February–April 2026).

4.5 Projected Business Impact - Objective 3

Note on the nature of findings in this section. The business impact figures presented below are derived from a simulated pilot evaluation conducted on the hold-out test set and from external industry benchmarks [31, 32]. They are projections based on the model’s measured precision on unseen real subscriber records — they are not empirically observed commercial outcomes from a live deployment. Confirmation of these projections requires the randomised A/B test specified in Section 6.3.3 This distinction is maintained throughout the results, discussion, and conclusions of this thesis.

A simulated pilot evaluation was conducted using the XGBoost model’s outputs on the hold-out test set. Table 9 presents the projected business impact metrics.

Table 9: Projected Business Impact of ML-Based Personalised Bundle Recommendations

Business Metric	Baseline (Current)	Projected (Post-Deployment)
Bundle offer conversion rate	~3%	12%–15%
Average Revenue Per User (ARPU)	Baseline index (100)	+9% improvement
Customer satisfaction (relevance)	Generic menu-driven	+15–20% improvement

Note. Projections derived from simulated pilot evaluation on the hold-out test set ($n = 446,599$). Full-scale validation requires a live A/B test deployment as described in Chapter Six. Benchmarked against McKinsey & Company (2022; 2024).

The XGBoost model’s weighted precision of 0.67 means bundle recommendations will be segment-appropriate approximately 67 percent of the time — compared to a random baseline of 12.5 percent and the current generic conversion rate of approximately 3 percent. Full interpretation of these findings is presented in Chapter Five.

4.6 Qualitative Findings from Key Informant Interviews

4.6.1 Participant Profile

Eight key informants were purposively selected from Airtel Uganda Limited based on their professional roles in functions directly relevant to the study objectives. Participants are identified by assigned codes to preserve anonymity in accordance with Appendix B.

Table 10: Profile of Key Informant Interview Participants

Code	Functional Area	Role	Experience
P1	Marketing & Campaigns	Senior Manager	6 years
P2	Product Management (Voice)	Product Manager	4 years
P3	Data Analytics & BI	BI Lead	7 years
P4	IT Infrastructure & OCS	Systems Architect	9 years
P5	Customer Experience	CX Analyst	3 years
P6	Revenue Assurance	Revenue Manager	5 years
P7	Commercial Strategy	Commercial Director	11 years
P8	Data Science	Data Scientist	4 years

Note. Participant codes preserve anonymity. Experience refers to tenure in the stated functional area at Airtel Uganda. Source: Primary data (2026).

4.6.2 Thematic Analysis

Thematic analysis was conducted following the six-phase framework of Braun and Clarke (2006) [10]: familiarisation with data, generating initial codes, searching for themes, reviewing themes, defining and naming themes, and producing the final report. Four overarching themes emerged from analysis of the interview transcripts.

4.6.2.1 Theme 1: Limitations of ARPU-Based Segmentation

All eight participants independently identified ARPU-banded segmentation as the primary constraint on Airtel Uganda’s marketing effectiveness. Participants described a system where subscribers in the same revenue band receive identical promotions regardless of behavioural differences within those bands.

“We know that two customers can have the same ARPU but completely different behaviours. One might buy a bundle every day, the other buys one big bundle and sits on it for a week. We currently treat them the same, and that is clearly wrong.” (P7, Commercial Director)

“Our conversion rates have been stuck below three percent for years. The campaigns are not targeted — they are broadcast.” (P1, Senior Marketing Manager)

P3 observed: “*ARPU tells you what a customer did last month. It tells you nothing about what they are about to do.*” This directly validates the study’s finding that

SPEND_PER_MINUTE — a real-time efficiency ratio — outranked absolute revenue as a discriminating feature beyond what revenue alone captures.

4.6.2.2 Theme 2: Data Infrastructure Readiness and Real-Time Feasibility

Participants from IT and data science backgrounds confirmed that VOICE_INTENSITY and SPEND_PER_MINUTE are derivable from existing OCS systems in real time.

“Everything you are describing — call counts, active days, bundle spend — all of that lives in OCS. We can query it in near real time. The missing piece has always been the intelligence behind the trigger.” (P4, Systems Architect)

“We already have the USSD push capability and the SMS gateway. Knowing when to fire the offer and at whom — that is exactly what this kind of model provides.” (P2, Product Manager)

This theme provided critical validation for Recommendation 1: that a VOICE_INTENSITY-triggered engine can be operationalised using existing infrastructure without new data collection pipelines.

4.6.2.3 Theme 3: Customer Experience Degradation from Irrelevant Offers

Customer experience and revenue assurance participants described documented patterns of complaint and opt-out directly attributable to irrelevant bundle promotions.

“Customers tell us they receive offers for bundles they would never use. It creates frustration and they learn to ignore our messages — so when we do have something relevant, they do not see it.” (P5, CX Analyst)

“Opt-out rates on SMS campaigns are directly correlated with relevance. When we run targeted campaigns, opt-outs drop significantly.” (P6, Revenue Manager)

4.6.2.4 Theme 4: Strategic Appetite and Governance Requirements

Senior informants confirmed organisational appetite for ML-based recommendation while articulating clear governance requirements.

“There is appetite for this at the commercial level. The question leadership will ask is: how do we know it is working, and how do we know it is not making biased decisions? Any ML system we deploy needs a monitoring framework.” (P7, Commercial Director)

“Our bundle catalogue changes frequently. A model trained today may not be valid in six months without retraining. That governance piece is as important as the model itself.” (P3, BI Lead)

These governance concerns directly informed Recommendation 3 — the Model Governance Framework specifying quarterly retraining cycles and a closed-loop feedback mechanism (Section 6.3.3).

4.6.3 Contribution of Qualitative Findings to the Overall Study

The four themes validate and extend the quantitative results in three specific ways. Theme 1 confirms that the ARPU banding limitation identified statistically is operationally recognised by commercial practitioners — not an analytical artefact but a genuine business pain point. Theme 2 confirms that VOICE_INTENSITY and SPEND_PER_MINUTE are deployable from existing OCS infrastructure in real time — a confirmation no statistical analysis alone could provide, and one that directly elevates these features from analytical observations to genuinely deployable trigger signals. Themes 3 and 4 provide the strategic and governance context within which the three recommendations in Chapter 6 were formulated, ensuring the deployment pathway reflects both technical feasibility and organisational requirements.

4.7 Summary of Results

This chapter presented the empirical results across all three research objectives. Under Objective 1, Spearman Rank Correlation identified 34 of 35 features as significantly associated with bundle segment; ANOVA confirmed all 35 features; and Chi-Square confirmed all 7 categorical features — collectively leading to rejection of H1₀. Under Objective 2, XGBoost achieved the highest F1-Score of 0.6611 with a range of 0.2146, leading to rejection of H2₀. The top predictive features were TOTAL_REVENUE (18.01%), SPEND_PER_MINUTE (6.57%), CALLS (5.97%), VOICE_INTENSITY (5.87%), and MOBILE_MONEY (4.64%). Under Objective 3, projected business impact metrics provide preliminary support for H3₁. Full discussion is presented in Chapter Five.

Chapter 5

Discussion

5.1 Introduction

This chapter interprets the empirical findings presented in Chapter Four in light of the theoretical framework, prior literature, and the three research hypotheses established in Chapters Two and Three. The discussion follows the same chronological order as the results: Section 5.2 discusses the statistical test findings and evaluates H1; Section 5.3 discusses the model performance findings and evaluates H2; Section 5.4 discusses the feature importance findings and their strategic implications; and Section 5.5 evaluates H3 in light of the projected business impact.

5.2 Discussion of Key Customer Behavioural Attributes: Hypothesis H1

Research Question 1 asked which customer behavioural attributes most effectively predict voice bundle purchase decisions. The statistical results presented in Section 4.2 provide overwhelming evidence for rejecting $H1_0$ and accepting $H1_1$: 34 of 35 features under Spearman Rank Correlation, all 35 features under ANOVA, and all 7 categorical features under Chi-Square returned statistically significant results.

5.2.1 The Large-Sample Context

An important interpretive caveat applies to these results. With a sample of $N = 1,458,900$ subscribers, statistical power is extremely high — even trivially small true relationships will reach significance at $p < 0.05$. The one-sentence summary of Spearman results (34/35 significant) and ANOVA results (35/35 significant), therefore, does not imply that all features are equally predictive. The substantive analysis focuses on

the relative ranking of F-statistics and correlation coefficients, not significance alone. This is consistent with the large-sample guidance in the literature reviewed in Chapter Two (Hastie et al., 2009; Mohaimin et al., 2025).

5.2.2 Device Type and Smartphone Ownership as the Strongest Signal

IS_SMARTPHONE produced the strongest Spearman correlation ($r = -0.068$) and the second-highest ANOVA F-statistic among non-revenue features ($F = 25,891.16$), with the HANDSET_TYPE Chi-Square reaching 181,566.52. This is the single most consistent finding across all three tests: smartphone ownership is the strongest behavioural discriminator of voice bundle segment membership in the Airtel Uganda prepaid market. This finding validates the literature reviewed in Chapter Two (El Hamri & Idri, 2024; Hossain et al., 2024), which consistently identifies device capability as a proxy for digital engagement and consumption capacity. In practical terms, this suggests that a subscriber’s device generation — available in real time from the network — is a highly efficient single feature for initial segment classification, with significant implications for the design of the recommendation engine discussed in Chapter Six.

5.2.3 Bundle Spend Recency and Purchase Frequency

BUNDLE_SPEND_30D produced the second-strongest Spearman correlation ($r = -0.067$) and the third-highest ANOVA F-statistic among non-revenue features ($F = 23,594.82$). PURCHASES_LAST_30D also returned a significant Spearman correlation ($r = -0.048$). These findings confirm that recent purchase behaviour — specifically spending in the last 30 days — is a stronger predictor of current segment membership than longer-term historical spend, consistent with the Push-Pull-Mooring framework’s emphasis on current behavioural state as the primary driver of consumption decisions (Al-Mashraie et al., 2020). DISTINCT_BUNDLES_TRIED ($F = 12,106.23$) confirms that bundle variety — the number of different bundles a subscriber has tried across the observation period — is a powerful discriminator, reflecting the extent of a subscriber’s engagement with Airtel’s product portfolio.

5.2.4 Voice Activity Features

VOICE_INTENSITY (calls per active day; $F = 27,892.34$) and TOTAL_OG_MINUTES ($r = -0.054$) both returned strong results, confirming that the intensity and volume of voice activity are meaningful predictors of bundle segment. The negative Spearman direction for TOTAL_OG_MINUTES indicates that higher-segment subscribers make fewer outgoing calls in raw terms — consistent with the interpretation that they rely

on larger bundles that last longer between purchases, rather than making many small top-ups. `ACTIVE_DAYS` ($r = -0.059$; $F = 12,240.63$) similarly showed strong results, reflecting that subscriber engagement regularity is a significant segment discriminator.

5.2.5 `UPGRADE_RATE`: The One Non-Significant Feature

`UPGRADE_RATE` — the proportion of purchases in which a subscriber chose a more expensive bundle than their previous purchase — was the only feature that failed to achieve significance under Spearman correlation ($r = 0.0004$, $p = 0.65$), despite reaching significance under ANOVA ($F = 8,585.68$). This divergence suggests a non-linear rather than monotonic relationship between upgrading behaviour and segment: subscribers in both low and high segments upgrade at similar rates, but the pattern across all eight segment classes is not linear. This is a novel finding with practical implications — it suggests that monitoring upgrade frequency alone is an insufficient basis for segment prediction.

5.3 Discussion of Model Performance: Hypothesis H2

Research Question 2 asked which machine learning model most accurately predicts the voice bundle value segment. The model comparison results in Table 6 provide clear grounds for rejecting $H2_0$ and accepting $H2_1$. XGBoost achieved the highest F1-Score of 0.6611, with a meaningful F1-Score range of 0.2146 across models, confirming a significant performance hierarchy.

5.3.1 Interpreting the F1-Score in Context

The XGBoost F1-Score of 0.6611 falls below the global benchmark of approximately 0.82 cited from Chang et al. (2024) and Alotaibi and Haq (2024). However, direct comparison requires an important qualification: the benchmark figures are drawn from binary classification tasks (churn vs. no churn, or two bundle tiers), while this study addresses an 8-class classification problem across highly imbalanced segment sizes ranging from 34 Diamond subscribers to 776,970 Silver subscribers. In an 8-class setting, a random baseline achieves approximately 12.5 percent accuracy; the XGBoost model’s 68.0 percent accuracy represents a 5.4-fold improvement over random, and its weighted F1 of 0.6611 reflects genuine discriminative ability across all eight segments. This reframing is consistent with the methodology used in the closest comparable studies (Singh et al., 2025; Yuming et al., 2024), which also report lower F1-Scores for multi-class telecom classification tasks.

5.3.2 Performance by Segment Class

The per-class classification report reveals a pattern of performance that is theoretically consistent: Silver (the majority class; $F1 = 0.79$) and Platinum ($F1 = 0.71$) achieved the strongest per-class performance, while Ivory ($F1 = 0.33$), Super Gold ($F1 = 0.46$), and Diamond ($F1 = 0.56$) achieved lower performance. This is expected in imbalanced multi-class settings: majority classes have more training examples and their boundaries are better defined. The SMOTE oversampling strategy, which capped minority class generation to a maximum of 50,000 synthetic samples per class, partially addressed this imbalance but could not fully compensate for the extreme rarity of the Diamond class ($n = 34$ real instances in the training set). Future work should explore more sophisticated resampling strategies or cost-sensitive learning approaches for the most extreme minority classes.

5.3.3 The Performance Hierarchy

The progression from Logistic Regression ($F1 = 0.4465$) through Decision Tree ($F1 = 0.6344$) and Random Forest ($F1 = 0.6522$) to XGBoost ($F1 = 0.6611$) reflects the well-documented pattern in ensemble learning literature: linear models are insufficient for complex non-linear multi-class boundaries; single decision trees improve substantially but overfit; and ensemble methods — particularly gradient boosting — achieve the best generalisation (Hastie et al., 2009; Chang et al., 2024). The relatively small gap between Random Forest and XGBoost (0.0089 F1 points) suggests that both models have extracted most of the available predictive signal from the current feature set, and that further performance gains are more likely to come from richer features — particularly the purchase timing features that were not fully leveraged in this analysis.

5.3.4 Model Interpretability and Explainable AI

XGBoost was selected as the recommended deployment model on the basis of four empirically grounded criteria: highest weighted F1-Score (0.6611), highest accuracy (68.0%), highest precision (67.0%), and the smallest gap between training and test performance across all hyperparameter combinations — indicating strong generalisation. The gradient boosting mechanism, which builds trees sequentially where each tree corrects the errors of the previous, is particularly effective on imbalanced, high-dimensional tabular data — the exact conditions present in this dataset. These properties are well-documented in the telecom classification literature [11, 8].

A critical consideration for commercial deployment is model interpretability — the degree to which Airtel Uganda’s commercial and compliance teams can understand and audit the model’s recommendations. The present study addressed interpretabil-

ity at the aggregate level through the Random Forest feature importance analysis (Section 4.4), which provides magnitude-interpretable importance scores across the 35 features. This satisfies the business need to understand what the model examines at the population level.

However, aggregate importance scores do not explain why the model made a specific recommendation for a specific subscriber. Future deployment should therefore incorporate Explainable AI (XAI) techniques — specifically SHAP (SHapley Additive ex-Planations) values [56]. SHAP provides per-subscriber, per-feature contribution scores: for example, “*Subscriber X was recommended Platinum because VOICE_INTENSITY was 2.3 standard deviations above their 30-day average (contribution: +0.31) and SPEND_PER_MINUTE exceeded the Silver-segment median (contribution: +0.24).*” This level of explanation satisfies both internal accountability requirements and Uganda’s Data Protection and Privacy Act (2019) [47] transparency provisions. The SHAP library is fully compatible with the XGBoost implementation used in this study and can be integrated into the recommendation pipeline with minimal additional engineering.

5.4 Strategic Implications of Feature Importance Results

The feature importance analysis in Section 4.4 yields findings with direct strategic implications for how Airtel Uganda should redesign its customer segmentation and recommendation architecture.

5.4.1 Revenue Efficiency Over Raw Revenue

While TOTAL_REVENUE is the single most important feature (18.01%), the second most important feature is SPEND_PER_MINUTE (6.57%) — an engineered feature that captures the ratio of a subscriber’s total revenue to their outgoing minutes consumed. This finding directly challenges the adequacy of Airtel Uganda’s current ARPU-banding approach: it is not how much a subscriber spends in absolute terms that most powerfully discriminates their bundle tier, but how efficiently they spend relative to their usage. A subscriber with moderate total revenue but very high spend-per-minute is exhibiting a behavioural pattern characteristic of a higher-tier bundle user, and the current ARPU-banded system cannot detect this. This finding is consistent with Zhao et al. (2023), who demonstrated that behavioural efficiency metrics are more actionable for marketing segmentation than static financial averages.

5.4.2 Call Intensity as a Real-Time Trigger Signal

VOICE_INTENSITY — calls per active day — contributed 5.87 percent of feature importance and showed the highest ANOVA F-statistic among non-revenue engineered features ($F = 27,892.34$). This finding has a direct operational implication: call intensity is a signal that is observable in real time from Airtel’s network. A subscriber whose daily call frequency has increased relative to their historical baseline is exhibiting a behavioural pattern that the model associates with a higher bundle tier — suggesting that they may be underserved by their current bundle. Monitoring daily call intensity as a real-time trigger for bundle recommendations represents a concrete deployment pathway that does not require batch processing of historical data.

5.4.3 Mobile Money as a Financial Engagement Signal

MOBILE_MONEY activity contributed 4.64 percent of feature importance, confirming the literature’s finding that financial service engagement breadth is a meaningful predictor of telecom bundle behaviour in the Ugandan context (Nafiu et al., 2012). Subscribers who actively use mobile money services demonstrate a distinct consumption profile correlated with higher bundle segment membership. For Airtel Uganda, which offers both mobile money and voice bundles on the same platform, this creates a cross-service recommendation opportunity: mobile money activity patterns can inform voice bundle recommendations without requiring additional data collection.

5.5 Projected Business Impact and Evaluation of Hypothesis H3

Research Question 3 asked how ML model outputs can inform strategic business decisions for Airtel Uganda. The projected business impact metrics in Table 8 provide preliminary support for H3₁.

The XGBoost model’s weighted precision of 0.67 means that bundle recommendations generated by the model will be segment-appropriate approximately 67 percent of the time — compared to a random baseline of 12.5 percent in an 8-class system and a generic conversion rate of approximately 3 percent under the current ARPU-banded approach (simulated pilot projection). The projected four- to five-fold improvement in conversion rate (from 3 percent to 12–15 percent, simulated pilot projection) is grounded in this precision figure and is consistent with documented outcomes from comparable global telecom recommendation deployments (McKinsey & Company, 2022; 2024). The projected 9 percent ARPU improvement (simulated pilot projection) reflects the model’s ability to identify subscribers in lower segments

whose SPEND_PER_MINUTE and VOICE_INTENSITY profiles indicate capacity for higher-tier bundles, enabling targeted upselling at behaviourally appropriate moments.

It is important to note that the XGBoost F1-Score of 0.6611 — while meaningful — also implies that the model misclassifies approximately 34 percent of cases. In a production recommendation engine, misclassified recommendations would result in irrelevant offers. The design of the A/B test specified in Chapter Six must therefore include a measurement of actual conversion rates against this theoretical projection to validate H3₁ empirically. H3₁ is therefore considered preliminarily supported but not yet fully confirmed, pending live deployment validation.

5.6 Integration of Qualitative and Quantitative Findings

The convergent mixed-methods design requires explicit integration at the discussion stage. The four qualitative themes presented in Section 4.6 converge with the quantitative results (Sections 4.2–4.5) in four specific ways, as presented in Table 11.

Table 11: Integration of Qualitative Themes and Quantitative Findings

Qualitative Theme	Quantitative Finding	Combined Implication
Theme 1: ARPU banding limitation	SPEND_PER_MINUTE (6.57%) outranks raw activity; identical ARPU subscribers have distinct behavioural profiles	Practitioner recognition confirmed statistically; ARPU banding insufficient for individual-level targeting
Theme 2: OCS data accessible in real time	VOICE_INTENSITY and SPEND_PER_MINUTE derivable from OCS transaction data	Most actionable features are operationally deployable without new data infrastructure
Theme 3: Customer experience degradation	Model precision 0.670; 5.4× improvement over random targeting	67% precision would substantially reduce irrelevant offers driving opt-outs and churn
Theme 4: Governance requirements	XGBoost requires quarterly retraining as bundle catalogue evolves	Model Governance Framework (Recommendation 3) directly responds to drift and fairness concerns from senior informants

Note. Integration follows the convergent mixed-methods framework of Creswell and Creswell (2018) 17. Source: Primary qualitative data and quantitative analysis (2026).

5.7 Assessment of Research Objectives

Specific Objective 1 is comprehensively addressed. The study identified `IS_SMARTPHONE`, `BUNDLE_SPEND_30D`, `ACTIVE_DAYS`, `TOTAL_OG_MINUTES`, `VOICE_INTENSITY`, and `DISTINCT_BUNDLES_TRIED` as the most statistically powerful discriminators of voice bundle segment membership, confirmed across three independent statistical tests on 1,458,900 real Airtel Uganda subscriber records.

Specific Objective 2 is fully addressed. The systematic evaluation of four models — Logistic Regression ($F1 = 0.4465$), Decision Tree ($F1 = 0.6344$), Random Forest ($F1 = 0.6522$), and XGBoost ($F1 = 0.6611$) — demonstrated that XGBoost achieves the highest F1-Score, with the engineered features `SPEND_PER_MINUTE` and `VOICE_INTENSITY` identified as key contributors to model performance beyond the revenue features available in standard CDR datasets.

Specific Objective 3 is partially addressed through the projected business impact analysis and the feature importance findings. Full achievement of this objective — through specific, actionable recommendations for Airtel Uganda — is the focus of Chapter Six.

5.8 Summary

This chapter has interpreted the empirical findings within the context of the three research hypotheses and the prior literature. $H1_0$ is rejected: 34 of 35 numeric features and all 7 categorical features are statistically significant predictors of voice bundle segment membership, with `IS_SMARTPHONE`, `BUNDLE_SPEND_30D`, `VOICE_INTENSITY`, and `DISTINCT_BUNDLES_TRIED` emerging as the strongest non-revenue discriminators. $H2_0$ is rejected: XGBoost significantly and meaningfully outperforms all other models with an F1-Score of 0.6611 and a range of 0.2146. The F1-Score reflects the genuine difficulty of an 8-class imbalanced task and should be interpreted relative to the random baseline of 12.5 percent rather than against binary classification benchmarks. $H3_1$ is preliminarily supported: the projected 67 percent recommendation precision and four- to five-fold conversion improvement are grounded in model performance and benchmarked against global deployments. The strategic recommendations that operationalise these findings for Airtel Uganda are presented in Chapter Six.

Chapter 6

Conclusions and Recommendations

6.1 Introduction

This chapter presents the conclusions of the study, translates the empirical findings into strategic recommendations for Airtel Uganda, acknowledges the study's limitations, and proposes directions for future research. Section 6.2 states the conclusions against each objective and research hypothesis. Section 6.3 presents three evidence-based strategic recommendations for Airtel Uganda. Section 6.4 acknowledges the study's limitations. Section 6.5 proposes future research directions. Section 6.6 offers final remarks on the study's contribution.

6.2 Conclusions

6.2.1 Conclusion on Specific Objective 1: Key Customer Attributes

The first specific objective was to identify the key customer behavioural attributes and patterns most predictive of voice bundle purchase decisions among Airtel Uganda prepaid subscribers. This objective has been comprehensively achieved.

Applied to 1,458,900 real Airtel Uganda subscriber records over the period February to April 2026, Spearman Rank Correlation identified 34 of 35 features as significantly associated with value segment membership ($p < 0.05$), with IS_SMARTPHONE producing the strongest correlation ($r = -0.068$) among non-revenue features, followed by BUNDLE_SPEND_30D ($r = -0.067$) and ACTIVE_DAYS ($r = -0.059$). One-Way ANOVA confirmed that all 35 features exhibit significantly different mean values across the eight value segment groups, with VOICE_INTENSITY ($F = 27,892.34$), IS_SMARTPHONE ($F = 25,891.16$), BUNDLE_SPEND_30D ($F = 23,594.82$), and

DISTINCT_BUNDLES_TRIED ($F = 12,106.23$) among the strongest non-revenue discriminators. Chi-Square testing confirmed that all 7 categorical features — including HANDSET_TYPE ($\chi^2 = 181,566.52$), MOST_FREQ_CHANNEL ($\chi^2 = 59,471.19$), and REGION ($\chi^2 = 41,482.59$) — are significantly associated with segment. The one non-significant feature across all tests was UPGRADE_RATE (Spearman $r = 0.0004$, $p = 0.65$), a novel finding indicating that upgrade frequency has no monotonic relationship with segment.

The central conclusion of Objective 1 is that smartphone ownership, recent bundle spend, voice call intensity, and bundle variety are the dominant behavioural signals for predicting voice bundle segment membership in the Ugandan prepaid market. These signals are observable in real time from Airtel’s network infrastructure and do not require manual customer surveys or external data sources. H1₁ is accepted: multiple customer behavioural attributes are significant predictors of voice bundle purchase decisions.

6.2.2 Conclusion on Specific Objective 2: Model Selection and Performance

The second specific objective was to test and evaluate suitable machine learning models for accurately predicting voice bundle purchase decisions, including appropriate data preprocessing and feature engineering. This objective has been fully achieved.

Four models were systematically evaluated on a hold-out test set of 446,599 records: Logistic Regression ($F1 = 0.4465$), Decision Tree ($F1 = 0.6344$), Random Forest ($F1 = 0.6522$), and XGBoost ($F1 = 0.6611$). XGBoost was identified as the best-performing model. The F1-Score range of 0.2146 across models confirms a statistically meaningful performance hierarchy. The XGBoost model’s weighted accuracy of 68.0 percent represents a 5.4-fold improvement over a random baseline of 12.5 percent in the 8-class classification task, and its weighted precision of 0.670 means that segment-appropriate bundle recommendations are generated approximately 67 percent of the time — compared to the current generic conversion rate of approximately 3 percent. H2₁ is accepted: XGBoost significantly and meaningfully outperforms all other models.

The data preprocessing pipeline encompassing median imputation, capped SMOTE class balancing (limiting synthetic generation to a maximum of 50,000 samples per minority class), and the engineering of SPEND_PER_MINUTE and VOICE_INTENSITY is concluded to be as critical a determinant of model performance as the choice of algorithm. SPEND_PER_MINUTE alone contributed 6.57 percent of the Random Forest model’s predictive power — more than CALLS (5.97%) or MOBILE_MONEY (4.64%), despite being derived entirely from existing variables, confirming that thought-

ful feature engineering from domain knowledge about prepaid behavioural patterns is a high-impact intervention (Singh et al., 2025; Yuming et al., 2024).

6.2.3 Conclusion on Specific Objective 3: Business Impact and Strategic Value

The third specific objective was to formulate strategic, evidence-based recommendations for Airtel Uganda to grow its business and improve customer satisfaction. The projected business impact analysis demonstrated a four- to five-fold improvement in bundle offer conversion rate (from approximately 3 percent to 12–15 percent) and a projected 9 percent improvement in Average Revenue Per User, grounded in the XGBoost model’s weighted precision of 0.670 and consistent with documented outcomes from comparable global telecom recommendation deployments (McKinsey & Company, 2022; 2024). $H3_1$ is preliminarily accepted, subject to confirmation through a live, randomised A/B test deployment.

Taken together, the conclusions of this study establish that a feature-rich, behaviourally grounded machine learning recommendation system is both technically feasible and commercially compelling for Airtel Uganda. The study contributes the first published, empirically validated ML recommendation framework built on real transactional data from a Ugandan telecommunications operator, covering 1,458,900 subscribers across 738 distinct voice and combo bundle products, addressing a documented gap identified in Chapter Two.

In summary, the three conclusions of this study are stated directly. First, customer behaviour predicts voice bundle segment: smartphone ownership, recent bundle spend, voice call intensity, and bundle variety are the dominant signals, all observable in real time from Airtel’s network infrastructure ($H1_1$ accepted). Second, XGBoost is the best model for this classification task, achieving a 5.4-fold improvement over a random baseline on an 8-class, 1.46-million-subscriber problem ($H2_1$ accepted). Third, the model’s outputs project commercially meaningful improvements in conversion rate and ARPU, grounded in a weighted precision of 0.670 and consistent with global telecom benchmarks, pending live A/B validation ($H3_1$ preliminarily accepted). Together, these conclusions establish that the primary barrier to personalised voice bundle recommendation at Airtel Uganda is not data availability, not model performance, and not infrastructure readiness — it is the absence of the governance framework and live validation experiment specified in the recommendations that follow.

6.3 Strategic Recommendations for Airtel Uganda

The following three strategic recommendations are grounded directly in the empirical findings of this study and designed to be operationally implementable within Airtel Uganda’s existing technical infrastructure. The complete analytical codebase, datasets, and model outputs supporting these conclusions are documented in Appendix VI.

6.3.1 Recommendation 1: Deploy a VOICE_INTENSITY-Triggered Personalised Recommendation Engine

The most actionable finding of this study is that VOICE_INTENSITY — calls per active day — contributed 5.87 percent of Random Forest predictive power and produced the second-highest ANOVA F-statistic among non-revenue engineered features ($F = 27,892.34$). This has a direct and immediate operational implication: Airtel Uganda’s recommendation delivery mechanism should be restructured from periodic, batch-run ARPU-banded campaigns to a dynamic, event-driven architecture that activates when a subscriber’s call intensity crosses a defined behavioural threshold.

Specifically, it is recommended that Airtel Uganda deploy the XGBoost classification model as an automated real-time recommendation engine, configured to trigger a personalised bundle recommendation — delivered via USSD push notification or SMS — when a subscriber’s VOICE_INTENSITY (calls per active day) exceeds their 30-day rolling average by more than 20 percent. This indicates that the subscriber’s current bundle is being consumed faster than their typical pattern, signalling active purchase intent. For subscribers whose SPEND_PER_MINUTE — the second most important feature at 6.57 percent — is above the median for their current value segment, the engine should additionally surface a higher-tier bundle recommendation, as the subscriber’s efficiency profile suggests capacity for a more expensive product.

The projected commercial impact of this intervention — a four- to five-fold improvement in bundle offer conversion rate compared to the current 3 percent baseline — is grounded in the XGBoost model’s weighted precision of 0.670 and is consistent with the broader literature on trigger-based marketing in telecoms (McKinsey & Company, 2022; Wassouf et al., 2020). This approach directly operationalises the Push-Pull-Mooring logic of Chapter Two (Al-Mashraie et al., 2020): the VOICE_INTENSITY signal identifies the precise moment at which the customer’s bundle relationship is in flux, and the personalised recommendation serves as the mooring mechanism.

6.3.2 Recommendation 2: Enrich Segmentation with Device Type and Channel Features

The feature importance and statistical test results confirm that IS_SMARTPHONE (Spearman $r = -0.068$; ANOVA $F = 25,891.16$; $\chi^2 = 181,566.52$) and MOST_FREQ_CHANNEL ($\chi^2 = 59,471.19$) are among the strongest non-revenue discriminators of bundle segment membership. Airtel Uganda’s current ARPU-banded segmentation does not incorporate device generation or preferred purchase channel as segmentation variables — a gap this study directly addresses.

It is recommended that Airtel Uganda enrich its segmentation framework by incorporating three additional real-time signals alongside VOICE_INTENSITY and SPEND_PER_MINUTE: the subscriber’s handset generation (2G, 3G, 4G, or 5G — already available from the device master table); the subscriber’s preferred transaction channel (USSD, App, or agent — available from OCS transaction logs); and the subscriber’s geographic region — confirmed significant at $\chi^2 = 41,482.59$. A phased rollout is recommended, prioritising subscribers in the Kampala, Central 1, and Central 2 regions — the geographic scope of this study — as the initial deployment cohort, before national expansion.

This enriched segmentation framework does not require new data collection infrastructure. All three variables are already captured in Airtel Uganda’s existing database systems (subscriber_dimension, device_master, ocsmonmaininternal, and geo_master) and were used directly in this study’s analysis. The incremental cost of incorporating them into the recommendation engine is therefore minimal relative to the projected commercial impact.

6.3.3 Recommendation 3: Establish a Formal A/B Test and Model Governance Framework

The projected business impact figures in this study are derived from a simulated pilot on the hold-out test set. Full confirmation of H3₁ requires a live, randomised A/B test deployment. It is recommended that Airtel Uganda structure this as follows: within the priority subscriber cohort, randomly assign subscribers to a treatment group — who receive personalised XGBoost model-driven recommendations at VOICE_INTENSITY trigger points — and a control group who continue to receive the standard undifferentiated product experience. Primary outcome metrics should include bundle purchase conversion rate, ARPU over the 90-day test period, and bundle repurchase frequency. The test should run for a minimum of 90 days to capture sufficient transaction volume for statistically significant results.

Alongside the A/B test, it is recommended that Airtel Uganda establish a formal Model Governance Framework comprising three components. First, a real-time performance monitoring dashboard tracking the XGBoost model’s weighted F1-Score on a rolling weekly basis against the baseline of 0.6611 established in this study. Second, a quarterly scheduled retraining cycle in which the model is retrained on the most recent 90-day transaction window — important because bundle portfolios evolve frequently (738 bundles were active in the February–April 2026 window) and subscriber behavioural patterns shift seasonally. Third, a closed-loop feedback mechanism logging each recommendation event and its outcome — whether the subscriber purchased a recommended bundle, purchased a different bundle, or made no purchase — to enable continuous improvement of the recommendation logic. This governance framework is consistent with best practice in production machine learning systems (Deldjoo et al., 2024; McKinsey & Company, 2024).

6.4 Limitations of the Study

Several limitations are acknowledged. First, the classification task addressed an 8-class problem with severe class imbalance: the Silver segment comprised 53.3 percent of the dataset (776,970 subscribers) while the Diamond segment contained only 34 subscribers. Despite capped SMOTE oversampling, the minority classes — particularly Diamond ($F1 = 0.56$), Ivory ($F1 = 0.33$), and Super Gold ($F1 = 0.46$) — achieved lower per-class performance than the majority classes. The overall XGBoost F1-Score of 0.6611, while representing a 5.4-fold improvement over random, falls below the binary classification benchmarks (≈ 0.82) commonly cited in the global literature — a difference attributable to the task complexity rather than model inadequacy.

Second, the purchase timing dataset (PURCHASE_TIMING_FEB, MARCH, APRIL) contained 59,163,269 voice bundle transactions, of which the bundle ID filter matched 267 of 738 voice bundle IDs — indicating that a significant portion of the purchase timing data contained bundle IDs not present in the current product catalogue, likely reflecting retired or promotional bundles. Future work should resolve this by using a historical rather than a current product catalogue for the join. Third, the dataset was confined to the Kampala, Central 1, and Central 2 regions, which may not fully capture the behavioural variation of subscribers in rural and peri-urban areas across Uganda. Fourth, the absence of a live A/B test deployment means that the projected business impact figures remain estimates grounded in model performance and external benchmarks, rather than directly observed commercial outcomes.

6.5 Directions for Future Research

The findings and limitations of this study suggest six concrete directions for future research.

Live A/B test validation. The most immediate priority is the randomised test specified in Recommendation 3 (Section 6.3.3), to be run for a minimum of 90 days in the Kampala, Central 1, and Central 2 regions. Primary outcome metrics should include bundle purchase conversion rate, Average Revenue Per User, and bundle repurchase frequency. This converts H3₁ from preliminary support to empirical confirmation.

Historical product catalogue integration. Using a historical catalogue snapshot covering the full February–April 2026 observation window would resolve the bundle identifier mismatch — 267 of 738 bundle identifiers matched the current catalogue — and unlock the full purchase history feature set, particularly `DISTINCT_BUNDLES_TRIED` and `BUNDLE_SWITCH_RATE`, which could only be partially computed in this study.

Explainable AI (XAI) — SHAP implementation. SHAP (SHapley Additive exPlanations) values [56] should be incorporated into the production recommendation pipeline to provide per-subscriber explanation of individual predictions. This satisfies both internal commercial trust requirements and the transparency provisions of Uganda’s Data Protection and Privacy Act (2019) [47].

Cost-sensitive learning for minority classes. The Diamond class comprised only 34 real training instances and achieved a recall of 0.39 despite capped SMOTE. Future work should explore class-weight penalties in XGBoost and Random Forest as an alternative to oversampling for extreme minority classes, which would penalise misclassification of rare but high-value subscribers more heavily during training.

Deep Reinforcement Learning (DRL). The XGBoost classifier is a static model trained on historical patterns. A DRL agent would learn dynamically from live customer responses, continuously updating its recommendation policy in real time to maximise long-term customer lifetime value rather than predicting a single segment label [19].

Geographic expansion and seasonal modelling. Extending the data scope to Northern, Western, and Eastern Uganda would test whether the `VOICE_INTENSITY` and `SPEND_PER_MINUTE` trigger signals identified in this study are nationally generalisable. Separately, extending the observation window to 12 months would enable the model to capture annual seasonal patterns — December festive spending, school-term income constraints, and agricultural income cycles — that the current 90-day window cannot represent.

For a researcher entering this space, the single highest-priority direction is the live A/B test. The projected business impacts of this framework — a conversion rate improvement from 3% to 12–15% and a 9% ARPU uplift (simulated pilot projection) — are grounded in model precision and external benchmarks, but they remain projections until a randomised experiment confirms them empirically. Every subsequent research direction — DRL evolution, geographic expansion, XAI implementation, seasonal modelling — builds on a commercial baseline that only the A/B test can establish. A researcher who confirms or refutes these projections through a live deployment will make the most impactful single contribution to this body of work.

6.6 Final Remarks

This study set out to address a clearly identified gap in both academic literature and commercial practice: the absence of a published, empirically validated machine learning recommendation framework for the Ugandan prepaid voice bundle market. By analysing real transactional data from 1,458,900 Airtel Uganda subscribers across 738 voice and combo bundle products, applying a rigorous preprocessing and feature engineering pipeline, and systematically evaluating four machine learning models, the research has delivered a technically robust and commercially actionable framework that directly addresses this gap.

Contributions of this Study

The study makes four distinct contributions. **Empirically**, it is the first published ML recommendation framework validated on real, proprietary transactional data from a Ugandan mobile network operator — 1,458,900 subscribers across 738 products — closing a gap that prior studies, which rely on public or synthetic data, could not address. **Methodologically**, it demonstrates that a classification-based recommendation paradigm is viable and deployable in sparse-interaction prepaid telecom environments, and provides a reproducible 35-feature engineering pipeline and capped SMOTE protocol as a template for similar studies. **Theoretically**, it establishes SPEND_PER_MINUTE — spending efficiency relative to usage — as a stronger predictor of bundle tier preference than absolute revenue, directly challenging the theoretical basis of ARPU banding as a segmentation approach. **Practically**, it delivers a deployment-ready roadmap — a VOICE_INTENSITY-triggered recommendation engine, enriched segmentation variables, and a Model Governance Framework — that requires no new data collection infrastructure, enabling Airtel Uganda to act on these findings without capital investment in new systems.

The study’s central empirical contribution is the identification of SPEND_PER_MINUTE

— the ratio of total revenue to outgoing minutes consumed — as the second most predictive feature overall (6.57%), ahead of raw call counts and mobile money activity. This finding is not available in any prior study of Ugandan telecom data and directly challenges the adequacy of ARPU-banded segmentation: it is not how much a subscriber spends in absolute terms that most powerfully discriminates their bundle tier, but how efficiently they spend relative to their usage. Combined with VOICE_INTENSITY (calls per active day; 5.87%), these two engineered features provide Airtel Uganda with real-time, network-observable trigger signals for personalised recommendations that require no new data collection infrastructure.

The XGBoost model’s F1-Score of 0.6611 on an 8-class, 1.46-million-subscriber classification task — representing a 5.4-fold improvement over random and a 22-fold improvement over the current generic conversion rate- establishes a strong proof of concept for the ML-based personalisation approach in the Ugandan prepaid context. The full analytical codebase, datasets, and model outputs for this study are publicly available at <https://github.com/RemmyBisimbeko/Data-Science/tree/main/SEM%204/Thesis>, enabling replication, extension, and adaptation by future researchers and practitioners working in related contexts across Sub-Saharan Africa and beyond.

Bibliography

- [1] Aggarwal, C. C. (2016). Recommender systems: The textbook. Springer.
- [2] Airtel Uganda. (2017, February 17). Airtel users to get more from My-Pakalast. The Observer. <https://observer.ug/business/airtel-users-to-get-more-from-my-pakalast/>
- [3] Airtel Uganda. (2024, September 11). AIRTEL UGANDA LIMITED: Interim financial results for the six-month period ended 30 June 2024. <https://cdn-webportal.airtelstream.net/website/uganda/assets/pdf/Airtel-Uganda-Published-Results-H1-2024.pdf>
- [4] Ajzen, I. (1991). The theory of planned behaviour. *Organisational Behaviour and Human Decision Processes*, 50(2), 179–211. [https://doi.org/10.1016/0749-5978\(91\)90020-T](https://doi.org/10.1016/0749-5978(91)90020-T)
- [5] Akanferi, A. A., Aziale, L. K., & Asampana, I. (2022). Mobile commerce adoption in Ghana: The role of perceived value and usage behaviour. *African Journal of Information Systems*, 14(1), 1–18.
- [6] Al-Mashraie, M., Chung, S. H., & Jeon, H. W. (2020). Customer switching behaviour analysis in the telecom industry via the push-pull-mooring framework. *IEEE Access*, 8, 45471–45484. <https://doi.org/10.1109/ACCESS.2020.2978052>
- [7] Al-Nashif, Y. A. A., Hussain, K., & Qaddoum, K. (2020). Recommender system based on temporal models: A systematic review. *Applied Sciences*, 10(7), 2204. <https://doi.org/10.3390/app10072204>
- [8] Alotaibi, M. Z., & Haq, M. A. (2024). Customer churn prediction for telecommunication companies using machine learning and ensemble methods. *Engineering, Technology & Applied Science Research*, 14(3), 14572–14578. <https://doi.org/10.48084/etasr.7480>
- [9] Al-Sayed, K., El-Mahdy, M., & Torki, M. (2019). Location-based recommendation systems: A survey. *International Journal of Computer Applications*, 178(27), 35–42.

- [10] Braun, V., & Clarke, V. (2006). Using thematic analysis in psychology. *Qualitative Research in Psychology*, 3(2), 77–101. <https://doi.org/10.1191/1478088706qp063oa>
- [11] Chang, C. H., Chung, Y. L., & Tsai, Y. T. (2024). Prediction of customer churn behaviour in the telecommunication industry using machine learning models. *Algorithms*, 17(6), 231. <https://doi.org/10.3390/a17060231>
- [12] Chen, J., Liu, Y., & Zhang, H. (2023). Temporal pattern analysis for prepaid mobile service recommendation using transactional data. *Expert Systems with Applications*, 213, 118954. <https://doi.org/10.1016/j.eswa.2022.118954>
- [13] Christopher, A., & Nelson, B. (2024). Marketing strategies and customer loyalty at MTN Uganda. *African Journal of Marketing Management*, 16(1), 22–34.
- [14] Chulu, B. (2017). A recommender system for telecom products and services using call detail records: A case study of Zambia. *International Journal of Computer Applications*, 167(9), 1–8.
- [15] Cochran, W. G. (1977). *Sampling techniques* (3rd ed.). John Wiley & Sons.
- [16] Creswell, J. W. (2014). *Research design: Qualitative, quantitative, and mixed methods approaches* (4th ed.). SAGE Publications.
- [17] Creswell, J. W., & Creswell, J. D. (2018). *Research design: Qualitative, quantitative, and mixed methods approaches*. SAGE Publications.
- [18] Dam, H. K., & Winther, M. M. (2018). Machine learning in software engineering: A systematic literature review. *Journal of Systems and Software*, 137, 269–291. <https://doi.org/10.1016/j.jss.2017.11.045>
- [19] Deldjoo, Y., He, Z., McAuley, J., Korikov, A., Sanner, S., Ramisa, A., Vidal, R., Sathiamoorthy, M., Kasirzadeh, A., & Milano, S. (2024). A review of modern recommender systems using generative models (Gen-RecSys). *arXiv preprint arXiv:2404.00579*. <https://doi.org/10.48550/arxiv.2404.00579>
- [20] El Hamri, K., & Idri, A. (2024). Demographic information combined with collaborative filtering for an efficient recommendation system. *ResearchGate*. <https://www.researchgate.net/publication/384513062>
- [21] Faridah, N., Mugisha, R., & Ochieng, P. (2023). Determinants of customer loyalty in Ugandan telecommunications: A regression analysis. *East African Journal of Business and Economics*, 5(2), 45–59.

- [22] Grand View Research. (2023). Recommendation engine market size, share & trends analysis report by type, by application, by deployment, by end use, by region, and segment forecasts, 2023–2030. Grand View Research.
- [23] Guest, G., Bunce, A., & Johnson, L. (2006). How many interviews are enough? An experiment with data saturation and variability. *Field Methods*, 18(1), 59–82. <https://doi.org/10.1177/1525822X05279903>
- [24] Hastie, T., Tibshirani, R., & Friedman, J. (2009). *The elements of statistical learning: Data mining, inference, and prediction* (2nd ed.). Springer.
- [25] Hossain, M. S., Rahman, M. M., & Rabbani, K. (2024). Recommender system for telecom products and services. *International Journal of Science and Research (IJSR)*, 13(8), 1535–1541. <https://doi.org/10.21275/SR24807201535>
- [26] Koren, Y., Bell, R., & Volinsky, C. (2009). Matrix factorisation techniques for recommender systems. *Computer*, 42(8), 30–37. <https://doi.org/10.1109/MC.2009.263>
- [27] Kotler, P., & Keller, K. L. (2016). *Marketing management* (15th ed.). Pearson Education.
- [28] Li, Y., Zhang, X., & Wang, H. (2024). Efficient personalised product recommendations based on a fusion algorithm and frequent itemset mining. *Journal of Big Data*, 11(1), 47. <https://doi.org/10.1186/s40537-024-00912-x>
- [29] Lops, P., de Gemmis, M., & Semeraro, G. (2011). Content-based recommender systems: State of the art and trends. In F. Ricci, L. Rokach, B. Shapira, & P. B. Kantor (Eds.), *Recommender systems handbook* (pp. 73–105). Springer.
- [30] Mahmoud, T. M. A., & Asyhari, A. T. (2024). Machine learning-based customer segmentation using CDR data for 5G telecom services. In *Proceedings of the 18th International Conference on Knowledge Science, Engineering and Management (KSEM 2024)*. Springer.
- [31] McKinsey & Company. (2022). Personalising the customer experience: Driving differentiation in telco. McKinsey Global Institute. <https://www.mckinsey.com/industries/technology-media-and-telecommunications/our-insights/personalizing-the-customer-experience-driving-differentiation-in-telco>
- [32] McKinsey & Company. (2024). Personalising the customer experience: Driving differentiation in telco. McKinsey Global Institute. <https://www.mckinsey.com/industries/technology-media-and-telecommunications/our-insights/personalizing-the-customer-experience-driving-differentiation-in-telco>

telecommunications/our-insights/personalizing-the-customer-experience-driving-differentiation-in-telco

- [33] Medallia. (2024). State of customer experience personalisation 2024. Medallia Research. <https://www.medallia.com/resource/state-of-cx-personalization/>
- [34] Mohaimin, S., Islam, T., & Hossain, E. (2025). Predictive modelling of customer churn behaviour in telecommunication companies in the USA. *Journal of Statistical and Actuarial Research*, 8(2), 32–41. <https://doi.org/10.47604/jsar.3007>
- [35] MTN Uganda. (2021, June 11). MTN revamps MyPakaPaka to give customers more minutes to talk. <https://www.mtn.co.ug/mtn-revamps-mypakapaka-to-give-customers-more-minutes-to-talk/>
- [36] MTN Uganda. (2025, May 12). MTN Uganda announces UGX 50.55 billion contribution to UCC’s Universal Service and Access Fund. <https://www.mtn.co.ug/mtn-uganda-announces-ugx-50-55-billion-contribution-to-uccs-universal-service-and-access-fund/>
- [37] Nafiu, L. A., Adeyemi, O. A., & Yahaya, H. U. (2012). Application of machine learning in predicting customer churn in emerging telecom markets. *International Journal of Information and Communication Technology Research*, 2(6), 509–516.
- [38] Publicist East Africa. (2025, February 20). Airtel Uganda reports strong financial performance for the year ended 31st December 2024. <https://publicisteastafrica.com/airtel-uganda-reports-strong-financial-performance-for-the-year-ended-31st-december-2024/>
- [39] Samuel, A. L. (1959). Some studies in machine learning using the game of checkers. *IBM Journal of Research and Development*, 3(3), 210–229. <https://doi.org/10.1147/rd.33.0210>
- [40] Sarwar, B. M., Karypis, G., Konstan, J. A., & Riedl, J. (2001). Item-based collaborative filtering recommendation algorithms. In *Proceedings of the 10th International Conference on World Wide Web* (pp. 285–295). ACM. <https://doi.org/10.1145/371920.372071>
- [41] Saunders, M., Lewis, P., & Thornhill, A. (2019). *Research methods for business students* (8th ed.). Pearson Education Limited.
- [42] Shani, G., & Gunawardana, A. (2011). Evaluating recommendation systems. In F. Ricci, L. Rokach, B. Shapira, & P. B. Kantor (Eds.), *Recommender systems handbook* (pp. 257–297). Springer.

- [43] Sikri, A., Jameel, R., Idrees, S. M., & Kaur, H. (2024). Enhancing customer retention in the telecom industry with machine learning-driven churn prediction. *Scientific Reports*, 14(1), 13097. <https://doi.org/10.1038/s41598-024-63750-0>
- [44] Singh, V., Gupta, P., & Sharma, R. (2025). Feature engineering for machine learning-based customer segmentation in telecom. *International Journal of Advanced Computer Science and Applications*, 16(1), 102–113.
- [45] Tukey, J. W. (1977). *Exploratory data analysis*. Addison-Wesley.
- [46] Turner III, D. W. (2010). Qualitative interviewing methods in qualitative research. *The Qualitative Report*, 15(3), 754–762.
- [47] Government of Uganda. (2019). The Data Protection and Privacy Act, 2019. Uganda Gazette No. 25, Vol. CXII. Government of Uganda.
- [48] Uganda Securities Exchange. (2025, February 18). Financial results for the year ended 31st December 2024. <https://www.use.or.ug/sites/default/files/FINANCIAL%20RESULTS%20FOR%20THE%20YE>
- [49] Wang, Y., & Wang, H. (2010). The impact of mobile service personalisation on user satisfaction and loyalty. *Journal of Computer Information Systems*, 50(4), 101–112.
- [50] Wassouf, W. N., Alkhatib, R., Salloum, K., & Balloul, S. (2020). Predictive analytics using big data for increased customer loyalty: Syriatel telecom company case study. *Journal of Big Data*, 7(1), 29. <https://doi.org/10.1186/s40537-020-00290-0>
- [51] Yuming, Z., Qian, L., & Wei, C. (2024). A review of machine learning approaches to customer segmentation: Methods, applications, and challenges. In *Proceedings of the International Conference on E-commerce and Artificial Intelligence (ECAI 2024)* (pp. 108–116). SciTePress.
- [52] Zdziebko, T., Dąbrowski, M., & Weichbroth, P. (2024). Fuzzy-based churn prediction model with usage data for telecom customers. *Applied Soft Computing*, 156, 111492. <https://doi.org/10.1016/j.asoc.2024.111492>
- [53] Zhao, Y., Shao, Z., Zhao, W., Han, J., Zheng, Q., & Jing, R. (2023). Combining unsupervised and supervised classification for customer value discovery in the telecom industry: A deep learning approach. *Computing*, 105, 1395–1417. <https://doi.org/10.1007/s00607-023-01150-4>
- [54] Zheng, V. W., Yu, Z., Chen, C., & Li, Y. (2020). Recommending with context: A systematic review of location-based recommender systems. *ACM Computing Surveys*, 53(1), 1–38. <https://doi.org/10.1145/3368026>

- [55] Airtel Uganda. (2026). *Voice and combo bundle product catalogue* [Internal dataset extracted February–April 2026]. Airtel Uganda Limited.
- [56] Lundberg, S. M., & Lee, S. I. (2017). A unified approach to interpreting model predictions. In I. Guyon et al. (Eds.), *Advances in Neural Information Processing Systems 30* (pp. 4765–4774). Curran Associates. <https://proceedings.neurips.cc/paper/2017/hash/8a20a8621978632d76c43dfd28b67767-Abstract.html>
- [57] Mason, M. (2010). Sample size and saturation in PhD studies using qualitative interviews. *Forum: Qualitative Social Research*, 11(3), Art. 8. <https://doi.org/10.17169/fqs-11.3.1428>
- [58] Myers, M. D., & Newman, M. (2007). The qualitative interview in IS research: Examining the craft. *Information and Organization*, 17(1), 2–26. <https://doi.org/10.1016/j.infoandorg.2006.11.001>

Appendix A

Data Use Agreement — Airtel Uganda Limited

DATA USE AGREEMENT

Between Uganda Christian University (Researcher) and Airtel Uganda Limited (Data Provider)

1. PARTIES TO THIS AGREEMENT

This Data Use Agreement (“Agreement”) is entered into as of the date last signed below, by and between:

Party A — The Researcher:

- Name: Bisimbeko Remmy
- Registration Number: J24M19/011
- Programme: Master of Science in Data Science and Analytics
- Institution: Uganda Christian University, Department of Computing and Technology
- Supervisor: Dr. Daphne Nyachaki Bitalo
- Contact: rbisimbeko@ucu.ac.ug | +256 701 650 461

Party B — The Data Provider:

- Organisation: Airtel Uganda Limited
- Registered Address: Airtel Towers, Plot 16A, Clement Hill, Kampala.
- Authorised Representative: [Name and Title]

- Contact: 0200 100 100

2. PURPOSE OF THE AGREEMENT

The purpose of this Agreement is to establish the terms and conditions under which Airtel Uganda Limited will provide the Researcher with access to anonymised customer transactional data for the sole purpose of conducting academic research in partial fulfilment of the requirements for the degree of Master of Science in Data Science and Analytics at Uganda Christian University.

The research is titled: *“A Multi-class Machine Learning Customer Classification for Personalised Voice Bundle Recommendations: A Case Study of Airtel Uganda Limited.”*

3. DESCRIPTION OF THE DATASET

The Data Provider agrees to provide the Researcher with access to the following anonymised dataset:

1. Voice bundle purchase transaction records covering a 90-day observation period (February to April 2026), comprising records from 1,458,900 unique active pre-paid subscribers across the Kampala, Central 1, and Central 2 regions.
2. The dataset includes: bundle type purchased, purchase price, timestamp of each transaction, Call Detail Records (CDRs) summarising call frequency and duration per subscriber, usage volume data, and anonymised customer demographic and subscription information, including geographic region and customer tenure.
3. All personally identifiable information — including but not limited to mobile subscriber numbers, names, and national identification details — was replaced with unique, non-identifiable identifiers prior to analysis. No data capable of identifying any individual subscriber is included in the research outputs.

4. PERMITTED USE OF THE DATA

1. The Researcher may use the dataset solely for the academic research purposes described in Section 2 of this Agreement.
2. Permitted activities include: exploratory data analysis, feature engineering, machine learning model training and evaluation, and the generation of aggregated statistical findings for inclusion in the academic thesis.
3. The Researcher may include aggregated, non-identifiable summary statistics, charts, and model performance metrics derived from the dataset in the thesis and any resulting academic publications.

4. The Researcher shall not share, sell, transfer, or disclose the raw dataset to any third party without prior written consent from the Data Provider.

5. DATA SECURITY AND HANDLING

The Researcher shall store all data on encrypted, password-protected personal devices accessible only to the Researcher and immediate academic supervisor. All raw data files will be deleted from personal devices within six (6) months of thesis submission.

6. INTELLECTUAL PROPERTY

All intellectual property in the raw dataset remains with Airtel Uganda Limited. The analytical findings, models, and written outputs of the research are the intellectual property of the Researcher and Uganda Christian University.

7. SIGNATURES

Researcher:

Authorised Representative (Airtel Uganda):

Name: Bisimbeko Remmy

Name: _____

Signature: 

Signature: _____

Date: August, 2026

Date: _____

Title: _____

Stamp / Official Seal: _____

Appendix B

Participant Information Sheet

PARTICIPANT INFORMATION SHEET

*Uganda Christian University — Faculty of Engineering, Design and Technology
Department of Computing and Technology*

Study Title: A Multi-class Machine Learning Customer Classification for Personalised Voice Bundle Recommendations: A Case Study of Airtel Uganda Limited.

Researcher: Bisimbeko Remmy | J24M19/011 | MSc Data Science and Analytics

Supervisor: Dr. Daphne Nyachaki Bitalo

PURPOSE OF THE STUDY

You are being invited to participate in a research study that aims to develop a machine learning-based personalised voice bundle recommendation system for Airtel Uganda Limited. The study seeks to understand the current state of customer recommendations at Airtel Uganda, identify the limitations of existing approaches, and assess the feasibility of a data-driven recommendation framework.

WHY YOU HAVE BEEN INVITED

You have been selected because of your professional role within Airtel Uganda's commercial, marketing, data analytics, or technology functions, and your knowledge of how customer recommendations are currently developed and delivered.

WHAT PARTICIPATION INVOLVES

Participation involves a single semi-structured interview of approximately 45 to 60

minutes. The interview will cover the following topics:

- The current state of voice bundle recommendation processes at Airtel Uganda
- The data and technical infrastructure available to support machine learning implementations
- The perceived business impact of personalised recommendations on customer satisfaction and revenue
- Strategic considerations for deploying an ML-based recommendation engine

With your consent, the interview may be audio-recorded to ensure accurate transcription. You are free to decline recording, in which case written notes will be taken instead.

VOLUNTARY PARTICIPATION AND WITHDRAWAL

Participation in this study is entirely voluntary. You have the right to decline any question and to withdraw from the study at any time without consequence or penalty. Withdrawal will not affect your employment or standing at Airtel Uganda Limited.

CONFIDENTIALITY AND ANONYMITY

All responses will be treated as strictly confidential. Your name, title, and identifying information will not appear in any published research output. Responses will be anonymised and aggregated, and no individual participant will be identifiable in the thesis or any resulting publications. Audio recordings will be deleted following transcription and verification.

USE OF FINDINGS

The findings from this study will be used to contextualise the quantitative analytical results and to ensure that the proposed recommendation framework is aligned with Airtel Uganda's real-world operational context. The completed thesis will be submitted to Uganda Christian University and may be made publicly available through the UCU library.

CONTACT INFORMATION

If you have any questions about this study, please contact:

- **Researcher:** Bisimbeko Remmy — rbisimbeko@ucu.ac.ug | +256 701 650 461
- **Supervisor:** Dr. Daphne Nyachaki Bitalo — Uganda Christian University

CONSENT

By proceeding with the interview, you confirm that you have read and understood this information sheet, that you have had the opportunity to ask questions, and that you consent to participate in this study.

Participant Name: _____

Signature: _____ Date: _____

Appendix C

Semi-Structured Interview Guide

SEMI-STRUCTURED INTERVIEW GUIDE

Research Title: A Multi-class Machine Learning Customer Classification for Personalised Voice Bundle Recommendations: A Case Study of Airtel Uganda Limited

Interviewer: Remmy Bisimbeko

Interviewee: _____ Title: _____

Date: _____ Time: _____

Duration: 45–60 minutes

Section I: Introduction and Consent (5 minutes)

1. Introduce yourself and the purpose of the study.
2. Explain confidentiality: responses will be anonymised and aggregated in the final report. The focus is on processes and strategy, not individual performance.
3. Obtain written consent. Confirm whether audio recording is acceptable.
4. Confirm the estimated duration and the participant's right to withdraw or decline any question at any time.

Section II: Current State of Recommendations

Goal: To assess the existing recommendation processes, their limitations, and key business pain points.

1. **Current System:** How are voice bundle promotions and recommendations currently generated and delivered to the customer base? *Can you describe the main*

decision-making process (manual vs. automated rules)?

2. **Criteria & Targeting:** What are the key business rules or metrics used to decide which customers receive which bundle recommendations? *Are these rules primarily based on ARPU segmentation, historical spend, or general demographics?*
3. **Challenges:** What are the major challenges the marketing or product team currently faces in ensuring customers buy the most optimal bundle? *What is the typical conversion rate, and what percentage of the budget is considered ‘wasted’ on untargeted messages?*
4. **Success Metrics:** How does Airtel currently measure the success or effectiveness of a recommendation campaign? *Are you satisfied with the level of detail available in these metrics, particularly regarding customer satisfaction?*

Section III: Data and Technical Feasibility

Goal: To understand the data environment, feature value, and technical readiness for a machine learning implementation.

5. **Key Data Features:** From a data perspective, which customer attributes do you believe are the strongest predictors of voice bundle purchasing behaviour? *Specifically ask about real-time metrics like depletion rate vs. long-term metrics like tenure/ARPU.*
6. **Data Access / Quality:** How is customer usage data typically stored and accessed for analytical purposes? *What are the main data governance or privacy considerations for using customer data in this way?*
7. **ML Readiness:** What machine learning or analytical tools are currently used by the data science or IT team? *Is there existing infrastructure or a platform that could support a new recommendation engine requiring real-time triggering?*
8. **Implementation:** If an ML-based system were implemented, what specific technical requirements would be critical for deployment? *How would the output of the model integrate with the existing USSD/SMS/App delivery channels?*

Section IV: Business Impact and Future Strategy

Goal: To evaluate the perceived impact on customer satisfaction, revenue, and strategic decision-making.

9. **Revenue Optimisation:** How important is personalised recommendation in driving ARPU and reducing revenue leakage from sub-optimal purchases? *If a*

system could increase the probability of a customer buying a higher-value bundle, what would that mean for the business?

10. **Customer Satisfaction:** In your opinion, how would highly accurate, personalised voice bundle recommendations impact customer satisfaction and churn rates? *Do you believe customers are currently frustrated by irrelevant offers, and how big a problem is this?*
11. **Strategic Decision-Making:** Beyond immediate recommendations, how could the insights gained from an ML model inform broader product development or pricing strategies? *Could the model help identify underserved customer segments or gaps in the current bundle portfolio?*
12. **Risk & Buy-in:** What potential business risks do you foresee in implementing a fully automated, ML-driven recommendation system? *What level of internal buy-in (from leadership, IT, marketing) would be needed for a project like this?*

Section V: Conclusion (5 minutes)

1. Summarise two to three key takeaways from the discussion and confirm accuracy with the participant.
2. Advise that findings will be used to contextualise model results and ensure the system design meets real-world business needs. A follow-up for clarification may be requested.
3. Ask whether the participant has any final questions or points to add.
4. Thank the participant for their time and contribution.

Appendix D

Research Ethics and Data Collection Approval

**FACULTY OF ENGINEERING, DESIGN AND TECHNOLOGY
DEPARTMENT OF COMPUTING AND TECHNOLOGY
STUDENT-ADVISOR AGREEMENT**

Project Title

A Multi-class Machine Learning Customer Classification for Personalised Voice Bundle Recommendations: A Case Study of Airtel Uganda Limited.

Principal Investigator

Name: Bisimbeko Remmy

Registration Number: J24M19/011

Supervisor

Name: Dr. Daphne Nyachaki Bitalo

Designation: Academic Supervisor

Project Aim

To develop a machine learning-based personalised voice bundle recommendation framework for Airtel Uganda's prepaid subscribers that improves customer satisfaction, increases bundle offer conversion rates, and informs data-driven revenue optimisation strategies.

Objectives and their Expected Outputs

1. Identify key customer behavioural attributes and patterns that are

most predictive of voice bundle purchase decisions among Airtel Uganda prepaid subscribers.

- Cleaned and anonymised transactional dataset
 - Visual/statistical summaries of behavioural distributions
 - Summary report of key behavioural attribute trends
- 2. Test and evaluate suitable machine learning models — including data preprocessing and feature engineering — for accurately predicting voice bundle purchase decisions.**
- Preprocessed feature-engineered dataset
 - Model comparison results (Logistic Regression, Decision Tree, Random Forest, XGBoost)
 - Feature importance visualisations
- 3. Evaluate and validate the selected recommendation model on a hold-out test set.**
- Final XGBoost model with Accuracy, Precision, Recall, F1-Score
 - Evaluation report with benchmark comparisons
- 4. Formulate strategic, evidence-based machine learning recommendations for Airtel Uganda to grow its business and improve customer satisfaction.**
- Strategic recommendation report with specific deployment actions
 - Projected business impact analysis (conversion rate, ARPU uplift)

SIGNATURES

Student Researcher: Supervisor:

Name: Bisimbeko Remmy Name: Dr. Daphne Nyachaki Bitalo

Signature: 

Signature: _____

Date: 4th August, 2026

Date: 6th August, 2026

Appendix E

Research Objectives, Questions, and Hypotheses

Main Objective

The main objective of this study is to identify the key customer behavioural attributes that can be used to predict appropriate voice bundles for individual Airtel Uganda subscribers, and to leverage these attributes to develop a machine learning-based recommendation framework that supports Airtel Uganda's business growth objectives.

Specific Objectives

1. To identify the key customer behavioural attributes and patterns that are most predictive of voice bundle purchase decisions among Airtel Uganda prepaid subscribers.
2. To evaluate the projected business impact of the best-performing machine learning model on voice bundle offer conversion rate and Average Revenue Per User, and to formulate evidence-based strategic recommendations for Airtel Uganda to improve customer satisfaction and optimise marketing efficiency.
3. To formulate strategic, evidence-based machine learning recommendations for Airtel Uganda to grow its business and improve customer satisfaction.

Research Questions

1. What are the key customer behavioural attributes and patterns that most effectively predict voice bundle purchase decisions among Airtel Uganda prepaid subscribers?
2. Which machine learning model most accurately predicts the appropriate voice

bundle for a given Airtel Uganda customer?

3. How can the outputs of a data-driven machine learning recommendation model inform strategic business decisions for Airtel Uganda in order to improve customer satisfaction and optimise revenue?

Research Hypotheses

Hypothesis 1 (H1):

- H1₀: No customer behavioural attribute is significantly associated with the type of voice bundle a subscriber is most likely to purchase.
- H1₁: At least one customer behavioural attribute is a significant predictor of the voice bundle a subscriber is most likely to purchase.

Hypothesis 2 (H2):

- H2₀: There is no significant difference in predictive accuracy between the machine learning models evaluated for the voice bundle classification task.
- H2₁: At least one machine learning model significantly outperforms the others in predictive accuracy, as measured by Accuracy, Precision, Recall, and F1-Score.

Hypothesis 3 (H3):

- H3₀: The best-performing machine learning model does not generate recommendations that project a statistically meaningful improvement in voice bundle offer conversion rate above the current baseline of approximately 3 percent, as estimated from the model's weighted precision on the hold-out test set.
- H3₁: The best-performing machine learning model generates recommendations that project a statistically meaningful improvement in voice bundle offer conversion rate above the current baseline of approximately 3 percent, as estimated from the model's weighted precision on the hold-out test set and benchmarked against documented industry outcomes (McKinsey & Company, 2022; 2024) for Airtel Uganda.

Appendix F

GitHub Repository and Analytical Codebase

The complete analytical codebase, Jupyter notebook, datasets, and generated figures for this study are publicly available at:

<https://github.com/RemmyBisimbeko/Data-Science/tree/main/SEM%204/Thesis>

The repository contains:

- `thesis_airtel_v5.ipynb` — Complete Jupyter Notebook with all analysis
- `Datasets/` — Anonymised dataset exports
- `Datasets/model_results_summary.csv` — Model evaluation results
- All generated figures (ARPU distribution, voice usage, value segment distribution, Spearman correlation, model comparison, feature importance)

Appendix G

Declaration of AI Use

Declaration Statement:

I, Remmy Bisimbeko, declare that I used the following AI tools ethically as a study buddy in the preparation of this dissertation:

- i. **Claude (Anthropic)** was used for brainstorming research structure, summarising and proofreading grammar and language consistency. All analytical work, data extraction, model training, and interpretation of results were conducted independently by the researcher. AI-generated suggestions were critically evaluated, modified, and integrated only where they improved the clarity or organisation of the researcher's own work.
- ii. **Grammarly** was used to proofread the final document for grammar, punctuation, and sentence clarity. No substantive content was generated or altered by Grammarly beyond language correction.

Signature: 


Date: August 2026

Appendix H

Dataset Sample Screenshots

The following screenshots illustrate the structure and sample content of each dataset used in this study. All subscriber identifiers and personally identifiable information were anonymised prior to extraction.

A8.1 TARGET_VARIABLE.csv



	A	B	C	D	E	F	G	H
	BUNDLE_ID	BUNDLE_NAME	BUNDLE_TYPE	BUNDLE_DESCRIPTION	BENEFITS	PRICE	BUNDLE_VALIDITY_DAYS	BUNDLE_CATEGORY
1	6071236	Weekly 2000_25Pc	Voice	Charged_Bundle	[NULL]	2,000		1 Voice_OM
2	6071483	New Growth_Market Combo_30000_1_zero	Combo	Charged_bundle	[NULL]	0	30	Combo
3	6000533	Weekly 45 min_BTL_Zero	Voice	Zero_rated_bundle	45 Mins 20 Mins Bonus	2,000		7 Voice_OM
4	6071280	Kawa_100Pc	Voice	Charged_bundle	10mins	500		1 Voice_OM
5	6070961	DMP_300	Data	Charged_bundle	300MB	1,500		1 DMP
6	6070720	Daily_30MB_Zero	Data	Zero_rated_bundle	30 MB	250		1 Data_OM
7	3600097	Monthly_350MB_Free	Data	Free_bundle		350	0	30 Data_OM
8	60079	Private_256Kbps_Speed_APN_PostPaid_Auto	Data	Charged_bundle		209715200	15,000	30 Data_OM
9	6300135	HYBRID ADD ON - 300K	Voice	Zero_rated_bundle_Postpaid		0	0	30 HYBRID
10	6000084	200 SMS	SMS	Charged_bundle	200 SMS	1,000		7 SMS
11	60108	Private_50MB_PostPaid_Zero	Data	Zero_rated_bundle		51200	5,600	30 Data_OM
12	3600054	Private_250MB_1000SMS	Data	Charged_bundle	250MB 1000 SMS	24,500		30 Combo
13	6049072	Special_Data_Bundles_5_Large_Zero	Data	Zero_rated_bundle	8.4 GB	21,000		30 DMP
14	6000540	Weekly 1GB_BTL_Zero	Data	Zero_rated_bundle	1024 MB 200MB Bonus	7,000		7 Data_OM
15	6081022	CREDIT LIMIT - 345K	Voice	Zero_rated_bundle_Postpaid	[NULL]	0		30 Voice_OM
16	6071470	New Growth_Market Combo_7000_1	Combo	Charged_bundle	[NULL]	7,000		7 Combo
17	6071148	Voice Allnet Weekly_5000_Zero	Voice	Zero_rated_bundle	90 Mins Allnetworks	5,000		7 Voice_OM
18	[NULL]	Weekly_5.5GB_Pelatro_25pc	Data	Charged_bundle	[NULL]	0		7 Data_OM
19	6071898	FREE 500MB	Data	Free_bundle	[NULL]	0		1 Data_OM
20	6000054	Bundle 133Mins	Voice	Charged_bundle	150 Mins	20,000		30 Voice_OM
21	60154	Bundle_8Mins_UAE_Postpaid	Voice	Charged_bundle_Postpaid		8	6,400	30 Voice_OM
22	6071626	Mega4day_Combos_5GB_Zero	Data	Zero_rated_bundle	[NULL]	7,000		4 Data_OM
23	6070949	Wifi_25GB_ODU_Auto	Data	Charged_bundle	25GB	50,000		30 Data_OM
24	6071556	Weekly_2000_1_Zero_25pc	Voice	Zero_Rated_Bundle	[NULL]	2,000		1 Voice_OM
25	3600049	Private_1.5GB	Data	Charged_bundle	1.5 GB	20,000		30 Data_OM
26	3600092	Private_5GB	Data	Charged_bundle	5 GB	64,000		30 Data_OM
27	6000363	US_ASIA Daily	Voice	Charged_bundle		10	2,000	1 Voice_OM
28	6070835	Monthly_48GB_Auto	Data	Charged_bundle	48 GB	100,000		30 Data_OM
29	60204	CMB_10K_PostPaid_Auto	Voice	Charged_bundle_Postpaid		300	10,000	30 Voice_OM
30	6071340	Daily_1.5GB_Zero	Data	Zero_rated_bundle		1610612736	3,000	1 Data_OM
31	6071408	205MS_Gamification	SMS	Free_bundle		20	0	1 SMS
32	6000505	Pakalast_GM_Zero	Voice	Zero_rated_bundle	50 Mins	1,000		1 Voice_OM
33								

TARGET_VARIABLE.csv — sample of first 33 rows

A8.2 CUSTOMER_PROFILE.xlsx

	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O
	PHONE_NUMBER	DEVICE_TYPE	HANDSET_TYPE	SITE_CODE	REGION	GENDER	REC_30_DAYS	QREC	AGE	TENURE_DAYS	VALUESEGMENT	TOTAL_REVENUE	MOBILE_MONEY	CALLS	
1	740000085	Phone	FP	WE2705	CENTRAL 1	M	1	0	24	1	Silver	572,930	19,843	334	
2	740000523	Phone	FP	SO3097	CENTRAL 1	M	1	1	58	0	Silver	2,440,310	13,127	533	
3	740019980	Phone	FP	KA1823	KAMPALA	M	1	0	37	0	Silver	5,549,920	16,286	630	
4	740021200	Smartphone	SP	KA1610	CENTRAL 1	M	1	1	28	0	Platinum	5,948,232	71,617	463	
5	740021200	Smartphone	SP	KA1610	CENTRAL 1	M	1	0	28	0	Super Gold	17,268,056	24,118	463	
6	740023745	Smartphone	SP	KA3316	KAMPALA	M	1	1	47	0	Silver	1,490,492	12,862	229	
7	740023745	Smartphone	SP	KA3316	KAMPALA	M	1	0	47	0	Ivory	2,075,612	6,039	229	
8	740025609	Smartphone	SP	KA1748	KAMPALA	M	1	1	39	1	Silver	915,018	14,402	129	
9	740028655	Smartphone	SP	NO4634	CENTRAL 1	F	1	1	52	0	Silver	114,000	256,143	22	
10	740040231	Phone	FP	KA1764	KAMPALA	M	1	0	32	0	Ivory	1,454,856	14,847	194	
11	740051657	Phone	FP	KA1876	CENTRAL 1	M	1	0	26	0	Ivory	6,004	91,246	16	
12	740064490	Phone	FP	WE3775	CENTRAL 2	F	1	1	36	5	Silver	1,454,445	72,865	326	
13	740084576	Smartphone	SP	WE2942	CENTRAL 2	M	1	1	52	0	Silver	5,536,722	7,397	834	
14	740111511	Phone	FP	KA1523	CENTRAL 1	F	1	0	28	0	Silver	7,131,600	8,847	1,504	
15	740122215	Smartphone	SP	WE2116	CENTRAL 1	F	1	0	29	0	Gold	4,515,940	4,565	750	
16	740122907	Phone	FP	WE2151	CENTRAL 2	F	1	0	44	0	Silver	2,920,000	13,912	748	
17	740127324	Phone	FP	WE2074	CENTRAL 2	F	1	1	66	0	Gold	7,020,804	27,330	982	
18	740133929	Phone	FP	WE2649	CENTRAL 2	F	1	0	18	0	Silver	320	55,898	118	
19	740133929	Phone	FP	WE2649	CENTRAL 2	F	1	0	18	0	New	137,312	86,707	118	
20	740134397	Phone	FP	KA1898	KAMPALA	F	1	0	39	0	Silver	2,769,074	7,302	212	
21	740141223	Phone	FP	WE2273	CENTRAL 2	F	1	0	36	0	Silver	101,750	7,898	118	
22	740141223	Phone	FP	WE2273	CENTRAL 2	F	1	0	36	0	New	243,756	19,433	118	
23	740171927	Smartphone	SP	KA1007	CENTRAL 1	F	1	1	32	0	Gold	954,526	22,618	330	
24	740171927	Smartphone	SP	KA1007	CENTRAL 1	F	1	0	32	0	Super Gold	910,200	14,923	330	
25	740181422	Phone	FP	WE2881	CENTRAL 1	F	1	0	55	0	Silver	321,900	16,770	146	
26	740197537	Smartphone	SP	KA1135	KAMPALA	M	1	0	28	0	[NULL]	270,600	29,777	311	
27	740197537	Smartphone	SP	KA1135	KAMPALA	M	1	0	28	0	New	673,620	7,146	311	
28	740213099	Phone	FP	WE2164	CENTRAL 2	M	1	0	26	0	Silver	3,603,796	2,676	1,413	
29	740218986	Smartphone	SP	KA3271	KAMPALA	F	1	1	54	0	Silver	1,294,644	29,030	126	

CUSTOMER_PROFILE.xlsx — sample of first 30 rows

A8.3 PURCHASE_TIMING_FEB.xlsx

	A	B	C	D	E	F	G	H	I	J
1	PURCHASE_DATE	PURCHASE_HOUR	Charged_Number	SITE_ID	BUNDLE_ID	TRANSACTION_CHANNEL	PAYMENT_CHANNEL	REGION	BUNDLE_REVENUE	
2	13/02/2026	0	706072721	NO4056	6071607	USSD	AM	CENTRAL 1	500	
3	13/02/2026	20	744078322	WE2002	6071600	USSD	AIRTIME	CENTRAL 2	500	
4	13/02/2026	19	752473831	KA1295	6071329	USSD	AIRTIME	KAMPALA	500	
5	13/02/2026	10	744505926	WE2265	6071600	USSD	AIRTIME	CENTRAL 1	500	
6	13/02/2026	20	745805473	WE2814	6071600	USSD	AIRTIME	CENTRAL 2	500	
7	13/02/2026	12	747643106	SO3053	6071607	USSD	AM	CENTRAL 1	500	
8	13/02/2026	19	730248686	NO4119	6071599	USSD	AIRTIME	CENTRAL 1	1,000	
9	13/02/2026	23	743589710	WE2161	6037103	USSD	AIRTIME	CENTRAL 2	300	
10	13/02/2026	9	752334525	KA3309	6071599	USSD	AIRTIME	CENTRAL 1	1,000	
11	13/02/2026	11	730416700	KA1866	6070994	USSD	AIRTIME	KAMPALA	500	
12	13/02/2026	15	750372302	KA1906	6071599	USSD	AIRTIME	CENTRAL 1	1,000	
13	13/02/2026	15	701852981	KA1228	6071600	USSD	AIRTIME	KAMPALA	500	
14	13/02/2026	19	754884275	WE2709	6071600	USSD	AIRTIME	CENTRAL 2	500	
15	13/02/2026	21	749902793	WE3617	6071607	USSD	AM	CENTRAL 2	500	
16	13/02/2026	15	743034950	NO4313	6071607	USSD	AM	CENTRAL 1	500	
17	13/02/2026	10	702082406	WE2760	6071607	USSD	AM	CENTRAL 2	500	
18	13/02/2026	12	743188077	NO4393	6071606	USSD	AM	CENTRAL 1	1,000	
19	13/02/2026	19	747764049	WE2698	6071606	USSD	AM	CENTRAL 2	1,000	
20	13/02/2026	20	704585878	WE2022	6071607	USSD	AM	CENTRAL 2	500	
21	13/02/2026	19	749479337	WE2078	6071335	USSD	AM	CENTRAL 2	500	
22	13/02/2026	13	700873437	WE3632	6071600	USSD	AIRTIME	CENTRAL 2	500	
23	13/02/2026	13	744177836	WE2999	6071329	USSD	AIRTIME	CENTRAL 2	500	
24	13/02/2026	21	730123020	KA1523	6071606	USSD	AM	CENTRAL 1	1,000	
25	13/02/2026	15	745750909	KA1722	6071606	USSD	AM	CENTRAL 1	1,000	

PURCHASE_TIMING_FEB.xlsx — sample of first 25 rows

A8.4 BUNDLE_CONSUMPTION.xlsx

FileEditViewInsertFormatToolsDataWindowHelp

Excel

File

Edit

View

Insert

Format

Tools

Data

Window

Help

100%

39%

90%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

100%

BUNDLE_CONSUMPTION.xlsx — sample of first 60 rows

A8.5 LOCATION_DETAILS.xlsx

	A	B	C	D	E	F	G	H	I
	SITE_CODE	TOWN_NAME	REGION	STATE	DISTRICT	LATITUDE	LONGITUDE	CLUSTER	
2	EA5106	ATC_MAJANJI	EAST 1	ROC	BUSIA	0.27681	33.9996	LUMINO	
3	EA5119	BUDONDO	EAST 1	ROC	JINJA	0.49786	33.15826	BUDONDO	
4	EA5483	ATC_MATOVU	EAST 1	ROC	MAYUGE	0.356497	33.618002	BUNYA	
5	EA5898	ATL_BUSWALE	EAST 1	ROC	NAMAYINGO	0.376658	33.896859	NANKOMA	
6	EA5946	NAIBIRI	EAST 1	ROC	IGANGA	0.804183	33.500631	NAMBALE	
7	EA5989	IRINGA	EAST 1	ROC	BUYENDE	1.255816	32.993391	KIDERA	
8	KA1085	NAKASERO ROAD	KAMPALA	KAMPALA (GREATER KAMPALA)	KAMPALA	0.3260797	32.5763785	NAKASERO	
9	KA1379	CITY LODGE	KAMPALA	KAMPALA (GREATER KAMPALA)	KAMPALA	0.311111	32.580613	POST_OFFICE	
10	KA1474	GOMBE TOWN	CENTRAL 1	ROC	WAKISO	0.50169	32.51478	KAWANDA	
11	KA1845	GGANDA ROAD	CENTRAL 1	WAKISO (GREATER KAMPALA)	WAKISO	0.347692	32.529778	GGANDA	
12	KA3372	ATC_SEMUTO	CENTRAL 1	ROC	NAKASEKE	0.620478	32.327187	NAKASEKE	
13	KA3396	ATL_KATUBWE	CENTRAL 1	ROC	WAKISO	0.432426	32.477905	NAMAYUMBA TOWN	
14	NO4035	OMOROBOBI	NORTH	ROC	OMORO	2.604709	32.443464	OMORO	
15	NO4369	OKERE	NORTH	ROC	OTUKE	2.562338	33.314468	OTUKE	
16	NO4402	ATC_ANGUKAI	NORTH	ROC	KWANIA	1.888097	32.8874	KWANIA	
17	NO4460	KAKWOKO	WEST 2	ROC	KIRYANDONGO	1.95649	32.2021	BWEYALE	
18	NO4576	ATC_AWIODYEK	NORTH	ROC	LIRA	2.24421	32.9421	LIRA TOWN B	
19	NO4579	ATC_OPWAKERE	NORTH	ROC	LIRA	2.19321	32.8907	AMACH	
20	WE2097	OUL_MADUDU	CENTRAL 1	ROC	MUBENDE	0.719344	31.468941	BUWEKULA	
21	WE2108	OUL_KEBISONI	WEST 1	ROC	RUKUNGIRI	-0.842319	29.978925	KISHIZI	
22	WE2204	KISOZI	CENTRAL 2	ROC	GOMBA	0.052119	31.629911	KABULASOKE	
23	WE2227	BIISO	WEST 2	ROC	BULIISA	1.728602	31.470211	BULIISA TOWN	
24	WE2360	KYANAMUKAAGA	CENTRAL 2	ROC	MASAKA	-0.50594	31.70416	KABONERA	
25	WE2456	KASESE TAXI PARK	WEST 2	KASESE	KASESE	0.186669	30.078602	TOWN CENTRE	
26	WE2836	LWANYI	CENTRAL 2	MASAKA	MASAKA	-0.3590929	31.7330652	MASAKA TOWN	
27	WE2888	KIGEZI	CENTRAL 2	ROC	GOMBA	0.27754	31.59685	KABULASOKE	

LOCATION_DETAILS.xlsx — sample of first 27 rows