mathematics

Article

A Fuzzy-XAI Framework for Customer Segmentation and Risk Detection: Integrating RFM, 2-Tuple Modeling, and Strategic Scoring

Gabriel Marín Díaz [1,2]

[1] Faculty of Statistics, Complutense University, Puerta de Hierro, 28040 Madrid, Spain; gmarin03@ucm.es
[2] Science and Aerospace Department, Universidad Europea de Madrid, Villaviciosa de Odón, 28670 Madrid, Spain

Abstract

This article presents an interpretable framework for customer segmentation and churn risk detection, integrating fuzzy clustering, explainable AI (XAI), and strategic scoring. The process begins with Fuzzy C-Means (FCM) applied to normalized RFM indicators (Recency, Frequency, Monetary), which were then mapped to a 2-tuple linguistic scale to enhance semantic interpretability. Cluster memberships and centroids were analyzed to identify distinct behavioral patterns. An XGBoost classifier was trained to validate the coherence of the fuzzy segments, while SHAP and LIME provided global and local explanations for the classification decisions. Following segmentation, an AHP-based strategic score was computed for each customer, using weights derived from pairwise comparisons reflecting organizational priorities. These scores were also translated into the 2-tuple domain, reinforcing interpretability. The model then identified customers at risk of disengagement, defined by a combination of low Recency, high Frequency and Monetary values, and a low AHP score. Based on Recency thresholds, customers are classified as Active, Latent, or Probable Churn. A second XGBoost model was applied to predict this risk level, with SHAP used to explain its predictive behavior. Overall, the proposed framework integrated fuzzy logic, semantic representation, and explainable AI to support actionable, transparent, and human-centered customer analytics.

Keywords: RFM model; fuzzy C-Means clustering; 2-tuple linguistic representation; analytic hierarchy process (AHP); customer segmentation; explainable AI (SHAP, LIME); churn prediction

MSC: 68T09; 68T37; 03E72; 62H30

check for updates

Academic Editor: Bingzhen Sun

Received: 22 May 2025
Revised: 18 June 2025
Accepted: 28 June 2025
Published: 30 June 2025

Citation: Marín Díaz, G. A Fuzzy-XAI Framework for Customer Segmentation and Risk Detection: Integrating RFM, 2-Tuple Modeling, and Strategic Scoring. *Mathematics* **2025**, *13*, 2141. https://doi.org/10.3390/math13132141

1. Introduction

In an increasingly competitive and dynamic market environment, understanding customer behavior is critical for ensuring sustained business growth, optimizing loyalty programs, and reducing churn. Among the most widely used frameworks in this domain is the RFM model (Recency, Frequency, Monetary), valued for its structured approach and direct connection to transactional data. RFM allows companies to quantify how recently and how often customers interact, as well as the financial value of those interactions, providing a solid basis for segmentation and targeting strategies [1].

Despite its popularity, however, the classical RFM approach presents several limitations. Traditional implementations often rely on rule-based scoring systems or discrete segmentations that fail to account for the vagueness and ambiguity inherent in human behavior. Customers frequently exhibit overlapping characteristics that challenge binary or crisp assignments. Moreover, the use of fixed thresholds and unweighted variables may obscure strategic business priorities or subtle behavioral nuances.

To overcome these limitations, recent research has explored more flexible and interpretable alternatives. Fuzzy logic, and in particular the Fuzzy C-Means (FCM) algorithm, has emerged as a relevant alternative to traditional crisp clustering methods like K-Means by allowing for partial membership between segments [2,3]. In parallel, the rise of explainable artificial intelligence (XAI), with techniques such as SHAP (SHapley Additive Explanations) and LIME (Local Interpretable Model-Agnostic Explanations) has brought much-needed transparency to complex machine learning models, especially in decision-critical business contexts [4,5].

At the same time, linguistic modeling techniques such as the 2-tuple representation have enabled more human-readable interpretations of numerical data [6]. These semantic layers bridge the gap between algorithmic computation and managerial insight, particularly when combined with scoring systems that reflect business priorities. Among such systems, the Analytic Hierarchy Process (AHP) is especially notable for its ability to incorporate expert judgments via pairwise comparisons, generating strategic weights for multi-criteria evaluation [7].

Despite the growing use of these techniques individually, there remains a lack of unified frameworks that combine fuzzy clustering, semantic modeling, explainable AI, and strategic scoring within a coherent pipeline. Existing approaches rarely offer a transparent and traceable path from raw behavioral data to actionable segmentation and prioritization of customers, especially in contexts where customer churn must be detected and mitigated proactively.

This paper addresses that gap by proposing a hybrid methodological framework that integrates the following:

- Fuzzy C-Means clustering for behavioral grouping;
- 2-tuple semantic modeling for transforming RFM variables into interpretable linguistic labels;
- SHAP and LIME for global and local model interpretability;
- AHP-based scoring for personalized customer ranking aligned with business strategy.

The proposed system allows for flexible and interpretable segmentation, as well as classifying customers according to their transactional patterns and explaining the reasons for each assignment. Furthermore, it helps to detect the risk of customer churn by combining fuzzy segmentation, behavioral markers (such as low recurrence with high spending), and strategic scoring. This research, thus, provides a novel, interpretable and operational framework for customer segmentation and retention that goes beyond static classification and enables transparent, data-driven decisions.

The contributions are structured as follows: Section 2 reviews the state of the art, analyzing prior work related to the methodologies adopted; Section 3 details the proposed methodological model, including fuzzy clustering, scoring, and explainability; Section 4 presents a real-world case study based on a pet product retail store in Spain; Section 5 discusses the results obtained and outlines future lines of research; and finally, Section 6 summarizes the main conclusions and implications of the study.

2. Related Work

Although each of the techniques integrated into the proposed framework has been previously studied in the context of customer analytics, existing approaches tend to treat them in isolation. Most works use RFM as a starting point for customer scoring or clustering, yet they often fail to capture strategic priorities and behavioral uncertainty, or provide interpretable justifications for their segmentations. In this section, the main lines of research related to each methodological component are reviewed: RFM scoring, AHP prioritization, fuzzy logic and 2-tuple representation, fuzzy clustering with FCM, and explainable AI, highlighting their strengths, limitations, and the gap addressed by this study.

2.1. RFM Modeling for Customer Analysis

The RFM model is one of the most widely used approaches to measuring customer behavior, mainly because of its simplicity and the fact that it is based on transactional data that is easy to extract through any management application [8]. The model quantifies customer value based on the recency of purchase (R), the frequency of purchase (F), and the amount spent (M), providing a solid basis for segmenting customers and promoting targeted marketing actions.

As shown in Figure 1, the number of studies indexed in Web of Science on this topic has grown significantly since 2019, with a particularly steep increase between 2020 and 2023. This evolution is supported by several of the most recent and frequently cited works included in Figure 1 [9–13], and reflects a broader shift toward data-driven customer strategies in the retail, banking, and services sectors [14].

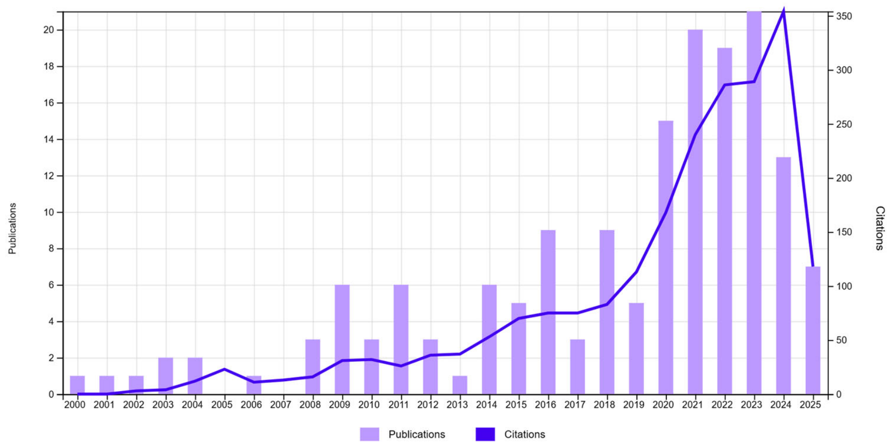

Figure 1. Publications (162) and citations (1478). TS = ("RFM" OR "Recency Frequency Monetary") AND TS = ("customer segmentation" OR "customer profiling").

Nevertheless, most classical applications of RFM rely on fixed scoring rules or percentile-based thresholds, which often fail to capture behavioral nuance or strategic business priorities. While these methods offer operational simplicity, they impose rigid segment boundaries and do not accommodate uncertainty or overlapping customer profiles. For example, two customers with similar spending may have very different purchase intervals or recency patterns, differences that traditional RFM scoring schemes may obscure.

The current literature reflects this dichotomy: a strong foundation of RFM-based segmentation approaches, but a growing demand for models that go beyond static classification. In response, recent works have started to explore hybrid strategies that combine RFM with machine learning techniques, fuzzy logic, or multi-criteria decision-making methods [15,16]. This paper follows that direction, using the RFM as an interpretable and behaviorally weighted input to a broader analytical pipeline.

2.2. Strategic Prioritization of RFM with AHP

The RFM model provides a solid basis for ranking customer behavior, but in most applications, the three components are usually given equal weights. In practice, organizations often prioritize one dimension over the others based on their strategic objectives.

To address this need for strategic differentiation, some studies have combined RFM analysis with the Analytic Hierarchy Process (AHP) [17]. AHP is a multi-criteria decision-making technique that translates expert judgment into weighted scores through pairwise comparisons [7]. When integrated with RFM, it allows for a customized score to be obtained for each client.

Although the number of publications on this topic remains limited, several applied studies have demonstrated the practical advantages of this hybrid approach in the retail, banking, and tourism sectors [18]. The integration of RFM and AHP has shown consistent academic interest over the last two decades. As illustrated in Figure 2, publications combining both techniques have appeared regularly since 2005, with a marked increase in citation impact from 2016 onward. This trend highlights the importance of complementing transactional behavioral models with expert-based prioritization frameworks, especially in industries where customer valuation must be closely aligned with organizational strategy. This pattern is supported by the most recent and influential studies in the field, as represented in references [19–22].

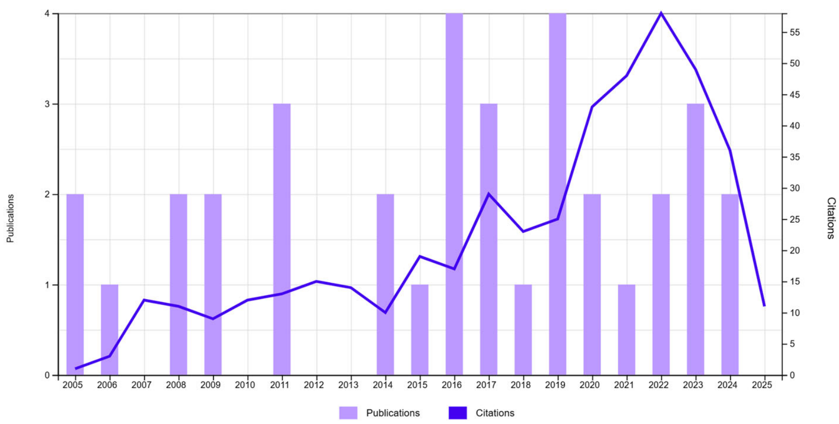

Figure 2. Publications (35) and citations (417). TS = ("RFM" OR "Recency Frequency Monetary") AND TS = ("Analytic Hierarchy Process" OR "AHP").

In this paper, AHP is used as a complementary scoring mechanism for each customer and applied after fuzzy segmentation and model interpretability. This allows for deriving individual strategic scores that reflect the business relevance of each customer within its respective segment.

2.3. Fuzzy Logic and 2-Tuple Linguistic Modeling

Fuzzy logic has long been used as a tool for modeling uncertainty and imprecision in real-world decision-making. Unlike binary logic, it allows variables to adopt degrees of membership in different categories, a feature particularly useful for modeling human judgement, which rarely conforms to rigid thresholds. In the context of customer analysis, fuzzy systems offer an elegant way to capture behavioral ambiguity, especially when customer profiles overlap or evolve gradually.

Within this paradigm, the fuzzy 2-tuple linguistic fuzzy model introduced by Herrera and Martinez has emerged as one of the most robust and interpretable approaches to semantic representation [23]. Instead of reducing fuzzy evaluations to numerical values, the 2-tuple model maintained the linguistic granularity of human input, while ensuring consistency and computational accuracy. Each value is expressed as a pair (s, α), where s is a linguistic term from a predefined set (e.g., Low, Medium, High), and α is a numerical translation representing the deviation from the central term.

Originally developed for linguistic decision making in multi-criteria problems, the 2-tuple model has gained wide application in areas such as educational assessment, recommendation systems, environmental analysis, and expert systems. As shown in Figure 3, research interest in this model has grown steadily over the last decade, with a peak in citations between 2018 and 2021, reflecting both its mathematical soundness and interpretive appeal. This evolution is supported by several of the most representative and frequently cited contributions, as reflected in references [24–27].

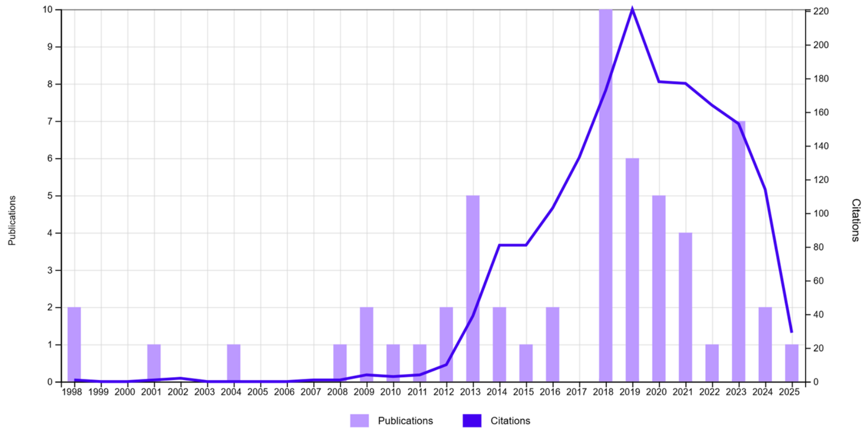

Figure 3. Publications (57) and citations (1419). TS = ("2-tuple" OR "two tuple") AND TS = ("fuzzy logic" OR "linguistic decision making").

The approach presented in this paper integrates 2-tuple modeling to express the normalized values of RFM indicators (Recency, Frequency, Monetary) in a five-level linguistic scale (Very Low, Low, Medium, High, Very High), allowing both customers and analysts to interpret patterns in terms closer to natural language.

2.4. Fuzzy C-Means Clustering for Customer Segmentation

Segmenting customers into groups that reflect behavioral patterns is a critical task for tailoring communication strategies, loyalty and product offerings to specific profiles. One of the most widely used techniques is the use of unsupervised algorithms, K-Means, due to its simplicity and low computational cost [28]. But it offers a rigid grouping system, which forces each customer to be classified into a single segment, oversimplifying reality, especially in contexts where behavior overlaps or evolves over time.

Fuzzy clustering, and in particular the Fuzzy C-Means (FCM) algorithm, offers a more nuanced alternative. Instead of assigning individuals exclusively to one group, FCM allows for degrees of membership in multiple clusters. This partial membership provides a truer reflection of customer variability and is especially important when behavioral patterns are ambiguous or transition between segments [29].

Interest in applying FCM to customer segmentation has grown steadily in recent years. As shown in Figure 4, the number of related studies has increased since the mid-2010s, accompanied by a significant rise in citations, particularly between 2020 and 2023. This trend is documented in some of the most influential and widely cited studies in the field, as listed in references [30–34].

This study uses FCM to identify fuzzy behavioral segments based on normalized RFM values. The objective is twofold: in addition to assigning customers to groups, it is essential to understand the strength and nature of their affinity with each cluster. These degrees of membership are interpreted using explainability tools, making the segmentation robust and accessible to decision-makers.

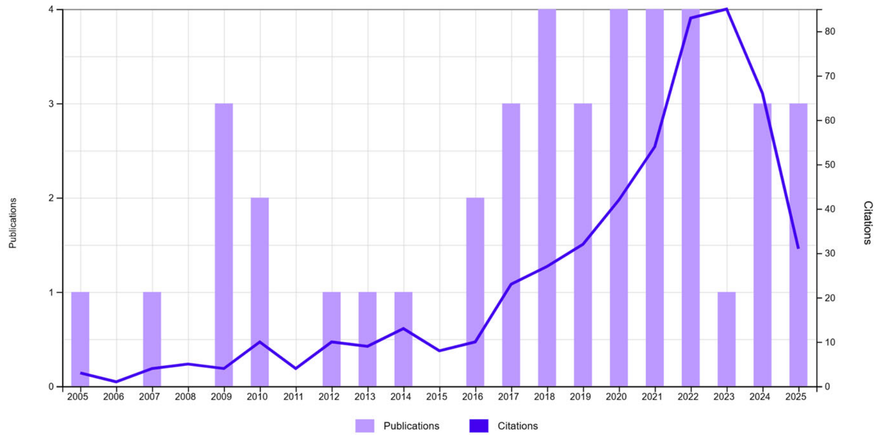

Figure 4. Publications (41) and citations (486). TS = ("Fuzzy C-Means" OR "fuzzy clustering") AND TS = ("customer segmentation").

2.5. Explainable AI (SHAP, LIME) in Customer Behavior and Churn

As machine learning models become increasingly common in customer analytics, the issue of interpretability has taken on a central role, especially in contexts where the results affect strategic decisions. This is particularly relevant in claim prediction scenarios, where simply identifying a customer as being at risk is insufficient unless the underlying reasons for that prediction are also understood [35].

Explainable Artificial Intelligence (XAI) techniques [36], such as SHAP and LIME, allow us to explore which variables contribute most to a particular classification and how they interact in each individual case. Unlike traditional feature importance metrics, SHAP and LIME provide explanations at both global and local levels, especially relevant in customer-centric environments [37].

In recent years, there has been a clear increase in the number of studies applying XAI to customer segmentation and churn detection [38–42], as shown in Figure 5. This growing body of work reflects a clear shift toward models that incorporate explainable AI (XAI) in customer segmentation and churn detection. Predictability, transparency, and justification are increasingly prioritized, as organizations demand accurate predictions and interpretable reasoning behind them. The ability to explain why a customer is assigned to a specific profile or identified as at risk of churn, has become a fundamental requirement in decision-making environments where data must inform action.

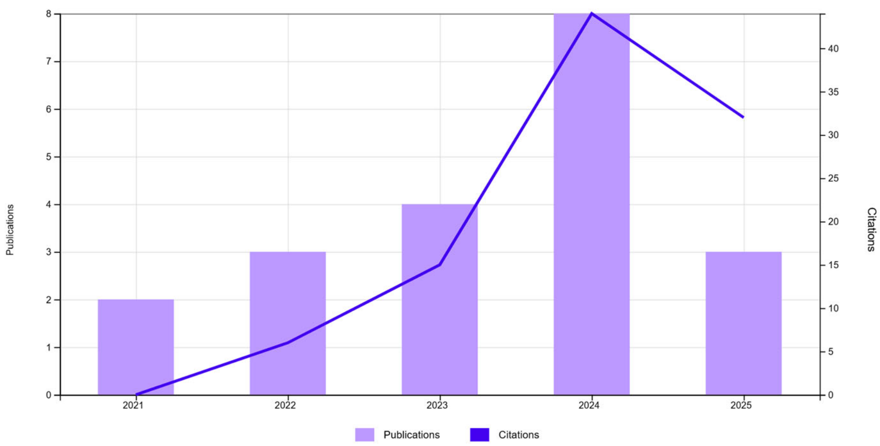

Figure 5. Publications (20) and citations (92). TS = ("Explainable AI" OR "XAI") AND TS = ("customer segmentation" OR "customer churn").

2.6. Toward a Unified Framework

Despite the increasing use of RFM models, fuzzy clustering, strategic scoring methods like AHP, and explainable AI tools in isolation, there is a clear lack of integrated approaches that bring all these components together into a coherent decision-making framework. Most existing studies focus either on segmentation or churn prediction, but rarely both, let alone through a process that incorporates fuzzy membership, semantic modeling, and transparent evaluation. No work published to date proposes a production line that performs customer segmentation using Fuzzy C-Means, interprets those profiles with SHAP and LIME, and complements the analysis with a strategic scoring layer using AHP. The present study addresses that gap by offering a methodological proposal that connects behavior, segmentation, interpretability, and business relevance, while also identifying and prioritizing potential churn cases through interpretable evidence.

To complement the state of the art and contextualize our methodological choices, we acknowledge several limitations inherent to the selected techniques. Fuzzy C-Means (FCM) is known for its sensitivity to initialization and dependence on the fuzzifier parameter, which can affect cluster stability [2]. AHP, while valuable for incorporating strategic priorities, relies on subjective expert judgment and may lead to inconsistency if not carefully validated [43]. SHAP and LIME, despite their interpretability benefits, can be computationally demanding and may face scalability challenges, especially in real-time or high-dimensional scenarios [44]. These limitations are considered in our design decisions and discussed in the context of practical implementation.

3. Methodology

The methodological approach combines clustering, semantic modeling, and explainable techniques to build a transparent and operational segmentation pipeline. Unlike traditional approaches that rely on rigid thresholds or black-box classifiers, the proposed pipeline combines interpretability and flexibility at each stage of the analysis.

The process begins with the extraction and normalization of RFM variables (Recency, Frequency, Monetary), which were used to characterize customer behavior. These variables were then mapped to a five-level linguistic scale using the 2-tuple fuzzy model, adding a semantic layer that enhances interpretability while maintaining mathematical precision.

Next, Fuzzy C-Means (FCM) clustering was applied to identify latent behavioral segments. Rather than assigning each customer to a single group, FCM generated degrees of membership across clusters, allowing for overlapping and hybrid profiles to emerge. The resulting centroids were then interpreted both numerically and linguistically, and representative customers were selected for further analysis.

To evaluate the internal consistency of the clustering, a supervised classifier (XGBoost) was trained to predict cluster membership based on RFM inputs. Explainable AI techniques, specifically SHAP and LIME, were used to interpret the model's decisions and assess the explanatory power of each variable in defining customer profiles.

Following this, the Analytic Hierarchy Process (AHP) was introduced to derive a global strategic score for each customer. Based on pairwise comparisons reflecting business priorities, AHP assigned different weights to the RFM components. These scores were then transformed into 2-tuple expressions and compared to the clustering results, allowing for alignment or contrast between fuzzy groupings and strategic valuation.

Finally, the model was extended to identify potential churn cases. Customers with low Recency, high Frequency and Monetary value, and low AHP scores were flagged as at risk. A classification model was built to validate this behavior, and explainability tools were applied once again to support actionable insights.

This section details each of these steps, from data preparation to modeling and interpretability, outlining how they combine into an integrated, explainable decision support system (Figure 6).

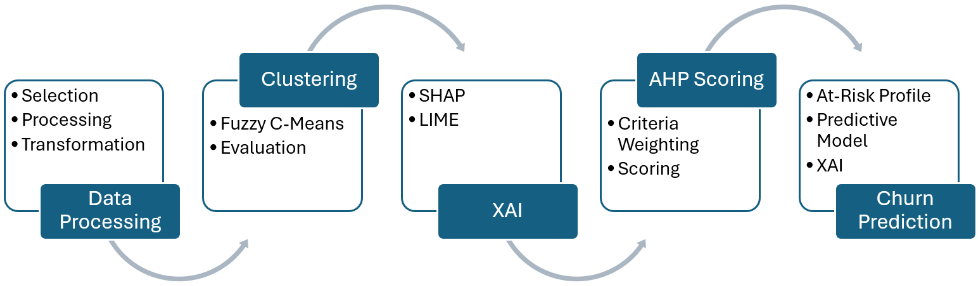

Figure 6. Methodological workflow for customer segmentation and churn prediction.

3.1. Data Processing

The data processing phase served as the foundation of the proposed methodology. Its objective was to transform raw transactional information into structured, normalized variables that are both interpretable and analytically useful. This step was anchored in the RFM model.

From the transaction dataset, three behavioral indicators were calculated for each customer:

Variable Selection

- Recency (R): the number of days since the customer's most recent purchase;
- Frequency (F): the total number of purchases made within a defined period;
- Monetary (M): the cumulative monetary value of all purchases.

For each customer i, the behavior vector was constructed:

$$r_i = (R_i, F_i, M_i) \tag{1}$$

Normalization and Scaling

Since RFM variables operate on different scales, min.–max. normalization was applied to ensure comparability between dimensions:

$$x_i^{(j)} = \frac{x_i^{(j)} - min\left(x^{(j)}\right)}{max\left(x^{(j)}\right) - min\left(x^{(j)}\right)},\ \ j \in \{R, F, M\} \tag{2}$$

This yielded a normalized behavioral vector:

$$r_{\mathrm{i}}{}^{(\mathrm{norm})} = \left(Percent_{Recency_{\mathrm{i}}},\ Percent_{Frequency_{\mathrm{i}}},\ Percent_{Monetary_{\mathrm{i}}}\right) \tag{3}$$

Normalization prepares the data for fuzzy analysis and semantic interpretation in the subsequent stages.

Linguistic Transformation Using 2-Tuple Model

To introduce an interpretable semantic layer, each normalized value was translated into a five-level linguistic scale:

$$S = \{Very\ Low, Low, Medium, High, Very\ High\} \tag{4}$$

Each normalized score $x_i^{(j)} \in [0, 1]$ was mapped to a linguistic term using the 2-tuple fuzzy linguistic representation model, which expresses the value as

$$(s, \alpha), \quad s \in S, \quad \alpha \in \left[-\frac{0.5}{g}, \frac{0.5}{g}\right) \tag{5}$$

where $g = |S| - 1 = 4$, and α is the symbolic translation representing the deviation from the closest linguistic label.

The choice of a five-level linguistic scale was based on literature, suggesting that this granularity offers a good trade-off between interpretability and expressiveness in managerial contexts [6,23].

This transformation preserves the mathematical fidelity of the data while adding interpretability that is essential for expert-based decision-making and customer analysis.

3.2. Clustering

3.2.1. Fuzzy Clustering

Once the normalized RFM variables were prepared and semantically mapped, the next step was to identify latent customer segments based on behavioral similarity. To this end, we applied Fuzzy C-Means (FCM) clustering, an unsupervised learning technique that allows for each data point to belong to multiple clusters with varying degrees of membership [16].

FCM generates a membership matrix that reflects the partial affiliation of each customer to each cluster. This feature is especially valuable in customer analytics, where behavioral boundaries are often diffused.

The FCM algorithm minimizes the following objective function:

$$J_m = \sum_{i=1}^{N} \sum_{k=1}^{c} u_{ik}^m \cdot \mid x_i - c_k \mid^2 \tag{6}$$

where

- N is the number of customers;
- c is the number of clusters;
- $x_i \in R^n$ is the vector of normalized RFM variables for customer i;
- c_k is the centroid of cluster k;
- $u_{ik} \in [0, 1]$ is the membership degree of customer i to cluster k;
- $m > 1$ is the fuzzification coefficient, typically set to $m = 2$.

The fuzzification coefficient was set to $m = 2$, a common and empirically stable value in FCM applications. This choice was validated by testing values between 1.8 and 2.5, and observing consistent cluster structures and silhouette scores.

The degrees of membership were updated iteratively using

$$u_{ik} = \left(\sum_{j=1}^{c} \left(\frac{\mid x_i - c_k \mid}{\mid x_i - c_j \mid} \right)^{\frac{2}{m-1}} \right)^{-1} \tag{7}$$

Each cluster centroid was recomputed as a weighted average

$$c_k = \frac{\sum_{i=1}^{N} u_{ik}^m \cdot x_i}{\sum_{i=1}^{N} u_{ik}^m} \tag{8}$$

The algorithm iterates between the membership and centroid updates until convergence, typically measured by the change in the objective function or the centroids being below a given threshold.

The result of the FCM algorithm is

- A set of fuzzy centroids $\{c_1, \ldots, c_c\}$;
- A membership matrix $U = [u_{ik}] \in R^{N \times c}$.

These values were retained for the interpretability phase, where both the centroids and customer memberships were analyzed linguistically. Customers were assigned a dominant cluster (the one with the highest membership), but the full fuzzy profile was preserved for later semantic analysis.

3.2.2. Evaluation Clustering Model

Selecting the appropriate number of clusters and validating their structure are essential to ensure that the segmentation process reflects meaningful behavioral patterns. In the context of fuzzy clustering, evaluation involves both structural and interpretability aspects. In this study, the internal validation index was applied, Silhouette Coefficient, to assess the quality and coherence of the fuzzy partition.

The Silhouette Coefficient evaluates how similar an object is to its own cluster compared to other clusters [45]. For each customer i, the silhouette score is defined as

$$s(i) = \frac{b(i) - a(i)}{max\{a(i), b(i)\}} \tag{9}$$

where

- $a(i)$ is the average distance between i and all other points in the same cluster;
- $b(i)$ is the minimum average distance between i and all points in the nearest different cluster.

The overall silhouette score was the mean of all individual scores:

$$S = \frac{1}{N} \sum_{i=1}^{N} s(i) \tag{10}$$

The silhouette score ranges from -1 to 1. Higher values indicate better-defined, more coherent clusters.

Although the silhouette score was originally designed for crisp clustering, it can be applied to fuzzy clustering by assigning each point to its most dominant cluster (i.e., the one with the highest membership degree).

To determine the optimal number of clusters c, the FCM algorithm was executed for a range of values (e.g., $c = 2$ to $c = 6$), and both validation indices were computed. The best value of c was selected based on the following criteria:

- Maximize the Silhouette Coefficient;
- Ensure interpretability: the chosen number of clusters must allow for meaningful segmentation that can be explained in the business context.

In the case study developed in Section 4, both K-Means and FCM were tested for several values of c, and the validation metrics were analyzed. The results confirmed that FCM with $c = 4$ offered a good balance between intra-cluster compactness and inter-cluster separation, while also aligning well with the behavioral diversity observed in the data.

The fuzzy membership degrees and centroids obtained were retained for the next stage, where linguistic interpretation and explainability were applied.

3.3. Interpretability

Understanding the reasoning behind the assignment of customers to behavioral clusters is essential for ensuring transparency and enabling actionable strategies. While Fuzzy C-Means (FCM) provides degrees of membership, it does not offer a clear explanation of why a given customer exhibits affinity to a particular segment. To address this limitation, Explainable Artificial Intelligence (XAI) tools, SHAP and LIME, were incorporated to interpret the clustering structure from a supervised modeling perspective.

To bridge unsupervised clustering with explainability techniques, we trained a supervised classifier, XGBoost, to predict the dominant cluster (i.e., the cluster with the highest membership degree) using the normalized RFM variables as features, Equations (2) and (3).

The model's high accuracy (as demonstrated in Section 4) confirmed that the clusters identified by FCM were reproducible and separable in the RFM feature space. Once trained, this classifier served as the basis for explainability.

3.3.1. SHAP (SHapley Additive Explanations)

SHAP values are computed to assess the global and local contribution of each variable to the classification output [4]. For each prediction, SHAP assigns an additive value to each feature, reflecting how much it pushes the prediction toward or away from a given cluster.

Formally, for a given prediction $f(x)$, the SHAP decomposition is

$$f(x) = \phi_0 + \sum_{j=1}^{M} \phi_j \tag{11}$$

where

- ϕ_0 is the model's expected value;
- ϕ_j is the SHAP value for feature j.

These values are visualized using summary plots and force plots, which provide insight into the most influential features across the entire dataset or for a specific customer.

3.3.2. LIME (Local Interpretable Model-Agnostic Explanations)

LIME complements SHAP by generating local surrogate models that approximate the classifier's behavior in the vicinity of a given instance [5]. For each customer x_i, LIME fits a simple, interpretable model g^* that mimics the complex classifier f in a neighborhood $\mathcal{N}(x_i)$:

$$g^* = \underset{g \in G}{argmin}[\mathcal{L}(f, g, \pi_{x_i}) + \Omega(g)] \tag{12}$$

where

- $\mathcal{L}$ is a loss function measuring the fidelity of $g\ to\ f$;
- π_{x_i} defines the neighborhood weighting;
- $\Omega(g)$ penalizes complexity.

This approach helps decision-makers understand why a specific customer has been assigned to a cluster, based on simple, human-readable rules derived from local approximations.

3.4. AHP Scoring

While fuzzy clustering enables the identification of behavioral patterns among customers, it does not account for the strategic priorities of the business. Two customers with similar purchase behaviors may not have the same relevance from a managerial perspective. To address this limitation, we incorporated the Analytic Hierarchy Process (AHP) as a complementary scoring mechanism that assigns an individual strategic value to each customer based on weighted RFM criteria.

AHP is a multi-criteria decision-making technique introduced by Saaty [7], which relies on expert judgment to construct a pairwise comparison matrix. In this context, the goal is to assign relative importance to the three RFM dimensions: *Recency, Frequency, Monetary*.

A decision maker provides pairwise comparisons of a_{ij} for all combinations of criteria $i,\ j$, where

$$A = \begin{bmatrix} 1 & a_{12} & a_{13} \\ 1/a_{12} & 1 & a_{23} \\ 1/a_{13} & 1/a_{23} & 1 \end{bmatrix} \tag{13}$$

The weight vector $w = (w_R, w_F, w_M)$ is obtained by calculating the principal right eigenvector of matrix A:

$$A \cdot w = \lambda_{max} \cdot w \tag{14}$$

where λ_{max} is the maximum eigenvalue of A.

The consistency index (CI) and consistency ratio (CR) are computed to ensure the reliability of the judgments:

$$CI = \frac{\lambda_{max} - n}{n - 1},\ \ CR = \frac{CI}{RI} \tag{15}$$

where

- $n = 3$ is the number of criteria;
- RI is the random index, e.g., $RI = 0.58$ for $n = 3$.

A consistency ratio $CR < 0.1$ is generally considered acceptable.

Once the weight vector w is computed, each customer's RFM vector is scored according to

$$AHP\_Score_i = w_R \cdot x_i^{(R)} + w_F \cdot x_i^{(F)} + w_M \cdot x_i^{(M)} \tag{16}$$

where $x_i^{(j)} \in [0, 1]$ are the normalized RFM values for customer i.

This results in a strategic score in the range $[0, 1]$, where higher values indicate customers that better align with the company's defined priorities.

To improve interpretability, the AHP score is also mapped to the five-level 2-tuple linguistic scale:

$$Score_i \to (s, \alpha),\ s \in \{VL, L, M, H, VH\} \tag{17}$$

where

- s is the nearest linguistic label to the score;
- α is the symbolic translation indicating the deviation from the central value of the label.

This semantic representation allows managers to easily interpret the strategic importance of each customer beyond numerical scoring.

By comparing the AHP score with the FCM cluster assignment, we can assess the following:

- Whether high-value customers (strategically) fall within high-intensity behavioral clusters;
- Potential mismatches between observed behavior and business-defined importance.

This dual-layer perspective supports more nuanced decisions, allowing for both profile-based actions (via clustering) and priority-based interventions (via AHP scoring).

3.5. Explainable Churn Prediction

The final stage of the proposed framework addresses a common and critical challenge in customer management: identifying clients at risk. While segmentation and scoring provide a multidimensional view of customer behavior and strategic value, they must ultimately support actionable insights. This section outlines how the outputs of the previous phases, fuzzy clustering, AHP scoring, and semantic interpretation can be leveraged to detect and explain churn risk.

Based on empirical analysis and business logic, a customer is at potential risk if they meet the following combined conditions:

- Low Recency (i.e., the customer has not purchased recently);
- High Frequency and Monetary values (indicating prior engagement and economic value);
- Low AHP Score (reflecting strategic devaluation based on company-defined priorities).

To facilitate classification, Recency is used to label each customer as

$$Recency\_Status_i = \begin{cases} Active,\ if\ x_i^{(R)} > 0.75 \\ Latent,\ if\ 0.4 < x_i^{(R)} \le 0.75 \\ Likely\ Churn,\ if\ x_i^{(R)} \le 0.4 \end{cases} \tag{18}$$

When combined with high F and M values, and a low AHP score, the segment *Likely Churn* identifies a profile of dormant customers who were previously valuable but are currently disengaged.

To validate and automate the detection process, an XGBoost classifier is trained using the RFM variables as features and the churn risk label as the target. The classifier achieved high performance, confirming that behavioral features are sufficient to distinguish at-risk customers.

This step closes the loop of the proposed methodology by moving from descriptive segmentation to predictive intervention, supporting human-centered decisions through both semantic and statistical justification, and enabling proactive actions such as personalized offers, loyalty incentives, or follow-up communication with high-risk customers.

4. Case Study and Results

4.1. Dataset Description

The dataset used in this study originates from a retail store specialized in pet-related products operating in Spain. The data reflect real customer interactions, encompassing transactions over a complete one-year period. The dataset covers a complete calendar year, from 1 January to 31 December 2023, allowing for seasonally representative customer behavior analysis. To ensure anonymity and data protection compliance, all personal identifiers were removed, and the data were processed solely at the behavioral level using transactional variables.

The initial database consisted of 40,677 customer records for a total of 1,217,148 transactions. After applying cleaning procedures, such as removing incomplete entries, standardizing date formats, and filtering customers with at least two purchases to ensure behavioral consistency, a final sample of 33,523 unique customer profiles was retained for analysis.

Table 1 summarizes the descriptive statistics of the three RFM variables used for profiling. These variables were normalized to a [0, 1] range and semantically transformed using a five-level 2-tuple scale to allow for interpretability and alignment with the fuzzy modeling approach employed in subsequent stages.

Table 1. Descriptive statistics of RFM variables for the 33,523 customer profiles.

Statistics	Monetary	Frequency	Recency
Count	33,523.00	33,523.00	33,523.00
Mean	234.18	10.07	143.19
Std	274.64	9.63	106.64
Min	−576.80	1.00	0.00
25%	41.85	2.00	48.00
50%	126.57	6.00	122.00
75%	324.35	15.00	230.00
Max	1467.94	40.00	358.00

Table 2 shows a representative sample of five customer profiles after the normalization and semantic transformation of RFM variables using the 2-tuple model. For each case, the normalized percentage values of Recency, Frequency, and Monetary (ranging in [0, 1]) are accompanied by their corresponding fuzzy linguistic representations. These expressions combine a linguistic label (e.g., *VL*—Very Low, *VH*—Very High) and a symbolic deviation, capturing the nuance within each interval. This dual representation enhances interpretability, allowing business analysts to associate behavioral patterns with meaningful categories.

Table 2. Normalized and fuzzy (2-tuple) representations of RFM model.

Customer	Percent_Recency	Percent_Frequency	Percent_Monetary	R_2tuple	F_2tuple	M_2tuple
10,089	0.333786	0.954628	0.890732	(L, 0.084)	(VH, −0.045)	(VH, −0.109)
10,588	0.488799	0.076873	0.031859	(M, −0.011)	(VL, 0.077)	(VL, 0.032)
11,034	0.075426	0.904245	0.889747	(VL, 0.075)	(VH, −0.096)	(VH, −0.11)
11,574	0.373087	0.207604	0.375235	(L, 0.123)	(L, −0.042)	(M, −0.125)
11,723	0.583883	0.076873	0.145706	(M, 0.084)	(VL, 0.077)	(L, −0.104)

For instance, customer 10089 is identified as having a *low* Recency (i.e., purchased relatively recently), but *very high* Frequency and Monetary value, placing them among the most active and valuable segments. In contrast, customer 10588 shows *very low* Frequency and spending, combined with a *medium* Recency, suggesting a low-engagement or potentially dormant profile. These contrasts highlight the descriptive power of fuzzy semantics over purely numerical classification.

Figure 7 reveals a strong positive correlation between Frequency and Monetary value ($r = 0.75$), indicating that more frequent customers tend to generate higher revenue. Conversely, Recency shows weak negative correlations with both Frequency ($r = -0.25$) and Monetary value ($r = -0.23$). This inverse relationship, based on the convention that lower Recency indicates more recent activity, suggests that customers who purchased more recently tend to be more engaged and more valuable.

These relationships help contextualize both the clustering and churn prediction tasks, particularly in identifying customers with high historical value but recent disengagement, which are critical cases in retention strategies.

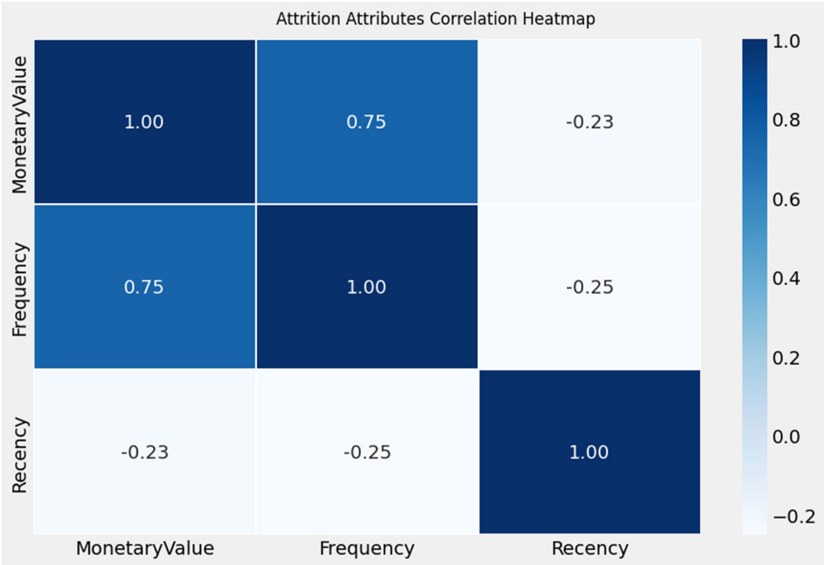

Figure 7. Correlation matrix; RFM model.

4.2. Cluster Profiles

To determine the most appropriate number of customer segments for the fuzzy clustering stage, we applied two complementary techniques: the elbow method [46] and the Silhouette coefficient [45]. The elbow method captures diminishing returns in compactness, while the Silhouette coefficient quantifies cohesion and separation. Using both helps balance interpretability and structural validity, especially in fuzzy clustering where overlap between segments is expected.

First computed the within-cluster sum of squared errors (WCSSs) for values of k ranging from 1 to 19. The results were plotted to identify the point where the marginal gain in clustering quality significantly drops, the so-called *elbow point*. As shown in Figure 8, there is a marked inflection at $k = 4$, after which the improvement in intra-cluster compactness becomes less significant. This suggests that four clusters may provide a reasonable balance between under- and over-segmentation.

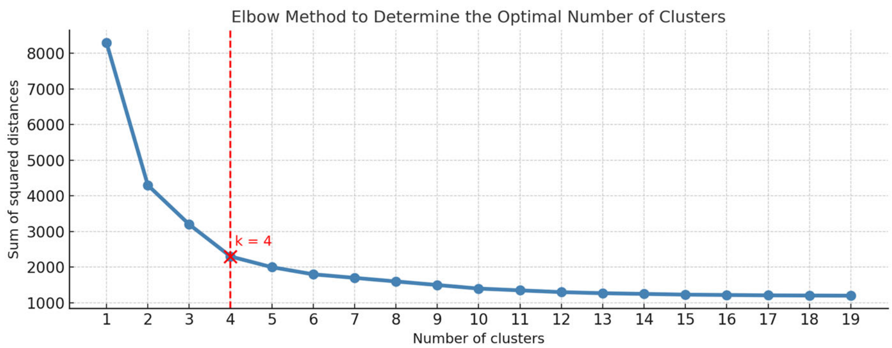

Figure 8. Number of clusters; elbow method.

As a complementary approach to the elbow method, Figure 9 presents the average Silhouette Coefficient obtained for different values of k (from 2 to 6). This metric quantifies how well each data point fits within its assigned cluster compared to other clusters. Higher values indicate more cohesive and better-separated clusters.

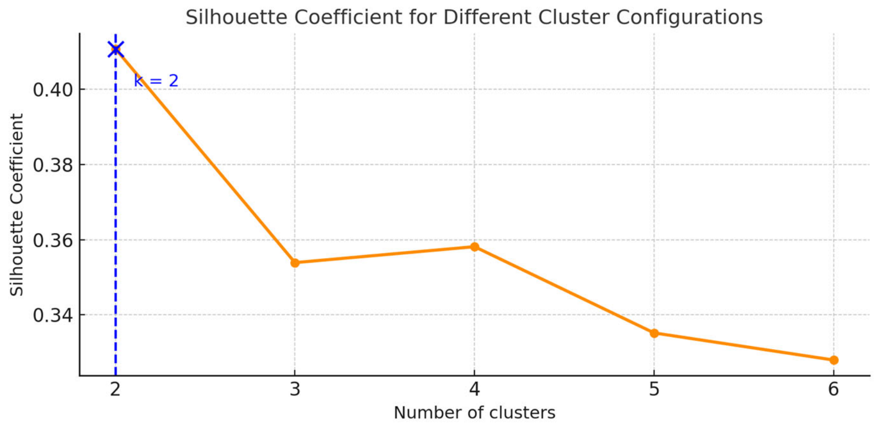

Figure 9. Number of clusters; Silhouette method.

Just like the elbow method, this technique helps identify the optimal number of clusters. While the highest Silhouette value was observed for $k = 2$, further interpretation suggests that this configuration may not capture the diversity of customer profiles required for effective segmentation strategies.

After reviewing both configurations, the elbow method with $k = 4$ provides a better balance between cohesion, interpretability and strategic segmentation of customer typologies. The four-cluster model captures key patterns such as inactive users, stable buyers and high-value loyal customers, allowing for more targeted and effective marketing actions.

Once the Fuzzy C-Means algorithm was applied, each customer was assigned a degree of membership to each cluster, rather than a single hard label. This allows us to identify overlapping behaviors and potential transitions between segments over time.

Table 3 presents the centroid-aligned cluster assignment along with the fuzzy membership degrees for the first five customers in the dataset. These values reflect the strength of association with each cluster, offering a nuanced interpretation of customer positioning across segments.

Table 3. Fuzzy cluster membership degrees for the first five customers.

Customer	Cluster_FCM	%Cluster_0	%Cluster_1	%Cluster_2	%Cluster_3
10,089	3	0.042802	0.025435	0.121633	0.810130
10,588	1	0.364265	0.505453	0.078514	0.051768
11,034	3	0.047339	0.023567	0.071197	0.857897
11,574	0	0.837035	0.089516	0.043208	0.030241
11,723	1	0.221777	0.687498	0.057584	0.033141

Such fuzzy partitions are particularly useful for identifying borderline profiles, which may require tailored interventions or more flexible marketing strategies than strictly defined groups.

The centroids obtained from the Fuzzy C-Means algorithm summarize the average behavior of each cluster across the RFM dimensions. To enhance interpretability, we translated these numeric values into 2-tuple linguistic representations, which provide both a qualitative descriptor (e.g., *low*, *high*) and a symbolic translation indicating the deviation from the central value of the linguistic term.

Table 4 displays the normalized RFM centroid values for each of the four clusters, alongside their corresponding 2-tuple approximations. This hybrid representation facilitates a more intuitive understanding of customer segments for decision-makers without sacrificing numerical accuracy.

Table 4. Fuzzy cluster centroids.

#Customers	Cluster	CRecency	CFrequency	CMonetary	CR_2tuple	CF_2tuple	CM_2tuple
7937	0	0.304069	0.316502	0.313132	(L, 0.054)	(L, 0.067)	(L, 0.063)
8574	1	0.773419	0.195468	0.198944	(H, 0.023)	(L, −0.055)	(L, −0.051)
8232	2	0.691267	0.666233	0.667036	(H, −0.059)	(H, −0.084)	(H, −0.083)
8780	3	0.230602	0.808758	0.805514	(L, −0.019)	(H, 0.059)	(H, 0.056)

The following strategic profiles were derived from the centroids:

- Cluster 0—Passive or Occasional Buyers

  RFM profile: Moderate Recency, low Frequency, low Monetary value.
  These customers purchase sporadically and show limited engagement.
  Strategic action: Design reactivation campaigns or low-threshold offers to increase visit frequency.

- Cluster 1—Recent but Uncommitted

  RFM profile: Very high Recency, low Frequency, low Monetary value.
  Likely to be new or trial customers with minimal follow-up activity.
  Strategic action: Implement onboarding and retention strategies, such as personalized product suggestions and follow-up messaging.

- Cluster 2—High-Value Loyalists

  RFM profile: high Frequency, Monetary, and Recency.
  These are the most engaged and profitable customers.
  Strategic action: Prioritize through loyalty programs, exclusive offers, or early access initiatives to reinforce brand attachment.

- Cluster 3—Dormant Champions

  RFM profile: Low Recency, very high Frequency and spending.
  Previously high-value customers who may be at risk of churn.
  Strategic action: Launch targeted win-back campaigns using previous purchase history or personalized incentives to restore engagement.

This multidimensional and fuzzy approach enables the identification of both high-value segments and strategically fragile groups, allowing for differentiated, data-driven interventions in customer lifecycle management.

4.3. Explanatory Insights with SHAP and LIME

4.3.1. Predictive Validation of Fuzzy Segments with XGBoost

To assess the coherence and learnability of the fuzzy clustering results, we trained a supervised classification model using the well-established XGBoost algorithm [47]. The model was configured to predict the cluster assignment (Cluster_FCM) based on the three normalized RFM dimensions: Recency, Frequency, and Monetary Value.

XGBoost was selected for its strong predictive performance, robustness to multicollinearity, and suitability for non-linear patterns. Its native compatibility with SHAP and LIME also facilitates integration into explainability-focused pipelines, making it ideal for validating clusters and supporting transparent interpretation.

The dataset was split into training (80%) and testing (20%) subsets to evaluate generalization performance, with hyperparameters set as follows: learning rate = 0.1, maximum depth = 3, and number of estimators = 100. The classifier was trained with a multiclass objective (multi: softprob) suitable for the four-cluster configuration.

The resulting model achieved an accuracy of 99.09% on the test set, indicating a remarkably strong alignment between the RFM features and the learned cluster labels. Table 5 presents the classification metrics for each segment.

Table 5. Precision, recall, and F1-score for each fuzzy cluster predicted by the XGBoost classifier.

Cluster	Precision	Recall	F1-Score	Support
0	0.99	0.99	0.99	1544
1	0.99	0.99	0.99	1723
2	0.99	0.99	0.99	1679
3	0.99	0.99	0.99	1759

The macro and weighted averages for precision, recall, and F1-score were also 0.99, confirming that the model performs consistently across all clusters without bias toward majority classes.

These results confirm that the fuzzy clusters are linearly separable with high reliability using only three RFM dimensions. This further supports the internal validity of the clustering process and justifies its use for downstream interpretability using SHAP and LIME.

4.3.2. Global Feature Importance with SHAP

To enhance the interpretability of the supervised XGBoost model trained to predict cluster membership, we applied SHAP (SHapley Additive exPlanations) to evaluate the average contribution of each feature to the model's output. The resulting summary plot (Figure 10) displays the mean absolute SHAP value for each RFM feature across the four predicted clusters.

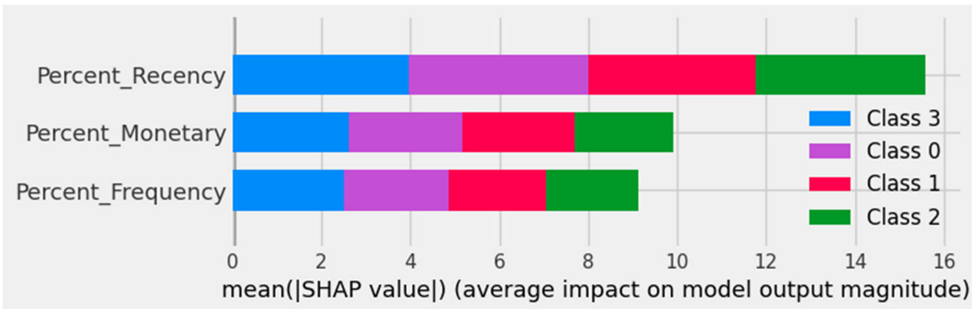

Figure 10. Global SHAP feature contributions per cluster.

SHAP values offer an additive decomposition of the model's decision process, enabling both global feature importance and per-class attribution. In this case, the color bars represent the cumulative impact of each variable on the prediction for each cluster (class), revealing the role that Recency, Frequency, and Monetary value play in distinguishing between customer profiles.

SHAP-Based Interpretation by Cluster:

- Cluster 0 (Passive or Occasional Buyers): The model relies most on Recency, followed by Frequency. This is coherent with the cluster's centroid profile (moderate Recency, low Frequency and value), confirming that temporal separation plays a key role in identifying these customers.
- Cluster 1 (Recent but Uncommitted): Recency and Monetary value are the dominant features. This mirrors the centroid structure where customers have very high Recency but low economic engagement, identifying them as new or trial users.

- Cluster 2 (High-Value Loyalists): All three RFM features are highly influential, especially Monetary value and Recency. The centroid of this cluster shows high frequency and spending, aligning with SHAP's attribution and confirming this as the most strategically important group.
- Cluster 3 (Dormant Champions): The dominant factor is again Recency, but from the opposite end. These are customers who used to be very active but have not interacted recently. The SHAP pattern supports the fuzzy segmentation and highlights the need for reactivation strategies.

Overall, the SHAP explanations closely match the fuzzy cluster centroids, confirming that the XGBoost model has captured the underlying behavioral patterns embedded in the unsupervised segmentation. This coherence strengthens the validity of the combined approach and provides transparent justification for subsequent strategic decisions.

4.3.3. Local Interpretability with LIME

To complement the global interpretation of feature contributions provided by SHAP, we used LIME to generate local explanations for a representative customer from each cluster. The goal was to validate the classifier's decisions at the individual level, and to ensure that predictions are not only accurate, but also semantically aligned with the RFM-based profiles discovered during the fuzzy clustering phase.

Figures 11–14 shows the prediction probabilities and local explanation rules produced by LIME for one customer in each cluster. These explanations are based on a linear approximation of the model's decision boundaries around each instance, highlighting the most influential RFM features and their thresholds.

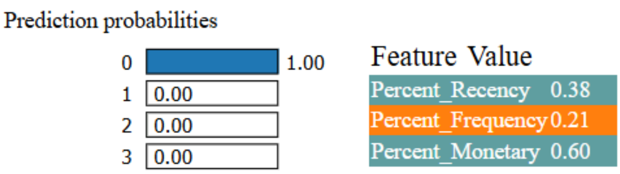

Figure 11. Local LIME explanations for fuzzy cluster 0.

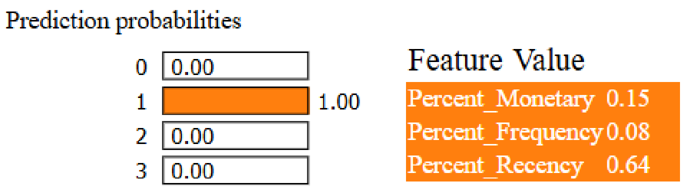

Figure 12. Local LIME explanations for fuzzy cluster 1.

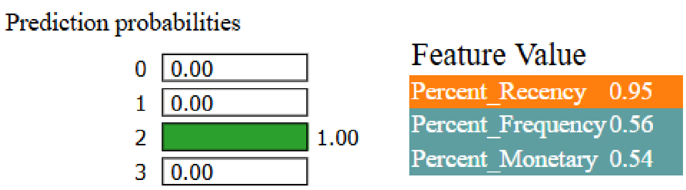

Figure 13. Local LIME explanations for fuzzy cluster 2.

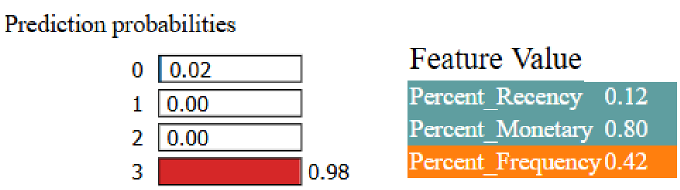

Figure 14. Local LIME explanations for fuzzy cluster 3.

The following interpretation summarizes each case:

- Cluster 0—Occasional Buyer: This customer had moderate Recency (0.38) and low Frequency (0.21). LIME highlighted both as strong indicators for Cluster 0, aligning with the profile of an infrequent but not dormant customer.
- Cluster 1—Recent but Low Engagement: The model classified this customer with 100% confidence into Cluster 1. LIME confirmed that this was due to high Recency (0.64) but very low Frequency (0.08) and Monetary value (0.15). This pattern is typical of new or uncommitted buyers.
- Cluster 2—High-Value Loyalist: LIME identified Recency > 0.75, Frequency > 0.47, and Monetary value > 0.50 as the key decision rules. The input values supported this, confirming the profile of a consistently active and profitable customer.
- Cluster 3—Dormant Champion: With a Recency of 0.12 and high Monetary value (0.80), LIME attributed the classification to past value but recent inactivity, perfectly matching the definition of a dormant high spender.

These local explanations validate the logic of the fuzzy segmentation, showing that the supervised model accurately replicates the assignments while relying on coherent and transparent feature combinations. They also enable individual-level interventions, such as reactivation offers or loyalty incentives, tailored to each customer's behavioral profile.

4.4. Strategic Scoring with AHP

To complement the fuzzy segmentation and enhance strategic decision-making, we computed an AHP-based scoring system for each customer. The aim was to quantify the relative strategic value of customers based on their RFM profile, integrating managerial priorities into the assessment.

The Analytic Hierarchy Process (AHP) was used to define the relative importance of the three RFM dimensions. Based on expert judgment, the following preference order was established: Recency > Frequency > Monetary value. The pairwise comparison matrix was developed in collaboration with a panel of three internal marketing and analytics professionals. Their consensus reflects the company's strategic emphasis on recent customer engagement, favoring Recency over Frequency and Monetary in the AHP prioritization, as shown in Equation (13) in Section 3.4:

$$\begin{pmatrix} 1 & 2 & 3 \\ 1/2 & 1 & 2 \\ 1/3 & 1/2 & 1 \end{pmatrix}$$

From this matrix, the normalized principal eigenvector was calculated, yielding the AHP weights: Recency (R): 0.540; Frequency (F): 0.297; Monetary value (M): 0.163. Consistency Ratio $CR = 1.0\%$.

These weights reflect a strategic emphasis on Recency of interaction, followed by engagement frequency, and lastly by transactional value.

Each customer's AHP score was computed as the weighted sum of their normalized RFM values:

$$AHP\ Score_i = 0.540 \cdot R_i + 0.297 \cdot F_i + 0.163 \cdot M_i$$

To improve semantic interpretability, each score was also translated into a 2-tuple linguistic representation, providing a qualitative label and a symbolic translation component. This hybrid representation facilitates clearer communication with non-technical stakeholders. It is also worth noting that AHP-based outcomes may vary depending on the selected criteria weights, indicating that sensitivity analysis could be a valuable direction for future research.

Table 6 presents the AHP scores and their 2-tuple equivalents for the five most representative customers in each cluster.

Table 6. AHP scores and 2-tuple, five customers in each cluster.

Customer	Cluster	AHP_Score	AHP_Score_2tuple	R_2tuple	F_2tuple	M_2tuple
CL2200297005	0	0.537385	(M, 0.037)	(M, 0.089)	(H, −0.122)	(L, 0.004)
CL0000663069	0	0.533952	(M, 0.034)	(M, 0.089)	(H, −0.094)	(L, −0.067)
CL2100157511	0	0.530266	(M, 0.03)	(M, 0.101)	(M, 0.058)	(L, 0.046)
CL2200212972	0	0.529033	(M, 0.029)	(M, 0.095)	(M, 0.058)	(L, 0.061)
CL2200389725	0	0.527643	(M, 0.028)	(M, 0.114)	(M, 0.019)	(L, 0.06)
CL0000018850	1	0.756649	(H, 0.007)	(VH, −0.008)	(M, 0.058)	(L, 0.092)
CL2200247178	1	0.755446	(H, 0.005)	(VH, −0.002)	(H, −0.122)	(L, −0.061)
CL2200013634	1	0.755076	(H, 0.005)	(VH, −0.005)	(M, 0.058)	(L, 0.073)
CL2100359099	1	0.754776	(H, 0.005)	(VH, −0.002)	(H, −0.122)	(L, −0.065)
CL2000982282	1	0.753110	(H, 0.003)	(VH, −0.012)	(M, 0.058)	(L, 0.084)
CL2200483478	2	1.000000	VH	(VH, −0.012)	(VH, −0.002)	(VH, −0.008)
CL2200041673	2	0.994955	(VH, −0.005)	(VH, −0.008)	(VH, −0.021)	(VH, −0.016)
CL2300014703	2	0.989778	(VH, −0.01)	(VH, −0.016)	(VH, −0.032)	(VH, −0.0)
CL0000392524	2	0.989467	(VH, −0.011)	(VH, −0.016)	(VH, −0.002)	(VH, −0.055)
CL0000412492	2	0.989282	(VH, −0.011)	(VH, −0.016)	(VH, −0.021)	(VH, −0.023)
CL2300356692	3	0.780117	(H, 0.03)	(M, 0.106)	(VH, −0.007)	(VH, −0.036)
CL2200010282	3	0.773729	(H, 0.024)	(M, 0.097)	(VH, −0.026)	(VH, −0.007)
CL2300288384	3	0.772610	(H, 0.023)	(M, 0.099)	(VH, −0.032)	(VH, −0.01)
CL0000550697	3	0.770827	(H, 0.021)	(M, 0.099)	(VH, −0.038)	(VH, −0.009)
CL2200237431	3	0.767832	(H, 0.018)	(M, 0.089)	(VH, −0.007)	(VH, −0.05)

Below, we present one representative customer from each cluster, along with their AHP score and strategic recommendations.

- Cluster 0—Passive or Occasional Buyers:

  Customer ID: CL2200297005
  AHP Score: 0.537 (M, 0.037)
  RFM Profile: (Recency: M, Frequency: H, Monetary: L)
  Cluster-level Strategy: Low priority; reactivation optional.
  Personalized Action: Although this client shows high purchase frequency, the monetary value is consistently low, suggesting price sensitivity or low-value baskets. A possible tactic would be personalized cross-selling suggestions to increase cart value or tiered discount campaigns to encourage incremental spending.
- Cluster 1—Recent but Uncommitted

  Customer ID: CL0000018850
  AHP Score: 0.757 (High, 0.007)
  RFM Profile: (Recency: VH, Frequency: M, Monetary: L)
  Cluster-level Strategy: Potential for loyalty development.
  Personalized Action: This profile is ideal for a welcome-back or first-month engagement campaign, possibly combining educational content (about the brand, offers, or services) with small loyalty rewards to reinforce early-stage conversion. The key is to increase touchpoints before disengagement occurs.
- Cluster 2—High-Value Loyalists

  Customer ID: CL2200483478
  AHP Score: 1.000 (Very High)
  RFM Profile: (Recency: VH, Frequency: VH, Monetary: VH)
  Cluster-level Strategy: Maximum priority; loyalty reinforcement.

Personalized Action: This client should be included in VIP or premium loyalty programs, with exclusive access to new products, birthday or anniversary incentives, and possibly early renewal offers. Personalized communication through high-value channels (e.g., app notifications, direct email) is recommended to maintain a strong relationship and perceived value.

- Cluster 3—Dormant Champions

  Customer ID: CL2300356692
  AHP Score: 0.780 (High, 0.03)
  RFM Profile: (Recency: M, Frequency: VH, Monetary: VH)
  Cluster-Level Strategy: Reactivation of valuable profiles.
  Personalized Action: This profile signals high past engagement but recent inactivity. A reactivation campaign using personalized incentives (e.g., discount on favorite products, win-back messages referencing past behavior) could be effective. It is critical to act before complete disengagement occurs.

These representative cases illustrate how the AHP-enhanced segmentation framework enables a dual-level strategy: it aligns marketing actions to cluster-wide behavioral logic, while also delivering personalized communication plans based on individual scoring and potential. This integrated approach strengthens the connection between data-driven insights and business impact.

4.5. Churn Risk Analysis

As a preliminary step in identifying customers at risk of churn, we applied a rule-based classification based on the normalized Recency score. Customers were segmented into three behavioral categories:

- Active: Percent_Recency > 0.75. Customers who purchased very recently.
- Latent: Percent_Recency between 0.4 and 0.75. Customers with moderate Recency.
- Probable Churn: Percent_Recency < 0.4. Customers who have not purchased in a long time.

This classification enables a quick behavioral overview of the customer base and aligns with the temporal component of engagement, which is heavily weighted in the AHP model. Table 7 shows the distribution of customers across these Recency categories.

Table 7. Distribution of customers by Recency-based status.

Status	# Customers
Probable Churn	13,493
Latent	11,662
Active	8368

To assess the robustness of this rule-based segmentation and enable scalable detection, we trained a binary XGBoost classifier to predict the Churn label (1 = probable churn, 0 = otherwise), using the following features: Percent_Recency, Percent_Frequency, Percent_Monetary and AHP_Score.

The model achieved perfect classification performance on the test set:

- Accuracy: 1.00.
- Precision, Recall, F1-score: 1.00 for both classes.
- Confusion Matrix: No misclassifications.

These results demonstrate a strong alignment between the churn rule definition and the input features, particularly highlighting the influence of Recency on the AHP Score.

While this near-perfect performance is coherent with the deterministic nature of the churn labeling, it should be interpreted with caution, as such results may not generalize to more complex or dynamic churn scenarios.

To interpret the decision logic of the classifier, we applied SHAP. As illustrated in Figure 15, Percent_Recency is by far the most influential variable, confirming that the model prioritizes recent activity in determining churn risk.

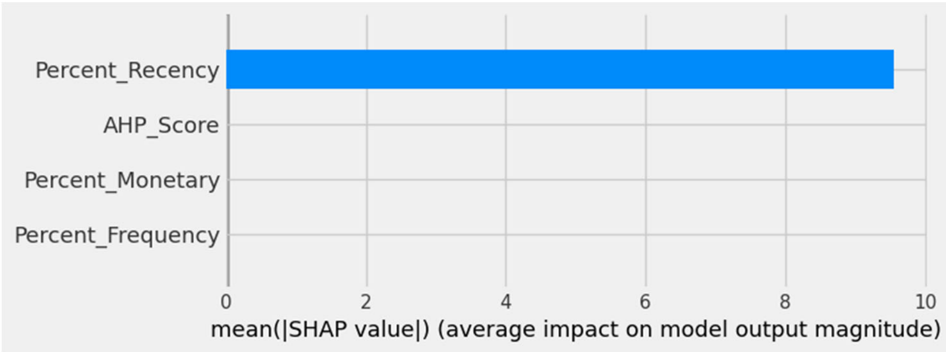

Figure 15. Mean SHAP values indicating feature contributions to churn prediction.

While this outcome confirms a strong alignment between the input features and the target label, it also reveals a potential limitation: the model may be reproducing the rule itself rather than uncovering more complex or emergent behavioral patterns. The comparatively lower importance of AHP_Score and the other RFM components suggests that Recency alone may suffice to explain churn behavior in this particular context, although this may not generalize well to more dynamic or multi-factorial definitions of disengagement.

Although the dominance of Recency is expected, since it was embedded in both the rule-based labeling and the AHP weighting, SHAP adds interpretive value by confirming that the model's decision logic is transparent, aligned with strategic priorities, and suitable for future automated monitoring.

5. Discussion

The proposed approach enhances customer segmentation by ensuring interpretability, personalization, and alignment with business priorities. The methodology enhances traditional RFM models by introducing personalized scoring, semantic interpretability, and local/global explainability, thus bridging the gap between technical robustness and managerial applicability.

By embedding fuzzy logic and 2-tuple linguistic representations, the model captures nuances in customer behavior that are often lost in conventional segmentation schemes. Furthermore, the inclusion of AHP allows for the explicit weighting of business priorities, offering a scoring system that is both data-driven and aligned with strategic objectives.

5.1. Summary of Results and Contributions

The combination of fuzzy clustering and AHP scoring enabled the identification of interpretable and overlapping customer segments, confirmed through both supervised classification (XGBoost) and local explanations (LIME). The clusters obtained, ranging from passive buyers to loyal champions, showed semantic coherence and practical relevance.

The AHP is employed not as a data-driven statistical technique, but as a structured decision-support method to reflect strategic business priorities. It complements explainable AI by incorporating expert judgment into customer evaluation, offering a human-centered scoring mechanism that enhances managerial alignment. Within this framework, the AHP scoring layer enabled fine-grained prioritization by assigning individualized scores that complemented the cluster structure. This allowed for the design of targeted actions beyond segment-level campaigns, such as the following:

- Personalized reactivation strategies for dormant champions (Cluster 3 with high past value but low Recency);
- Loyalty reinforcement for high-scoring profiles (Cluster 2);
- Differentiated onboarding paths for recent but low-commitment clients (Cluster 1);
- And cross-selling or value-enhancement strategies for passive or occasional buyers (Cluster 0), of which the frequency may indicate routine engagement but whose spending levels remain low.

In addition, the application of rule-based heuristics and predictive modeling revealed customer profiles at risk of churn. This analysis, based on Recency degradation and declining AHP scores, was confirmed by an XGBoost classifier with perfect accuracy, and further validated through SHAP, which highlighted Recency as the dominant predictive factor.

Importantly, this work demonstrates how the AHP framework can evolve from a global weighting tool into a mechanism for individual strategic assessment, supporting personalized marketing, service tiering, and retention decisions.

5.2. Model Limitations and Opportunities for Improvement

While the proposed framework demonstrates robustness and strategic value, several improvement areas can be explored to enhance its precision, adaptability, and scalability:

- Dataset-Specific Bias: The dataset used reflects a specific commercial context (Spanish retail of pet products), which may introduce behavioral patterns that are not generalizable. We encourage future validation in other domains (e.g., telecom, banking) to assess portability and mitigate contextual bias.
- Weight sensitivity in AHP: The model relies on expert-defined weights for Recency, Frequency, and Monetary dimensions. Although justifiable, this approach may introduce subjectivity. Future iterations could incorporate data-driven weight optimization (e.g., via entropy methods or sensitivity analysis) to validate or adjust expert judgment.
- Dynamic Recency modeling: The current model uses a static cut-off for Recency-based churn prediction. In practice, customer inactivity may follow variable patterns depending on segment or industry. Introducing time series-based or adaptive thresholds could improve detection of behavioral shifts.
- Integration with qualitative signals: The model currently focuses on transactional behavior. Extending the input features to include qualitative feedback, satisfaction ratings, or sentiment analysis (e.g., from surveys or support interactions) could offer a more holistic view of customer health and loyalty.
- Explainability versus complexity trade-off: While the framework prioritizes interpretability, more complex models (e.g., neural networks or ensemble hybrids) could improve predictive performance. A hybrid approach balancing performance and transparency could be considered for high-stakes decision contexts.
- Operational deployment challenges: Finally, transitioning from analytical models to real-time operational tools requires addressing integration with CRM platforms, automated scoring pipelines, and user-friendly dashboards for business users.
- Risk of Overfitting: The exceptionally high accuracy observed in both clustering prediction and churn classification models suggests a potential risk of overfitting. This may be partly due to the deterministic nature of the Recency threshold used to label churn and should be revisited using more sophisticated or longitudinal churn definitions.

These improvement lines do not challenge the validity of the current approach but rather point to natural extensions that could further enhance its business value and adaptability to diverse operational environments.

5.3. Future Work

Building on the contributions of this study, several promising directions for future research can be identified to extend the framework both conceptually and practically:

- Extension to alternative valuation models: The methodology developed here could be adapted to work with broader customer valuation approaches such as the Customer Engagement Value (CEV) model proposed by Kumar et al. [1]. These models consider dimensions like customer influence, advocacy, and knowledge contribution, which could be assessed through fuzzy linguistic variables and integrated into an extended AHP structure.
- Incorporation of the RFID model: A key direction involves integrating the RFID model (Recency, Frequency, Importance, and Duration), proposed by the author, to evaluate customer profiles based on interaction data with the contact center [28]. This would enrich the RFM framework with relational and experiential dimensions, enabling a more complete picture of customer engagement beyond purchases.
- Real-time scoring and automation: The current model is designed for batch processing and periodic analysis. Future work could focus on deploying the framework in real-time environments, with dynamic updating of scores and segments as new behavioral data arrives. This would be particularly useful in CRM systems or digital platforms requiring immediate decision support.
- Cross-sector generalization: While this study used data from the retail sector, the approach could be validated in other industries, such as telecom, banking, or e-commerce. Testing the model under different customer life cycles and business rules could help refine its generalizability and sector-specific adaptations.
- Integration with behavioral and attitudinal data: Finally, the model could be enriched by incorporating data from customer surveys, satisfaction scores, or Net Promoter Scores (NPS) [48]. Combining behavioral scoring with attitudinal indicators would support a more comprehensive customer-centric strategy.

6. Conclusions

This study presents a novel, integrated methodology for customer segmentation and strategic scoring that addresses the limitations of traditional RFM models by incorporating fuzzy logic, AHP prioritization, and explainable artificial intelligence (XAI). Unlike classical segmentation frameworks that rely on rigid clusters and fixed rules, our approach introduces semantic flexibility, individualized strategic scoring, and behavioral transparency, offering a more operationally meaningful representation of customer profiles.

The combined use of Fuzzy C-Means, 2-tuple linguistic modeling, and AHP-based prioritization enables the detection of interpretable and behaviorally overlapping customer segments, and also the fine-grained ranking of individual clients in accordance with business priorities. This evolution from segment-level targeting to personalized strategic assessment allows companies to tailor marketing actions, service tiers, and retention strategies with a much higher level of precision.

A key advantage of the model lies in its explainability by design. Through the use of SHAP and LIME, both global and local decision processes are made transparent, reinforcing managerial trust and supporting informed, customized interventions. This focuses on interpretability and alignment with strategic decision-making positions the framework as a valuable tool not only in marketing but across various operational domains, including customer service, loyalty management, and churn prevention.

In contrast with conventional approaches, the methodology supports multi-criteria evaluation, accommodates fuzzy and relational dimensions, and is designed to evolve toward real-time, adaptive customer intelligence systems. Its modular structure enables extensions into advanced scoring models such as Customer Engagement Value (CEV) and the RFID model (Recency, Frequency, Importance, Duration), which incorporates customer interaction data.

Furthermore, the proposed framework is grounded in a broader research trajectory where Fuzzy C-Means and XAI techniques have proven effective across diverse domains, including ethical decision-making, astrophysical classification, and business intelligence. This consistent performance underscores the robustness and scalability of the methodology.

In this regard, the study contributes to a conceptual reference model for strategic and interpretable customer analytics. Its design is generalizable across sectors such as retail, banking, and telecommunications, enhancing its relevance as both a research instrument and a practical decision-making tool.

Overall, this work advances the field of human-centered data science, offering a powerful, explainable, and strategically aligned methodology for turning behavioral data into actionable insights. By bridging methodological innovation and operational relevance, it fosters a common language between data science and managerial decision-making.

Funding: This research received no external funding.

Data Availability Statement: The dataset used in this study was fully anonymized and provided by a Spanish retail company for academic demonstration purposes. Due to internal confidentiality policies, it cannot be publicly shared. All data processing was conducted in alignment with GDPR-compliant practices. Further inquiries can be directed at the corresponding author.

Conflicts of Interest: The author declares no conflicts of interest.

References

1. Kumar, V.; Aksoy, L.; Donkers, B.; Venkatesan, R.; Wiesel, T.; Tillmanns, S. Undervalued or overvalued customers: Capturing total customer engagement value. *J. Serv. Res.* **2010**, *13*, 297–310. [CrossRef]
2. Bezdek, J.C.; Ehrlich, R.; Full, W. FCM: The fuzzy c-means clustering algorithm. *Comput. Geosci.* **1984**, *10*, 191–203. [CrossRef]
3. Díaz, G.M.; Medina, R.G.; Jiménez, J.A.A. Integrating Fuzzy C-Means Clustering and Explainable AI for Robust Galaxy Classification. *Mathematics* **2024**, *12*, 2797. [CrossRef]
4. Lundberg, S.M.; Lee, S.-I. A Unified Approach to Interpreting Model Predictions Scott. *Nips* **2012**, *16*, 426–430.
5. Ribeiro, M.T.; Singh, S.; Guestrin, C. 'Why Should I Trust You?' Explaining the Predictions of Any Classifier. In Proceedings of the NAACL-HLT 2016-2016 Conference of the North American Chapter of the Association for Computational Linguistics: Demonstrations, San Diego, CA, USA, 12–17 June 2016; pp. 97–101. [CrossRef]
6. Herrera, F.; Martinez, L. An approach for combining linguistic and numerical information based on the 2-tuple fuzzy linguistic representation model in decision-making. *Int. J. Uncertain. Fuzziness Knowl.-Based Syst.* **2000**, *8*, 539–562. [CrossRef]
7. Saaty, T.L. *The Analytic Hierarchy Process: Planning, Priority Setting, Resource Allocation LK-New York*; SE-xiii, 287 p. Illustrations; 24 cm; McGraw-Hill International Book Co.: New York, NY, USA, 1980; Available online: https://ucm.on.worldcat.org/oclc/911278091 (accessed on 17 March 2025).
8. Kumar, V. A Theory of Customer Valuation: Concepts, Metrics, Strategy, and Implementation. *J. Mark.* **2018**, *82*, 1–19. [CrossRef]
9. Wilbert, H.J.; Hoppe, A.F.; Sartori, A.; Stefenon, S.F.; Silva, L.A. Recency, Frequency, Monetary Value, Clustering, and Internal and External Indices for Customer Segmentation from Retail Data. *Algorithms* **2023**, *16*, 396. [CrossRef]
10. Chen, X.; Gan, W.; Chen, Z.; Zhu, J.; Cai, R.; Yu, P.S. Toward Targeted Mining of RFM Patterns. *IEEE Trans. Neural Netw. Learn. Syst.* **2025**, 1–14. [CrossRef]
11. Wang, G. Customer segmentation in the digital marketing using a Q-learning based differential evolution algorithm integrated with K-means clustering. *PLoS ONE* **2025**, *20*, e0318519. [CrossRef]
12. Ling, S.S.; Too, C.W.; Wong, W.Y.; Hoo, M.H. Customer Relationship Management System for Retail Stores Using Unsupervised Clustering Algorithms with RFM Modeling for Customer Segmentation. In Proceedings of the 2024 IEEE 14th Symposium on Computer Applications & Industrial Electronics (ISCAIE), Penang, Malaysia, 24–25 May 2024; Volume 2024, pp. 550–555. [CrossRef]
13. Handojo, A.; Pujawan, N.; Santosa, B.; Singgih, M.L. A multi layer recency frequency monetary method for customer priority segmentation in online transaction. *Cogent Eng.* **2023**, *10*, 2162679. [CrossRef]

14. Khajvand, M.; Zolfaghar, K.; Ashoori, S.; Alizadeh, S. Estimating customer lifetime value based on RFM analysis of customer purchase behavior: Case study. *Procedia Comput. Sci.* **2011**, *3*, 57–63. [CrossRef]
15. Martínez, R.G.; Carrasco, R.A.; García-Madariaga, J.; Gallego, C.P.; Herrera-Viedma, E. A comparison between Fuzzy Linguistic RFM Model and traditional RFM model applied to Campaign Management. Case study of retail business. *Procedia Comput. Sci.* **2019**, *162*, 281–289. [CrossRef]
16. Carrasco, R.A.; Blasco, M.F.; Herrera-Viedma, E. A 2-tuple fuzzy linguistic RFM model and its implementation. *Procedia Comput. Sci.* **2015**, *55*, 1340–1347. [CrossRef]
17. Liu, D.R.; Shih, Y.Y. Integrating AHP and data mining for product recommendation based on customer lifetime value. *Inf. Manag.* **2005**, *42*, 387–400. [CrossRef]
18. Díaz, G.M.; Salvador, J.L.G.; Hernández, J.J.G. Smart Cities and Citizen Adoption: Exploring Tourist Digital Maturity for Personalizing Recommendations. *Electronics* **2023**, *12*, 3395. [CrossRef]
19. Monalisa, S.; Nadya, P.; Novita, R. Analysis for customer lifetime value categorization with RFM model. *Procedia Comput. Sci.* **2019**, *161*, 834–840. [CrossRef]
20. Martínez, R.G.; Carrasco, R.A.; Sanchez-Figueroa, C.; Gavilan, D. An rfm model customizable to product catalogues and marketing criteria using fuzzy linguistic models: Case study of a retail business. *Mathematics* **2021**, *9*, 1836. [CrossRef]
21. Lang, L.; Zhou, S.; Zhong, M.; Sun, G.; Pan, B.; Guo, P. A Big Data Based Dynamic Weight Approach for RFM Segmentation. *Comput. Mater. Contin.* **2023**, *74*, 3503–3513. [CrossRef]
22. Chang, T.H.; Hsu, K.Y.; Fu, H.P.; Teng, Y.H.; Li, Y.J. Integrating FSE and AHP to Identify Valuable Customer Needs by Service Quality Analysis. *Sustain.* **2022**, *14*, 1833. [CrossRef]
23. Herrera, F.; Martínez, L. A 2-tuple fuzzy linguistic representation model for computing with words. *IEEE Trans. Fuzzy Syst.* **2000**, *8*, 746–752. [CrossRef]
24. Díaz, G.M.; González, R.A.C. Fuzzy Logic and Decision Making Applied to Customer Service Optimization. *Axioms* **2023**, *12*, 448. [CrossRef]
25. Abbas, F.; Ali, J.; Mashwani, W.K.; Syam, M.I. An integrated group decision-making method under q-rung orthopair fuzzy 2-tuple linguistic context with partial weight information. *PLoS ONE* **2024**, *19*, e0297462. [CrossRef]
26. Yan, L.; Pei, Z. A novel linguistic decision-making method based on the voting model for large-scale linguistic decision making. *Soft Comput.* **2022**, *26*, 787–806. [CrossRef]
27. Díaz, G.M.; Salvador, J.L.G. Group Decision-Making Model Based on 2-Tuple Fuzzy Linguistic Model and AHP Applied to Measuring Digital Maturity Level of Organizations. *Systems* **2023**, *11*, 341. [CrossRef]
28. Díaz, G.M.; Carrasco, R.A.; Gómez, D. RFID: A Fuzzy Linguistic Model to Manage Customers from the Perspective of Their Interactions with the Contact Center. *Mathematics* **2021**, *9*, 2362. [CrossRef]
29. Jipkate, B.R.; Gohokar, V.V. A Comparative Analysis of Fuzzy C-Means Clustering and K Means Clustering Algorithms. *Int. J. Comput. Eng. Res.* **2012**, *2*, 737–739.
30. Christy, A.J.; Umamakeswari, A.; Priyatharsini, L.; Neyaa, A. RFM ranking—An effective approach to customer segmentation. *J. King Saud Univ.-Comput. Inf. Sci.* **2021**, *33*, 1251–1257. [CrossRef]
31. Alghamdi, A. A Hybrid Method for Big Data Analysis Using Fuzzy Clustering, Feature Selection and Adaptive Neuro-Fuzzy Inferences System Techniques: Case of Mecca and Medina Hotels in Saudi Arabia. *Arab. J. Sci. Eng.* **2023**, *48*, 1693–1714. [CrossRef]
32. Ling, L.S.; Weiling, C.T. Enhancing Segmentation: A Comparative Study of Clustering Methods. *IEEE Access* **2025**, *13*, 47418–47439. [CrossRef]
33. Chen, W.; Tong, S.; Yuan, Z.; Fang, X. Module Configuration of Rail Freight Transportation with Both Customer Segmentation and Product Family Genealogy in China. *Mathematics* **2024**, *12*, 3947. [CrossRef]
34. Antonius, V.H.; Fitrianah, D. Enhancing Customer Segmentation Insights by using RFM + Discount Proportion Model with Clustering Algorithms. *Int. J. Adv. Comput. Sci. Appl.* **2024**, *15*, 902–911. [CrossRef]
35. Díaz, G.M.; Galán, J.J.; Carrasco, R.A. XAI for Churn Prediction in B2B Models: A Use Case in an Enterprise Software Company. *Mathematics* **2022**, *10*, 3896. [CrossRef]
36. Molnar, C. *Interpretable Machine Learning. A Guide for Making Black Box Models Explainable*, 3rd ed.; 2025; ISBN 978-3-911578-03-5. Available online: https://christophm.github.io/interpretable-ml-book (accessed on 18 March 2025).
37. Gramegna, A.; Giudici, P. SHAP and LIME: An Evaluation of Discriminative Power in Credit Risk. *Front. Artif. Intell.* **2021**, *4*, 752558. [CrossRef]
38. Szeląg, M.; Słowiński, R. Explaining and predicting customer churn by monotonic rules induced from ordinal data. *Eur. J. Oper. Res.* **2024**, *317*, 414–424. [CrossRef]
39. Văduva, A.G.; Oprea, S.V.; Niculae, A.M.; Bâra, A.; Andreescu, A.I. Improving Churn Detection in the Banking Sector: A Machine Learning Approach with Probability Calibration Techniques. *Electronics* **2024**, *13*, 4527. [CrossRef]

40. Joy, U.G.; Hoque, K.E.; Uddin, M.N.; Chowdhury, L.; Park, S.B. A Big Data-Driven Hybrid Model for Enhancing Streaming Service Customer Retention Through Churn Prediction Integrated with Explainable AI. *IEEE Access* **2024**, *12*, 69130–69150. [CrossRef]
41. Asif, D.; Arif, M.S.; Mukheimer, A. A data-driven approach with explainable artificial intelligence for customer churn prediction in the telecommunications industry. *Results Eng.* **2025**, *26*, 104629. [CrossRef]
42. Díaz, G.M.; Medina, R.G.; Jiménez, J.A.A. A Methodological Framework for Business Decisions with Explainable AI and the Analytic Hierarchical Process. *Processes* **2025**, *13*, 102. [CrossRef]
43. Ishizaka, A.; Labib, A. Review of the main developments in the analytic hierarchy process. *Expert Syst. Appl.* **2011**, *38*, 14336–14345. [CrossRef]
44. Bhatt, U.; Xiang, A.; Sharma, S.; Weller, A.; Taly, A.; Jia, Y.; Ghosh, J.; Puri, R.; Moura, J.M.F.; Eckersley, P. Explainable machine learning in deployment. In Proceedings of the FAT* 2020-2020 Conference on Fairness, Accountability, and Transparency, New York, NY, USA, 27–30 January 2020; pp. 648–657. [CrossRef]
45. Rousseeuw, P.J. Silhouettes: A graphical aid to the interpretation and validation of cluster analysis. *J. Comput. Appl. Math.* **1987**, *20*, 53–65. [CrossRef]
46. Syakur, M.A.; Khotimah, B.K.; Rochman, E.M.S.; Satoto, B.D. Integration K-Means Clustering Method and Elbow Method for Identification of the Best Customer Profile Cluster. *IOP Conf. Ser. Mater. Sci. Eng.* **2018**, *336*, 012017. [CrossRef]
47. Chen, T.; Guestrin, C. XGBoost: A scalable tree boosting system. In Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, New York, NY, USA, 13–17 August 2016; 13–17, pp. 785–794. [CrossRef]
48. Baehre, S.; O'Dwyer, M.; O'Malley, L.; Lee, N. The use of Net Promoter Score (NPS) to predict sales growth: Insights from an empirical investigation. *J. Acad. Mark. Sci.* **2021**, *50*, 67–84. [CrossRef]