

Customer-Centric Retail Analytics Using RFM Segmentation and Segment-Level FP-Growth for Product Recommendation

Muhammad Tegar Juleo Wicaksana*, Anteng Widodo, Andy Prasetyo Utomo

Department of Information System, Faculty of Engineering, Universitas Muria Kudus, Indonesia

Email: ^{1,*}tegarwicsleo@gmail.com, ²anteng.widodo@umk.ac.id, ³andy.prasetyo@umk.ac.id

Correspondence Author Email: tegarwicsleo@gmail.com*

Submitted: 14/08/2026; Accepted: 09/09/2026; Published: 30/09/2026

Abstract—Retail transaction data are often used only for sales recording, leaving customer value and product affinity disconnected from managerial decisions. This study develops a web-based retail analytics framework that integrates RFM, K-Means, and global and segment-level FP-Growth to generate product-bundle, availability-check, and shelf-placement recommendations for SRC Inka Jaya. The analysis used 198,879 non-tobacco transaction lines representing 67,284 invoices, 2,565 products, and 1,257 identified customers from 1 January 2025 to 22 June 2026. RFM variables were transformed using $\log(1 + x)$, standardized, and evaluated for $k=2-6$. A four-cluster solution was selected as a trade-off between statistical quality and managerial usefulness, with a silhouette score of 0.371 and a Davies-Bouldin index of 0.934. Champion customers represented 28.48% of identified customers but contributed 54.35% of their monetary value. The primary FP-Growth scenario produced 963 rules, including 188 global and 775 segment-level rules, consolidated into 86 bundles, 46 availability-priority products, and 69 shelf-placement recommendations. Temporal validation achieved a 60.61% rule-match hit rate. Combining store-wide and segment-specific patterns produced more contextual recommendations and supported traceable delivery through an active-dataset dashboard.

Keywords: Retail Analytics; RFM; K-Means; FP-Growth; Product Bundles

1. INTRODUCTION

Minimarkets record thousands of routine purchases, yet this data is not always used as a source of business knowledge. In many small-scale retail businesses, cashier reports mainly serve to recap sales and reconcile stock. Owners lack a measurable picture of high-value customers, groups that are beginning to lapse, or product pairs that repeatedly appear together on the same invoice. This condition limits the store's ability to design promotions, maintain the availability of complementary products, and arrange shelves based on actual transaction behavior.

This problem is relevant to SRC Inka Jaya. The store has kept point-of-sale transaction records for more than a year, but customer segmentation and product recommendations have not yet been built through an integrated analytical procedure. Digital innovation within the Sampoerna Retail Community ecosystem can help with sales administration and management [1]. However, the benefits of digitalization only increase once the stored data is turned into a basis for decisions. Forecasting research on small businesses also shows that historical data can support operational planning when processed with an appropriate method [2]. SRC Inka Jaya's needs relate not only to sales trends but also to two interrelated questions: which customers are most valuable, and which products tend to be purchased together by those customer groups.

Recency, frequency, and monetary (RFM) analysis summarizes customer behavior through the interval since the last purchase, the number of transactions, and spending value. Combining RFM with K-Means has produced actionable segmentation across various contexts. [3] highlights the importance of evaluating several internal indices, while [4] applies RFM and K-Means++ to e-commerce. [5] compares K-Means with K-Medoids and Fuzzy C-Means in a web-based application. [6] likewise obtains three customer profiles from normalized online retail data. [7] positions K-Means as the best algorithm based on a combination of the silhouette score, the Davies-Bouldin Index (DBI), and the Calinski-Harabasz Index (CHI), while [8] confirms that visualization helps users understand the characteristics of each group.

Choosing the number of clusters cannot rely on a single metric. Evidence from non-RFM contexts also shows that attributes and business objectives shape cluster structure. [9] forms five segments based on payment characteristics and obtains a DBI of 0.7878. [10] produces four profiles based on personality and consumption data and uses PCA to explain the differences between groups. K-Medoids can serve as an alternative when robustness to outliers is a priority [11], while RFM and K-Means segmentation can be directed toward customer relationship management strategies [12]. Statistical quality therefore needs to be read together with group size, the original profile, and managerial usefulness, as emphasized in multi-index evaluations of RFM research [3], [7].

Customer segmentation alone does not explain product relationships. Market basket analysis complements customer analysis through association rules assessed with support, confidence, and lift. FP-Growth builds an FP-tree structure without generating the full set of candidate itemsets the way Apriori does. [13] and [14] show that Apriori and FP-Growth can produce comparable patterns, but FP-Growth tends to be more stable as the search

space grows. The number of rules is also strongly influenced by the support and confidence thresholds [15]. In retail, this method has been used to find supermarket product combinations [16], support cross-selling at health stores [17], build bundles from sparse data under a controlled low threshold [18], and link basket patterns to sales and inventory decisions [19].

Association rules need to be filtered so that they do not remain an overly long list of patterns. [20] selects rules for promotional recommendations, while [21] translates product co-occurrence into layout proposals. Most studies still mine transactions globally. This approach treats customers as a homogeneous population, even though patterns from high-value groups may be obscured by store-wide aggregation. [22] shows that segment-level FP-Growth produces a different association context. [23] likewise combines RFM, K-Means, and FP-Growth for marketing strategy, but has not yet placed global and segment-level rules, temporal validation, and several operational outputs within a single pipeline. Web-based implementations have been used for promotional recommendations [24] and retail purchase-pattern analysis [25].

While prior integrated studies [22], [23] demonstrate the benefits of combining RFM segmentation with FP-Growth association mining, their frameworks present specific methodological boundaries. [22] focused primarily on segment-specific itemsets in an e-commerce setting without systematically contrasting store-wide global rules against segment-level rules. Similarly, [23] applied RFM, K-Means, and FP-Growth for promotional targeting, but did not incorporate absolute-count threshold filtering for smaller segments, temporal stability validation across transaction periods, or multi-operational output generation. To address these specific gaps, the analytical framework proposed in this study introduces a dual-tier execution structure that evaluates global and segment-level rules side by side, enforces absolute-count support limits to prevent pattern loss in smaller customer clusters, evaluates rule recurrence using a 20% temporal holdout dataset, and maps derived patterns into three distinct operational artifacts (product bundles, availability-priority lists, and shelf-placement layouts) within a multi-dataset dashboard environment.

Based on this problem, the research question is: how can RFM- and K-Means-based customer segmentation, together with global and segment-level FP-Growth, be integrated into a web-based retail analytics framework to generate recommendations for product bundles, availability-priority products, and differentiated shelf placement? This study aims to build customer value profiles, determine business-interpretable segmentation, compare global and segment-level patterns, test the temporal recurrence of patterns, and operationalize the results through a multi-dataset dashboard. The research contribution lies in using segmentation as context when computing product associations, and in the traceability of recommendations from source data through to managerial reports.

2. RESEARCH METHODOLOGY

2.1 Research Stages

This study uses an applied quantitative design with SRC Inka Jaya as the case study. The unit of analysis differs at each stage: transaction lines for data cleaning, customers for RFM segmentation, and invoices for FP-Growth. The research sequence consists of data collection, preprocessing, RFM formation, K-Means evaluation, behavioral labeling, global and segment-level FP-Growth mining, recommendation filtering, temporal validation, and dashboard implementation. This sequence is applied consistently so that every result can be traced back to its dataset and analysis parameters.

2.2 Data Source and Preprocessing

The research data comes from SRC Inka Jaya's original transaction records from 1 January 2025 to 22 June 2026. The raw data contains 210,301 rows, covering invoice number, transaction date, product description, quantity, unit price, line value, payment method, and customer identity when recorded. The cashier-operator column was empty in every row and was therefore excluded from the analysis. The payment-method values were uniform and were retained only for documentation, not as a model feature.

Preprocessing standardized column names, parsed Indonesian-format dates and numbers, normalized product names, checked invoice numbers, and recalculated line values where necessary. Rows without a valid invoice, date, product, or numeric value were separated from the modeling data. Products repeated on the same invoice were merged into a single basket item so that quantity would not artificially inflate itemset occurrence. A total of 11,422 tobacco-product rows, covering 34 products, were excluded because the scope of recommendations was directed at general merchandise. This decision was a managerial boundary, not an anonymization process.

The final dataset consisted of 198,879 non-tobacco transaction lines, 67,284 invoices, 2,565 products across 19 categories, and a transaction value of Rp 3,577,262,699. A total of 63,654 invoices contained at least two eligible products. After item-frequency filtering for the global analysis, 63,604 baskets were used in the primary scenario. The gap of 50 baskets arose because every item in those baskets failed to meet the product-eligibility threshold, not because of missing invoices.

Table 1. Data Characteristic and Analysis Parameters

Component	Operationalization
Data period	1 January 2025 – 22 June 2026
Raw / final rows	210,301 / 198,879 non-tobacco rows
Invoices / products / categories	67,284 invoices; 2,565 products; 19 categories
RFM customers	1,257 identified customers
RFM reference date	23 June 2026
RFM preprocessing	$\log(1 + x)$, then StandardScaler
K-Means	$k = 2-6$; random_state = 42; n_init = 10; Lloyd algorithm
Primary FP-Growth scenario	minimum support = 0.003; minimum confidence = 0.20
Recommendation filter	lift ≥ 1.10 ; confidence ≥ 0.20 ; support count ≥ 50 global or ≥ 15 per segment
Temporal validation	most recent 20% of baskets as holdout

To ensure ethical data governance and safeguard customer privacy, transaction records were accessed with store management permission and processed in compliance with data protection principles. Personally identifiable information (PII), such as customer names, was replaced with consistent pseudonymized identifiers, enabling longitudinal transaction aggregation without exposing personal identities. A total of 1,257 customers possessed sufficiently consistent identifier records to compute Recency, Frequency, and Monetary (RFM) metrics. These individuals are defined as 'identified customers' rather than formal registered members. All valid invoices were retained for global Frequent Pattern-Growth (FP-Growth) analysis, whereas segment-level FP-Growth strictly utilized invoices that could be mapped to clustered customers.

2.3 RFM Formation and K-Means Segmentation

For customer i , Recency (R_i) is calculated as the number of elapsed days between the reference date and the customer's last transaction. Frequency (F_i) represents the count of distinct valid invoices, while Monetary (M_i) denotes the cumulative monetary value of valid transactions. The operational mathematical definitions for these metrics are defined in Equations (1)-(3).

$$R_i = t_{ref} - \max(t_{ij}) \quad (1)$$

$$F_i = \text{number of unique invoices of customer } i \quad (2)$$

$$M_i = \sum \text{transaction value of customer } i \quad (3)$$

The Preliminary feature exploration revealed severe right-skewness and heavy-tailed distributions across all three RFM metrics. To normalize feature distributions and stabilize variance prior to clustering, a natural logarithmic transformation, $\log(1 + x)$, was explicitly applied to all three variables, including Recency. Transforming Recency via $\log(1 + R_i)$ is specifically justified because unadjusted recency values contain extreme dormant periods that disproportionately skew Euclidean distance metrics during K-Means partitioning. Logarithmic scaling compresses these large dormant values, ensuring non-linear sensitivity to active behavioral differences. For example, a customer with a recency value of 5 days is considered substantially more recently active than a customer with a recency value of 20 days, capturing a critical behavioral shift that carries greater relative weight in distance calculations than an equivalent 15-day difference between long-inactive customers (e.g., 200 versus 215 days). Additionally, the +1 offset handles zero-day recency observations to prevent undefined logarithms. Following logarithmic transformation, all features were standardized to a mean of zero and unit variance ($\mu = 0, \sigma = 1$) as expressed in Equation (4), ensuring equal weighting across dimensions during distance calculations.

$$z = (\ln(1 + x) - \mu) / \sigma \quad (4)$$

The K-Means clustering algorithm was evaluated across candidate cluster counts from $k = 2$ to $k = 6$. Model execution utilized the Lloyd algorithm with a fixed random seed (random_state = 42) and 10 random centroid initializations (n_init = 10) to ensure computational reproducibility. Cluster validity was assessed using a combination of internal validation metrics: Inertia, Silhouette Coefficient, Davies–Bouldin Index (DBI), and Calinski–Harabasz Index (CHI). Higher Silhouette and CHI values reflect superior cluster cohesion and separation, whereas a lower DBI indicates a more optimal partition. Rather than relying solely on a single mathematical index,

the selection of k integrated validation index trade-offs, cluster size distributions, domain-level RFM profiles, and operational decision utility. The analytical pipeline was implemented in Python using pandas, scikit-learn, and mlxtend, with all parameters and outputs stored to ensure process reproducibility.

Segment profiles were assigned business-oriented labels following an inspection of the original un-transformed mean RFM values per cluster. The group displaying the lowest Recency alongside the highest Frequency and Monetary metrics was designated as *Champion*. The cluster with secondary high-value characteristics was labeled *Potential*. Customers exhibiting recent transactions but lower purchase frequency and spending intensity were designated as *Regular*, whereas the cluster with the highest Recency and lowest overall transaction activity was categorized as *Passive*. These segment labels represent context-specific behavioral interpretations tailored to SRC Inka Jaya rather than absolute or universal customer classifications.

2.4 Global and Segment-Level FP-Growth

Each invoice was converted into a set of unique products. Single-item baskets were not used because they cannot form associations. Three scenarios were tested: conservative (support 0.005; confidence 0.25), primary (0.003; 0.20), and exploratory (0.0015; 0.10). Itemset length was capped at three products. Global products had to appear at least 50 times. The eligibility threshold per segment was adjusted to population size: 20 for Champion, 25 for Potential, and 15 for Regular and Passive. A tiered scenario design was necessary because the volume of patterns is sensitive to threshold changes [15]. Using a low support threshold remains defensible for heterogeneous retail data as long as it is accompanied by an absolute-count limit and an interpretability check [18].

Support measures the proportion of transaction baskets containing both items X and Y , Confidence indicates the conditional probability of purchasing Y given the presence of X , and Lift compares this Confidence against the baseline expected frequency of Y . These three evaluation metrics are formally defined in Equations (5)-(7). A Lift value strictly greater than 1.0 indicates a positive co-occurrence pattern between itemsets; however, it signifies a statistical association rather than a direct causal relationship.

$$\text{support}(X \rightarrow Y) = \frac{\text{count}(X \cup Y)}{N} \quad (5)$$

$$\text{confidence}(X \rightarrow Y) = \frac{\text{support}(X \cup Y)}{\text{support}(X)} \quad (6)$$

$$\text{lift}(X \rightarrow Y) = \frac{\text{confidence}(X \rightarrow Y)}{\text{support}(Y)} \quad (7)$$

The primary recommendation set retained rules with a lift of at least 1.10, a confidence of at least 0.20, and a support count of at least 50 for global rules or 15 for segment-level rules. Reciprocal rules and equivalent structures were merged into bundle candidates. Products that recurred across many rules were used as an availability-check indicator, while strong pairs became shelf-adjacency candidates. The list of availability-priority products is not a best-seller ranking and does not determine restock quantities. Restocking decisions still require physical stock, sales velocity, supplier lead time, expiry dates, and safety stock.

2.5 Temporal Validation and System Architecture

Rule recurrence was tested both descriptively across all baskets and temporally. The most recent 20% of baskets, based on transaction time, was set aside as a holdout. Rules were mined from the earlier period and matched against the holdout without using future transactions during pattern formation. Reported metrics include antecedent coverage, basket hit rate, rule-match hit rate, and mean rule hit rate. Rule-match hit rate measures the proportion of antecedent occurrences followed by the consequent. This value is not equivalent to a promotional conversion rate, since customers did not receive any intervention.

The analytics pipeline was connected to a Laravel 11 application on PHP 8.2, a Blade-Tailwind-Alpine interface, a MySQL database, and Python analytical processes. Each upload has its own input, output, log directory, and dataset_id. Only one dataset is marked active for dashboard queries at a time, but earlier results remain stored. The pipeline produces 13 web-ready outputs 10 of the main CSV files are mapped to database tables. This structure prevents new results from overwriting a completed dataset and preserves traceability between source data, processing, and recommendations.

3. RESULT AND DISCUSSION

3.1 Transaction Profile and Preprocessing Results

Preprocessing produced 198,879 non-tobacco rows from 210,301 initial rows. The average number of product lines per invoice was about 2.96. This structure shows that SRC Inka Jaya's baskets are relatively small, but most invoices still contain more than one product. A total of 63,654 invoices met the initial requirement for association analysis. Single-item invoices were still used in revenue and RFM calculations but cannot generate product rules. The monthly invoice count fluctuated over the observation period, as shown in Figure 1. June 2026 is a partial month because the data ends on 22 June 2026.




Figure 1. Number of invoices per month over the observation period

Using operational data increases ecological relevance because the patterns come from actual shopping behavior. Consequently, the data contains inconsistent product names, long-tail products, small baskets, and customer identities that are not always recorded. The cleaning strategy favored traceability over aggressive imputation. This approach preserves the population distinction: identified customers become the RFM unit, while all valid invoices contribute to the global pattern. Anonymous transactions are therefore not forced into any single segment.

The final transaction value reached Rp 3,577,262,699. This figure was used to form customer monetary value and summarize the data's coverage, not to state profit. The data does not contain cost of goods sold or per-product margin. The resulting recommendations must therefore be combined with margin information before the store sets discounts or price bundles.

3.2 K-Means Evaluation and Segment Profiles

The candidate cluster evaluation results are presented in Table 2 and Figure 2. Mathematically, $k = 2$ achieved the strongest internal statistical separation, yielding a Silhouette Coefficient of 0.482, a Davies-Bouldin Index (DBI) of 0.781, and a Calinski–Harabasz Index (CHI) of 1,659.21. In comparison, $k = 4$ produced a Silhouette Coefficient of 0.371, a DBI of 0.934, and a CHI of 1,317.94. Selecting $k = 4$ represents a deliberate methodological trade-off rather than an optimization based solely on a single mathematical index. Although $k = 2$ provides stronger statistical separation, it merely divides customers into broad high-value and low-value groups, severely limiting differentiated marketing strategies. The $k = 4$ solution was selected because it generated four behaviorally distinguishable profiles with meaningful differences in Recency, Frequency, and Monetary characteristics. These profiles provide sufficient analytical granularity to support targeted retention, customer development, routine engagement, and reactivation strategies while avoiding excessive cluster fragmentation. Furthermore, the resulting four customer segments were reviewed with SRC Inka Jaya management. The discussion focused on the practical relevance of the identified profiles, namely Champion, Potential, Regular, and Passive customers, for customer retention, promotion planning, and reactivation strategies. The management confirmed that the four-segment structure provided actionable customer differentiation to support operational decision-making in the retail context.

Table 2. Internal evaluation of K-Means candidates

k	Inertia	Silhouette	DBI	CHI	Decision
2	1.623,976	0,482	0,781	1.659,21	—
3	1.195,293	0,409	0,950	1.351,11	—
4	907,482	0,371	0,934	1.317,94	Selected
5	752,435	0,336	1,004	1.255,68	—
6	636,429	0,347	0,933	1.232,30	—




Figure 2. Silhouette scores for candidate numbers of clusters

This choice is consistent with applied segmentation research that places business interpretation alongside internal indices [3], [7], [10]. A silhouette score lower than in some other studies [4], [6] does not automatically indicate model failure. Retail customers vary along a continuous spectrum, so boundaries between groups can overlap. The PCA projection in Figure 3 shows this overlap, while the heatmap in Figure 4 confirms differences in average recency, frequency, and monetary value across groups.

Table 3. Behavioral profile of the four customer segments

Segment	Customers	Proportion	Mean R	Mean F	Mean M	Contribution to M
Champion	358	28,48%	6,21	47,56	Rp2.857.198	54,35%
Potential	290	23,07%	25,56	30,93	Rp1.679.925	25,89%
Regular	248	19,73%	8,90	13,87	Rp666.570	8,78%
Passive	361	28,72%	66,75	11,62	Rp572.596	10,98%





Figure 3. PCA projection of the four customer segments**Figure 4.** RFM profile heatmap for the four customer segments

The Champion segment contains 358 customers, representing 28.48% of the identified population. A mean Recency of 6.21 days, Frequency of 47.56 invoices, and Monetary value of Rp2,857,198 reflect high purchase activity and strong transaction proximity. This group contributes 54.35% of the total spending value among identified customers. The primary operational implication is to prioritize customer retention and ensure the availability of routine products. Aggressive discounting is unnecessary because purchasing behavior is already robust.

The Potential segment comprises 290 customers with a recency of 25.56 days, frequency of 30.93, and monetary value of Rp 1,679,925. Its 25.89% contribution indicates substantial room for growth. The store can test cross-category bundles or repurchase reminders to raise frequency. Champion and Potential together generate more than four-fifths of the identified customers' value, so a uniform strategy for all customers would overlook this concentration of value.

The Regular segment consists of 248 customers who have transacted only recently but with lower intensity and spending. Back-to-school bundles, meal-need bundles, or clearly priced complementary products are more relevant than complex offers. The Passive segment comprises 361 customers with a recency of 66.75 days. This group requires selective reactivation based on already-familiar product combinations. The Passive label does not mean customers are permanently lost, only that their activity is relatively lower during the observation period.

These four labels are not a classification of members versus regular customers. Customer numbers only replace real names to protect privacy, while frequency and spending value are still calculated from all invoices sharing the same code. This distinction matters so dashboard users do not treat the labels as an administrative status.

3.3 Global and Segment-Level FP-Growth Results

In the primary scenario, the global analysis produced 188 rules from 63,604 filtered baskets. The segment-level analysis produced 207 rules for Champion, 193 for Potential, 187 for Regular, and 188 for Passive. The total of 963 rules consists of 188 global patterns and 775 segment-level patterns 80.48% of the patterns are segment-specific. The distribution of rule counts per group is shown in Figure 5. This finding reinforces research showing that segmentation changes the context of support and confidence calculations [22], [23]. Patterns that are weak at the store level can become strong when computed on a specific group's transactions.

Table 4. Representative global and segment-level rules

Scope	Rule	Support	Count	Confidence	Lift
Global	Minyak Goreng Sunco 1L → Beras Ramos Curah Toko 5KG	7,30%	4.642	75,02%	4,81
Global	Mie Sedaap Goreng 90G → Indomie Goreng 85G	7,11%	4.525	74,13%	7,69

Scope	Rule	Support	Count	Confidence	Lift
Champion	Sunlight Jeruk Nipis 230ML → Mama Lemon 230ML	2,19%	353	89,82%	34,85
Champion	Pepsodent 120G → Sikat Gigi Formula	1,92%	309	74,46%	28,34
Potential	Pampers Baby Dry M 20S → Zwitsal Baby Bath 200ML	1,93%	164	77,00%	27,13
Passive	Kornet Pronas 198G → Sarden ABC Saus Tomat 155G	1,85%	73	70,19%	26,18

The global rules in Table 4 have large transaction support. Cooking oil and rice co-occur in 4,642 baskets, while the two instant-noodle brands co-occur in 4,525 baskets. These patterns are worth considering for continuity of availability and physical proximity, but they do not automatically become discount bundles. Substitute products may be purchased together due to household preference, home-stock needs, or flavor variety. Owners therefore still need to assess margin and promotional purpose.

Segment-level rules have smaller support but higher lift. The Sunlight Mama Lemon relationship within Champion reaches a confidence of 89.82% and a lift of 34.85. This means that, within the Champion population meeting the antecedent, the consequent appears far more often than its baseline frequency. The Champion label indicates the source of the pattern and a more relevant initial target. The bundle may still be offered to other customers the analysis does not restrict sales to a single segment.

The Potential rule linking diapers and baby soap provides context for customer-value development. Conversely, the Passive pattern involving corned beef and sardines can become a low-risk reactivation offer. Segment-level interpretation is more operational than applying a single global rule list to all buyers. This principle aligns with the filtering of promotional recommendations [20] and the integration of segmentation with product association [22], [23].

Table 5. FP-Growth results for the primary scenario

Level	Segment	Baskets	Eligible Baskets	Support	Eligible Items	Itemset	Rules
Global	All	63.654	63.604	0,0030	169	121	188
Segment	Champion	16.118	16.100	0,0030	125	140	207
Segment	Potential	8.499	8.493	0,0030	55	126	193
Segment	Regular	3.243	3.238	0,0046	47	112	187
Segment	Passive	3.958	3.954	0,0038	49	115	188

The effective minimum support ($S_{\text{effective},k}$) for customer segment k was calculated by dividing the minimum itemset occurrence threshold ($T_{\min}$) by the total number of eligible transaction baskets within that segment (N_k), as formulated in Equation (8):

$$S_{\text{effective},k} = \frac{T_{\min}}{N_k} \quad (8)$$

where $T_{\min}$ represents the minimum number of itemset occurrences required to retain an association pattern and N_k represents the number of eligible transaction baskets in segment k . Consequently, differences in effective support values between the *Regular* and *Passive* segments resulted from variations in their transaction basket sizes. For example, the *Regular* segment required 15 occurrences from 3,238 eligible transaction baskets, resulting in an effective support of 0.00463, whereas the *Passive* segment required the same 15 occurrences from 3,954 eligible transaction baskets, resulting in an effective support of 0.00379. The exploratory scenario generated a larger number of low-support rules with limited practical interpretability, whereas the conservative scenario primarily captured highly frequent product combinations. The primary scenario was selected by jointly considering support, confidence, lift, and absolute occurrence counts. This adaptive thresholding strategy is consistent with previous findings that association rule characteristics and operational usefulness are influenced by transaction scale, threshold configuration, and basket density [14], [18].




Figure 5. Number of association rules per customer segment

After reciprocal rules and equivalent structures were merged, 86 bundle suggestions were obtained: four from global patterns and 82 from segment-level patterns. The system also compiled 46 availability-priority products and 69 shelf-placement suggestions. The number of patterns linked to a product is not the same as the number of restock units. For example, a product involved in 93 rules simply means it appears in 93 co-purchase relationships. The largest support count also counts supporting invoices, not units sold.

3.4 Temporal Validation and Recommendation Reliability

Descriptive validation across the full dataset yielded 100% rule coverage, a 99.88% basket hit rate, a 60.08% rule-match hit rate, and a 62.85% mean rule hit rate across 963 final rules. To ensure reproducibility and mathematical transparency, the Basket Hit Rate is defined as the proportion of evaluated transaction baskets containing at least one successful association rule match where both antecedent (X) and consequent (Y) items co-occur relative to the total evaluated basket population (N_{basket}). In the full dataset evaluation ($N_{\text{basket}} = 63,654$), a total of 63,577 baskets contained at least one matching rule pair, yielding a Basket Hit Rate of 99.88%. This metric reaches near-complete coverage because retail transaction baskets frequently contain multi-item combinations that match at least one generated rule among the 963 candidate rules. Consequently, the Basket Hit Rate measures aggregate system-level pattern coverage rather than single-bundle promotional success probability.

To evaluate rule-level recurrence, the Rule-Match Hit Rate measures the proportion of total antecedent occurrences where the corresponding consequent itemset also appears across all rule evaluations. Across 901,274 total antecedent evaluations in the full dataset, 541,449 resulted in successful consequent matches, establishing a baseline Rule-Match Hit Rate of 60.08%.

Table 6. Descriptive and temporal validation of the final rules

Validation	Rules	Basket	Coverage	Basket hit	Rule-match hit	Mean rule hit
Full Data	963	63.654	100,00%	99,88%	60,08%	62,85%
Temporal Holdout	954	12.714	100,00%	99,87%	60,61%	63,95%

The temporal holdout dataset comprised 12,714 multi-item transaction baskets from the subsequent evaluation period. Of the 963 final rules mined from the initial training period, 954 rules were evaluable in the temporal holdout split. The remaining 9 rules could not be evaluated because their specific antecedent item combinations did not occur within any holdout transaction basket ($N_{\text{antecedent}} = 0$), making empirical evaluation of consequent recurrence impossible for those rules. Among the 12,714 holdout baskets, 12,697 contained at least one successful rule match, yielding a holdout Basket Hit Rate of 99.87%. Furthermore, across 180,051 total antecedent evaluations in the holdout split, 109,122 resulted in successful consequent matches, producing a holdout Rule-Match Hit Rate of 60.61% and a Mean Rule Hit Rate of 63.95%. The close alignment between full-dataset and holdout validation metrics suggests that the identified item associations represent recurring purchasing patterns that remain observable across different transaction periods rather than transient noise.

It is important to emphasize that the 60.61% holdout rule-match hit rate reflects temporal consistency of historical co-occurrence patterns, not a commercial promotional conversion rate. This study did not manipulate bundle exposure, apply promotional discounts, or evaluate a randomized control group. The validation metric

strictly establishes that when an antecedent itemset appeared in a holdout basket, the consequent itemset co-occurred in 60.61% of instances. Measuring true causal sales uplift requires in-store field experiments, such as A/B testing comparing average order value, profit margin, and repurchase velocity before and after deploying recommendations to the interactive dashboard.

3.5 System Implementation and Managerial Implications

The analytical results were operationalized through a dashboard, rather than simply delivered as a spreadsheet. Admins manage the dataset process, segmentation details, model evaluation, purchase patterns, bundles, availability-priority products, and shelf suggestions. Owners are given a more concise interface: a main summary, store reports, user management, and system settings. This separation keeps business users focused on decisions, while technical detail remains available to data managers.

The active-dataset dashboard summarizes 67,284 invoices, 1,257 customers analyzed, 2,565 products, 963 purchase patterns, 86 bundle suggestions, 46 availability-priority products, 69 shelf-pairing suggestions, and four customer groups. All figures are read from the dataset marked active, not hard-coded. The two-page PDF report uses the same source, so the period, metrics, and recommendations change whenever the active dataset is switched.




Figure 6. Admin dashboard of the retail analytics system on the active dataset

The first implication is differentiated bundling. Global patterns with strong support can form the basis for general promotions, shelf adjacency, or stock checks. Segment-level patterns are more suitable as initial targets. A Champion bundle does not mean it may only be sold to Champion customers the label simply indicates that the pattern was found in that group's transactions. Personalized action requires mapping the code to identity data, which is stored privately. Without that mapping, the recommendations remain useful at the group level.

The second implication concerns supply continuity. Products that repeatedly appear as an antecedent or consequent can act as a link between multiple purchases. A stock-out of such a product could reduce complementary purchases. However, association rules are not an inventory model. The list of 46 products only directs inspection priority. Order quantities must still be calculated from physical stock, daily sales, lead time, safety stock, capacity, and expiry.

The third implication is shelf placement. A total of 69 pairs indicate products that frequently appear on the same invoice. Physical proximity can reduce search time and support basket completeness [20], [21]. Implementation must still consider category, food safety, shelf capacity, supplier contracts, and substitute products. The dashboard supports decisions; it does not replace the store owner's judgment.

Compared with earlier web-based implementations [24], [25], this system adds dataset linkage, process logs, result activation, global and segment-level analysis, and managerial reporting. Using Python and Laravel separates analytical computation from the operational interface. A web-based approach applied to restaurant data has likewise shown that FP-Growth can support bundling and inventory decisions [26], but this study extends that context through RFM segmentation, temporal validation, role restrictions, and the active-dataset mechanism.

3.6 Research Limitations

This study is limited to a single minimarket and a single observation period, so the segment names and product patterns cannot be generalized directly. RFM covers only customers with a consistent code, while anonymous transactions enter the global pattern. K-Means remains sensitive to scale, outliers, and cluster-center assumptions, even with transformation, standardization, and repeated initialization applied. Association rules show co-occurrence, not intent or causation. The model also does not yet incorporate margin, ongoing promotions, product substitution, supplier lead time, physical stock, or shelf capacity. Evaluation of user adoption and long-term financial impact has not been carried out.

4. CONCLUSION

This study presents an integrated analytics framework linking customer value profiling with product association mining in retail point-of-sale transactions. Following $\log(1 + x)$ transformation and z-score standardization, K-Means clustering established four operational customer segments (*Champion*, *Potential*, *Regular*, and *Passive*), prioritizing operational decision usefulness beyond a binary segmentation approach ($k = 2$). Executing segment-aware FP-Growth mining uncovered localized purchasing patterns masked during store-wide analysis. These rules were systematically consolidated into three operational artifacts product bundles, availability-priority lists, and shelf-placement layouts and embedded into an interactive multi-dataset dashboard. Temporal holdout validation showed that generated patterns remained observable across transaction periods without overclaiming commercial promotional conversion. Primary limitations include single-store scope, partial customer identification, and the absence of inventory constraints or promotional field experiments. Future research should evaluate commercial uplift via in-store A/B testing, incorporate profitability and stock constraints into rule ranking, analyze longitudinal segment migration, and test framework generalizability across multi-branch retail networks.

REFERENCES

- [1] P. Salsa Nabila, A. Indah Lestari Nasution, and M. Ikhsan Harahap, “Analisis Efektivitas Inovasi Digital SRC (Sampoerna Retail Community) Bagi Toko Kelontong Di Kecamatan Kuala Kabupaten Langkat,” *Jurnal Media Wahana Ekonomika*, vol. 20, pp. 355–366, Jul. 2023, [Online]. Available: <https://jurnal.univpgri-palembang.ac.id/index.php/Ekonomika/index>
- [2] A. Widodo and A. Abdul Chamid, “Peramalan Penjualan Genteng Menggunakan Metode Holt-Winters Exponential Smoothing,” *Jurnal SITECH: Sistem Informasi dan Teknologi*, vol. 8, no. 2, pp. 291–302, 2025, [Online]. Available: <https://jurnal.umk.ac.id/index.php/sitech/article/view/16765/6843>
- [3] S. F. Djun, I. G. A. Gunadi, and S. Sariyasa, “Analisis Segmentasi Pelanggan pada Bisnis dengan Menggunakan Metode K-Means Clustering pada Model Data RFM,” *JTIM : Jurnal Teknologi Informasi dan Multimedia*, vol. 5, no. 4, pp. 354–364, Feb. 2024, doi: 10.35746/jtim.v5i4.434.
- [4] A. L. M. Tampubolon, T. M. E. Y. Butar Butar, and S. Rochimah, “Segmentasi Pelanggan Majalah pada Situs Web E-Commerce dengan K-Means++ dan Metode RFM,” *Jurnal Teknologi Informasi dan Ilmu Komputer*, vol. 11, no. 6, pp. 1243–1252, Dec. 2024, doi: 10.25126/jtiik.2024118208.
- [5] R. M. Fauzan and G. Alfian, “Segmentasi Pelanggan E-Commerce Menggunakan Fitur Recency, Frequency, Monetary (RFM) dan Algoritma Klasterisasi K-Means,” *JISKA (Jurnal Informatika Sunan Kalijaga)*, vol. 9, no. 3, pp. 170–177, Sep. 2024.
- [6] A. Yusak, N. Rumapea, D. Pratiwi, and S. Sari, “Analisis Segmentasi Pelanggan Ritel Online Menggunakan K-Means Clustering Berdasarkan Model Recency, Frequency, Monetary (RFM),” *Jurnal Sains dan Teknologi*, vol. 6, no. 3, pp. 292–299, 2024, doi: 10.55338/saintek.v6i3.4607.
- [7] M. F. Fadhillah, A. S. A. Lovely, and I. Puspitasari, “Customer Segmentation with Clustering Algorithm Based on Recency, Frequency, and Monetary (RFM) Attributes,” *MALCOM: Indonesian Journal of Machine Learning and Computer Science*, vol. 5, no. 1, pp. 48–56, Jan. 2025, doi: 10.57152/malcom.v5v1.1491.
- [8] D. A. Imanuel and G. Alfian, “Visualisasi Segmentasi Pelanggan Berdasarkan Atribut RFM Menggunakan Algoritma K-Means Untuk Memahami Karakteristik Pelanggan pada Toko Retail Online,” *Jurnal Teknologi Informasi dan Ilmu Komputer*, vol. 12, no. 2, pp. 283–292, Apr. 2025, doi: 10.25126/jtiik.2025128619.
- [9] A. Syaharani, V. Ilhadi, and A. Karima, “ANALISIS CLUSTERING SEGMENTASI PELANGGAN PADA CV PRIMA JAYA UTAMA MENGGUNAKAN METODE K-MEANS,” *JATI (Jurnal Mahasiswa Teknik Informatika)*, vol. 9, no. 5, Oct. 2025.
- [10] I. Iqbal, N. Hidayat, D. P. Gevano, and A. P. R. Ilahi, “Segmentasi Pelanggan Menggunakan K-Means Clustering Berdasarkan Data Kepribadian dan Pola Konsumsi,” *Jurnal Teknik Informatika (Jutif)*, vol. 6, no. 5, pp. 3914–3924, Oct. 2025, doi: 10.52436/1.jutif.2025.6.5.5140.
- [11] E. W. Agustin, K. Uthami, A. I. Ulfa, L. Efrizoni, and R. Rahmaddeni, “Optimization of Customer Segmentation in the Retail Industry Using the K-Medoid Algorithm,” *MALCOM: Indonesian Journal of Machine Learning and Computer Science*, vol. 5, no. 3, pp. 766–775, Jun. 2025, doi: 10.57152/malcom.v5i3.1977.
- [12] Y. Syahra, A. Fadlil, and H. Yuliansyah, “Customer Segmentation Using RFM and K- Means Clustering to Support CRM in Retail Industry,” *Sinkron : Jurnal dan Penelitian Teknik Informatika*, vol. 9, no. 3, pp. 1108–1119, Jul. 2025, doi: 10.33395/sinkron.v9i3.14974.
- [13] I. Musdalifah and A. Jananto, “Analisis Perbandingan Algoritma Apriori Dan FP-Growth Dalam Pembentukan Pola Asosiasi Keranjang Belanja Pelanggan,” *Jurnal Ilmiah Komputer*, vol. 18, pp. 175–184, Aug. 2022.
- [14] S. R. Karimah and E. D. Udayanti, “Comparison of Apriori and FP-Growth Algorithms in Market Basket Analysis for Online Book Sales,” *Journal of Applied Informatics and Computing (JAIC)*, vol. 10, no. 1, pp. 566–574, Feb. 2026, [Online]. Available: <http://jurnal.polibatam.ac.id/index.php/JAIC>

- [15] U. M. Wulandari, A. T. Suseno, and M. Kholilurrahman, “Market Basket Analysis Using FP-Growth and Apriori on Distro Store Sales Transaction,” *MATICS: Jurnal Ilmu Komputer dan Teknologi Informasi (Journal of Computer Science and Information Technology)*, vol. 17, no. 1, pp. 12–18, Mar. 2025, doi: 10.18860/mat.v17i1.28820.
- [16] H. Hairani and J. Ximenes Guterres, “Exploring Customer Purchasing Patterns: A Study Utilizing FP-Growth Algorithm on Supermarket Transaction Data,” *International Journal of Engineering and Computer Science Applications (IJECSA)*, vol. 3, no. 1, pp. 33–42, Mar. 2024, doi: 10.30812/ijecsa.v3i1.3874.
- [17] F. Achmad, O. Nurdiawan, and Y. A. Wijaya, “ANALISA POLA TRANSAKSI PEMBELIAN KONSUMEN PADA TOKO RITEL KESEHATAN MENGGUNAKAN ALGORITMA FP-GROWTH,” *Jurnal Mahasiswa Teknik Informatika*, vol. 7, no. 1, pp. 168–175, Feb. 2023, doi: <https://doi.org/10.36040/jati.v7i1.6210>.
- [18] V. M. Christy and I. Maryati, “FP-Growth Implementation for Market Basket Analysis in Building Materials Store in Surabaya,” *Journal of Applied Informatics and Computing (JAIC)*, vol. 10, no. 2, pp. 1751–1759, Apr. 2026, [Online]. Available: <http://jurnal.polibatam.ac.id/index.php/JAIC>
- [19] N. Purwati and S. Karnila, “Strategi Peningkatan Penjualan Produk Menggunakan Market Basket Analysis,” *Jurnal Sistem Informasi Bisnis*, vol. 13, no. 2, pp. 96–103, Oct. 2023, doi: 10.21456/vol13iss2pp96-103.
- [20] I. Ismarmiaty and R. Rismayati, “Product Sales Promotion Recommendation Strategy with Purchase Pattern Analysis FP-Growth Algorithm,” *Sinkron : Jurnal dan Penelitian Teknik Informatika*, vol. 7, no. 1, pp. 202–211, Jan. 2023, doi: 10.33395/sinkron.v8i1.11898.
- [21] D. Alfitra, M. Afdal, M. Fronita, and E. Saputra, “Analisa Keranjang Belanja untuk Menentukan Tata Letak Barang Menggunakan Algoritma FP-Growth,” *Sistemasi: Jurnal Sistem Informasi*, vol. 13, no. 4, pp. 1651–1661, Jul. 2024, [Online]. Available: <http://sistemasi.ftik.unisi.ac.id>
- [22] C. Olyfiyany and H. Risma Ameilya, “Segmentasi Pelanggan Berbasis RFM dan Analisis Asosiasi Produk pada Olist Brazilian E-Commerce Menggunakan FP-Growth,” *Journal of Information Systems and Business Technology*, vol. 2, pp. 910–917, 2026, [Online]. Available: <https://journal.jci.co.id/jisbt>
- [23] S. Aulia, A. N. Muhammad, and A. Wibowo, “Optimization of Customer Segmentation with RFM, K-Means, and FP-Growth for Marketing Strategy,” *SINTECH JOURNAL*, vol. 8, pp. 163–177, Aug. 2025, doi: <https://doi.org/10.31598/sintechjournal.v8i2.1942>.
- [24] K. F. Handoko, N. Nugraha, and A. Sujjada, “Customer Segmentation-Based Promotion Recommendation Using RFM and K-Means Clustering on a Web Platform,” *bit-Tech*, vol. 8, no. 3, pp. 3661–3672, Apr. 2026, doi: 10.32877/bt.v8i3.3677.
- [25] S. A. Alfatikhah, S. Supriyono, and S. Adiyono, “IMPLEMENTASI METODE FP-GROWTH DALAM ANALISA POLA PEMBELIAN PELANGGAN PADA SWALAYAN PANTES PATI BERBASIS WEB,” *Rabit : Jurnal Teknologi dan Sistem Informasi Univrab*, vol. 11, no. 1, pp. 1535–1545, Jan. 2026, doi: 10.36341/rabit.v11i1.7401.
- [26] N. A. Khumaira, D. Priyanto, H. Hairani, G. H. Martono, Moch. Syahrir, and H. Husain, “Analisis Pola Pembelian Konsumen Menggunakan Algoritma FP-Growth pada Data Transaksi Restaurant Burger,” *Jurnal Teknologi Informasi dan Multimedia*, vol. 8, no. 3, pp. 497–509, Jul. 2026, doi: 10.35746/jtim.v8i3.983.