Data-Driven Optimization of Customer Acquisition Cost: A Business Model Based on Predictive and Prescriptive Analytics

Kabir Chibber[1], Anand Zutshi[2]

[1]Kunskapsskolan Gurgaon School, Grade 12, India
Email: *chibberkabir[at]gmail.com*

[2]Netaji Subhas University of Technology, Delhi
Email: *123anandji[at]gmail.com*

Abstract: ***This study presents a data-driven framework for optimizing Customer Acquisition Cost (CAC), a key metric for evaluating marketing and operational efficiency. Drawing on a comprehensive dataset from North American supermarkets, the research applies machine learning and segmentation analysis to identify the primary factors influencing CAC. The findings reveal that promotional design and media strategy significantly affect acquisition costs, while store format and amenities offer secondary leverage. A predictive business model is proposed, integrating store characteristics, geographic context, and promotional variables to guide marketing allocation. The results offer actionable insights for businesses aiming to reduce CAC, enhance scalability, and support financial planning through advanced analytics.***

Keywords: customer acquisition cost, predictive analytics, business strategy, machine learning, marketing optimization

1. Introduction

Customer Acquisition Cost (CAC) is widely recognized as a pivotal business metric. It quantifies the expense of acquiring one new customer and therefore directly influences measures of profitability and scalability. CAC insights help firms of all sizes optimize marketing efforts and refine unit economics. Keeping CAC low is critical to a strong return on investment: by revealing which campaigns and channels are most cost-effective, CAC guides strategic budgeting and pricing decisions. This is especially true for resource-constrained startups and small companies (Rosler, 2024). With limited capital and thin margins, a high CAC relative to customer value can quickly exhaust resources, whereas efficient acquisition spending supports sustainable scaling. For example, experts note (Long, 2025) that clear visibility into CAC enables better cash-flow forecasting and break-even planning. In short, CAC guides marketing strategy by supporting profitability, managing cash flow, and ensuring long-term viability.

However, CAC is notoriously difficult to predict and optimize (Ababyomi et al., n.d., 18). Marketing efforts today span many channels and touchpoints, making attribution complex. In practice, campaign data are often scattered across multiple systems – one platform may track social ads, another email, another offline promotions – so teams rarely have a unified view of spending and results. As one analysis observes, when marketing data is "scattered across 6–15 different sources," firms frequently "don't actually know what's driving results". These fragmented systems tend to create data silos and inconsistencies, which delay insights and hinder the ability to tie costs to customer conversions. On top of that, customer behavior and promotional effectiveness vary over time and geography, and external factors (seasonality, competitive shifts, economic conditions) can cause CAC to fluctuate. Indeed, industry commentary warns (Nair, 2025) that such external trends introduce additional challenges in accurately analyzing and predicting CAC. Compounding the problem, many organizations lack advanced analytics tools or real-time integration to model CAC dynamically, meaning they cannot swiftly adjust to changing market conditions.

The motivation for this research comes from the author's own interest and experience in marketing analytics. While working in SEO and digital advertising (e.g. Google Ads), it was repeatedly seen how hard it is for businesses to forecast their acquisition costs. Participation in Harvard's Undergraduate Ventures-TECH summer program further underscored CAC's critical role in sustaining business operations. These experiences revealed a recurring gap: many small firms simply cannot predict CAC accurately. This inspired an exploration of AI and machine learning as a solution. The aim is to apply AI-driven models to integrate diverse factors – from online ads to offline promotions and regional differences – into a unified predictive framework. In particular, comparing the cost-efficiency of various channels versus non-digital marketing through an ML model that takes into account a variety of seller and consumer factors to accurately forecast CAC.

The study has four primary objectives:
- Predict Customer Acquisition Cost (CAC) for businesses of all sizes using artificial intelligence and machine learning models.
- Provide actionable, data-driven insights on acquisition efficiency across different marketing channels and regions.
- Create sample business models and marketing strategies based on the predicted CAC patterns.
- Enable more accurate forecasting of profitability and costs by incorporating the CAC predictions into financial planning.

2. Dataset Description

We used a dataset from Kaggle (Ashik, 2023) with 60,428 customer records from Food Mart stores in Canada, the United States, and Mexico. It had 40 columns about products, promotions, media, stores, and customer details. The key columns included food_category, food_department, food_family, unit_sales, store_sales, store_cost, promotion_name, media_type, store_type, store_city, and store_sqft, alongside indicators for store amenities such as coffee bars, video stores, salad bars, prepared foods, and florists. This setup let us link what was sold (merchandising) to which promotions and media people saw, and study how store amenities, size, and marketing related to customer acquisition costs, revenue, and profitability.

Analysis

Preprocessing was performed in Python to ensure analytical integrity. First, we gave the column names a standard format so they were easier to work with. All money-related values were converted into numbers by removing currency symbols, and "yes/no" fields such as whether a store had a coffee bar or florist (known as boolean indicators) were changed into a consistent format. The store size (store_sqft) and department sizes were set as whole numbers (integers), and then each store was grouped into three categories: small (25,000 square feet or less), medium (25,000–35,000 square feet), and large (over 35,000 square feet). We checked the data to make sure that store sales were always higher than store costs and removed any cases where profit margins were unrealistic (for example, negative or excessively high). In addition, if the same promotion appeared in multiple media channels, we counted it only once to avoid inflating the results (deduplication). These cleaning steps made sure that later comparisons such as how promotions, media, and store amenities affected customer acquisition costs were fair and reliable.

Feature engineering expanded the dataset's interpretive power. Metrics such as revenue per square foot (Total store revenue over the period / Total store square footage) and gross margin rate ((Total store sales − Total store cost) / Total store sales) were computed to benchmark store format productivity. As seen in Fig. 1, Gourmet supermarkets and Supermarkets have the highest revenue per square foot of all store types, while small grocery stores have the least and are thus least productive by this metric. Fig. 2 shows that the gross margin rate of all the store types are roughly the same.

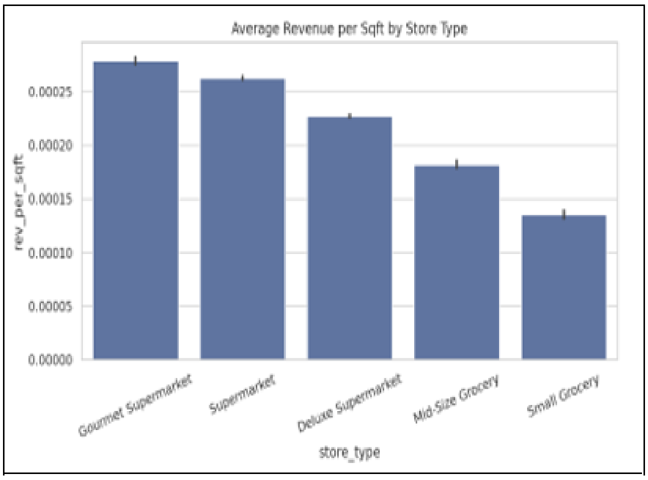

Figure 1: Average Revenue per Sqft

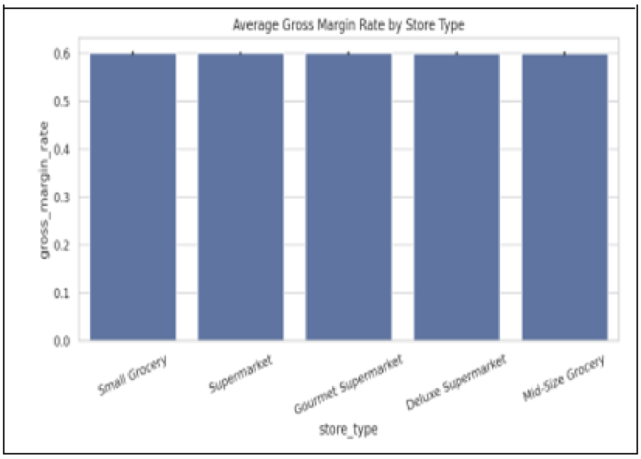

Figure 2: Average Gross Margine Rate by Store Type

Exploratory analysis revealed systematic patterns across store sizes and locations. Small stores consistently exhibited low CAC indices, especially when coffee bars and prepared food amenities were present, while medium stores performed strongly with a broader amenity mix. Large stores achieved the highest revenue per square foot when operating a full amenity bundle, though their CAC indices were higher, reflecting a tradeoff between efficiency and absolute EBITDA contribution (Stewart, 2019). Correlation analysis (fig. 1) showed weak overall relationships between CAC and most numerical variables, though certain amenities—such as coffee bars, florists, and prepared foods—were negatively associated with CAC, indicating that richer in-store experiences helped lower acquisition costs.

International Journal of Science and Research (IJSR)
ISSN: 2319-7064
Impact Factor 2024: 7.101

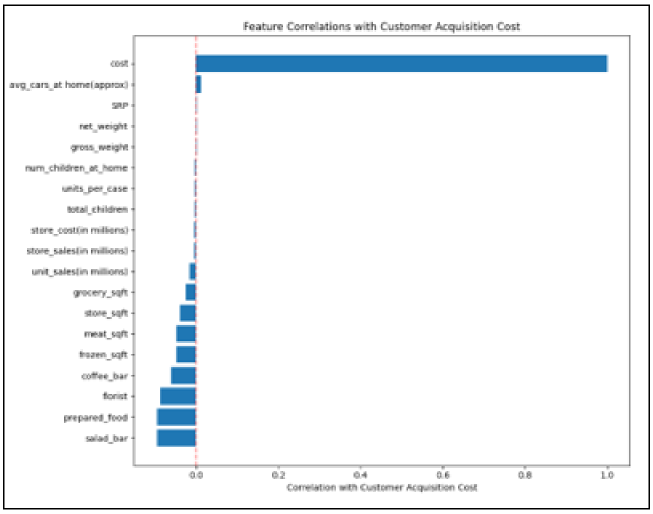

Figure 3: Feature Correlation with Customer Acquisition Cost

Random forest modeling confirmed these insights and provided predictive strength, achieving an R^2 above 0.998.

Feature importance scores emphasized that marketing strategy was far more predictive of CAC than demographic variables, with promotion name (0.44) and media type (0.24) emerging as the dominant drivers, followed by store size characteristics.

Based on the elbow method analysis (Fig. 4), 3 segments were selected because the inertia drop from k=2 to k=3 was substantial (k=2 to k=3: 120022), while subsequent drops were much smaller (k=3 to k=4: 51370), indicating diminishing returns beyond 3 clusters.

We clustered on a balanced mix of demographics (marital status, gender, total and at-home children, education, member card tier, occupation, houseownership, cars at home, and income brackets converted to numeric), store characteristics (store type, city/state, total and department square footage, and amenities like coffee bar, video store, salad bar, prepared foods, and florist), and marketing features (promotion name and media type, such as Daily Paper, Radio, and In-Store Coupon). This gives each cluster a distinct behavioral and contextual footprint, not just a demographic slice.

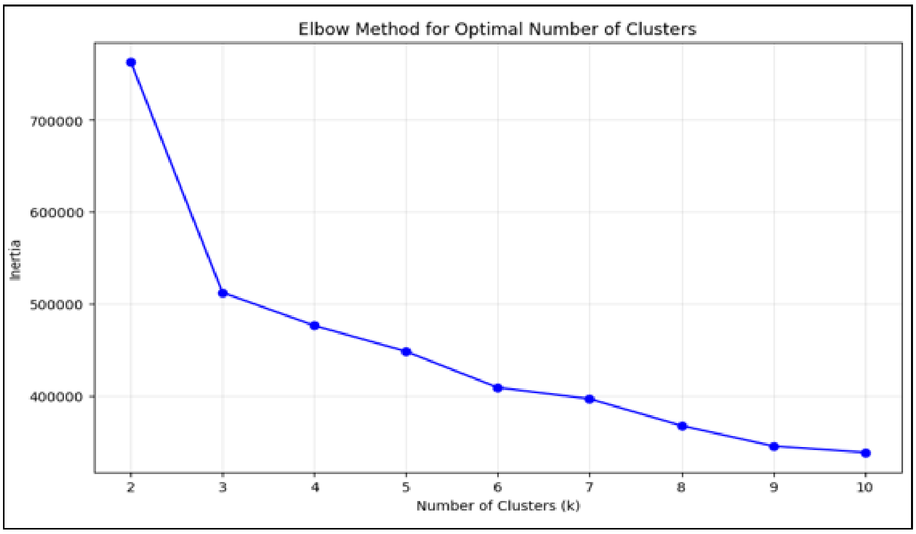

Figure 4: Elbow Method for Optimal Clusters

Segment 0 consists of busy households that respond to convenience and visible value. They tend to shop at larger supermarket formats with more in-store amenities, where family-friendly layouts and department breadth matter. Promotions with straightforward savings messaging perform well; they like the obvious, in-aisle triggers (endcaps, in-store signage) and respond to mass reach channels they already encounter in daily routines. In practice, in-store coupons and weekly circular-driven offers, amplified by Daily Paper and Radio, are reliable levers. Their average acquisition cost is around the low 100s, which is solid given their household size and likely larger basket potential. These customers fit best in suburban or larger-footprint stores with amenities like salad bar, prepared foods, or a florist which are features that make a one-stop trip more attractive. Positioning should emphasize family value, bulk or multi-pack savings, and time-saving options for households managing multiple schedules. Income-wise, they span modest to mid-level brackets, often with multiple cars at home and at least one child at home. Media that reinforces routine shopping (newspaper inserts, local radio) plus on-site nudges (in-store coupons, promoted endcaps) is ideal. Promotions like Bag Stuffers or Double Down Sale that clearly show instant savings work particularly well, especially when layered with practical store amenities that make trips efficient.

Segment 1 skews toward compact, convenience-focused shopping patterns, often urban or near-urban locations where store footprint is smaller and trips are quicker. They prioritize speed and clarity. Promotions that reduce decision friction (e.g., clear price points, limited-time deals) and channels that meet them on-the-go (radio, simple digital placements, and concise paper media) are effective. Their average acquisition cost sits in the low 100s, roughly on par with Segment 0, but their shopping missions are shorter, so messaging should be sharp and time-sensitive rather than sprawling value narratives. Income distribution here leans into mid to upper-mid brackets, with fewer children at home and tighter time windows. Amenities that speed the visit like prepared foods and grab-and-go options matter more than breadth. Media that reaches them during commutes (radio) or via quick-scanning formats (concise paper circulars, store signage) pairs well

International Journal of Science and Research (IJSR)
ISSN: 2319-7064
Impact Factor 2024: 7.101

with promotions like Cash Register Lottery or Double Down Sale where the incentive is obvious and immediate. These customers fit best in denser locations, smaller formats, and store types positioned for frequent, short visits.

Segment 2 delivers the lowest average acquisition cost, in the mid-to-high 90s, making them the most efficient to acquire. They are very responsive to clear promotions and consistent value cues. Straightforward savings messaging and well-labeled offers pay off, and they often reward regularity, i.e., if the store runs a dependable set of deals, they show up. Daily Paper and Radio work well for broad, dependable reach; In-Store Coupon is a strong in-aisle closer. They respond to classic promotions like Bag Stuffers and Double Down Sale, particularly when offers are simple to understand and redemption is seamless.

Income tends to be more budget-sensitive, with a meaningful share in lower brackets and fewer discretionary frills. They are flexible across store formats but perform best where promotions are easy to find and compare like standard or deluxe supermarkets with clean promotional wayfinding. Store amenities are nice-to-have rather than must-have; clarity and consistency matter more than extras. Think disciplined weekly ad cycles, sharp price points, and repeatable featured brands. This segment is a great target for scale: wide reach, predictable promotions, and reliable media channels produce the best ROI.

In Fig. 5, the clusters look reasonably separated, indicating the three-segment solution captures distinct structure in the data.

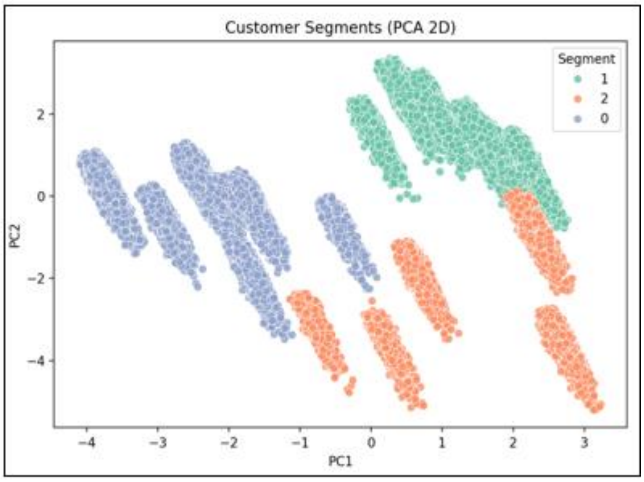

Figure 5: Customer Segment PCA 2D

Further analyses validated that promotion and media type combinations drove the greatest variance in CAC. This is seen in the drastic difference in variance between the combination of “Pick Your Savings” Promotion with “Radio” media type compared to “Big Promo” promotion with “Sunday paper, radio and TV” as media types (fig. 6).

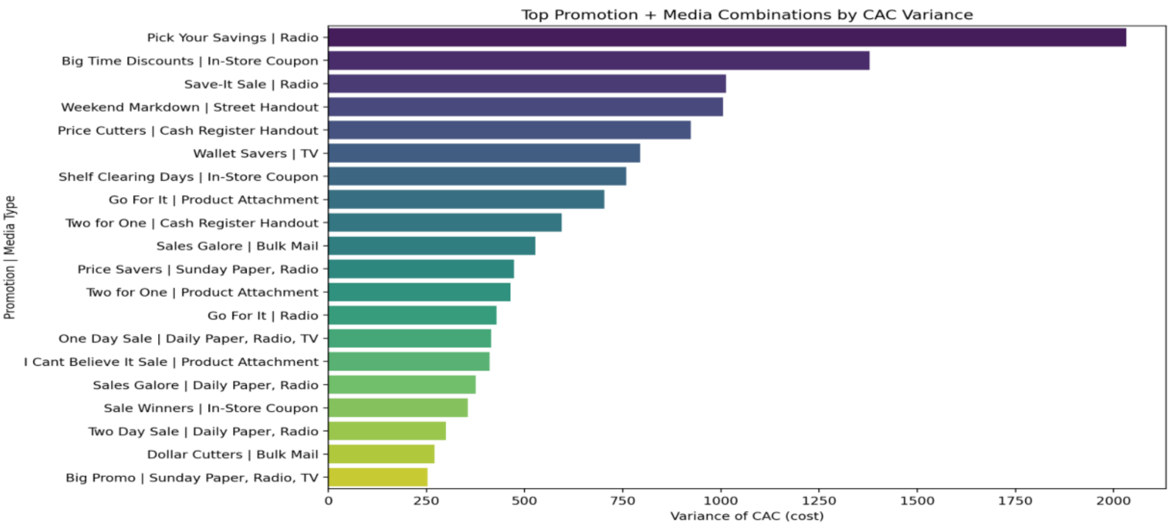

Figure 6: Top Promotion + Media Combination by CAC Variance

Promotions such as “Price Slashers,” “Wallet Savers,” and “Two for One” produced the lowest CAC values (fig. 7), while “Dollar Days” resulted in the highest, signaling inefficiency (fig. 8).

International Journal of Science and Research (IJSR)
ISSN: 2319-7064
Impact Factor 2024: 7.101

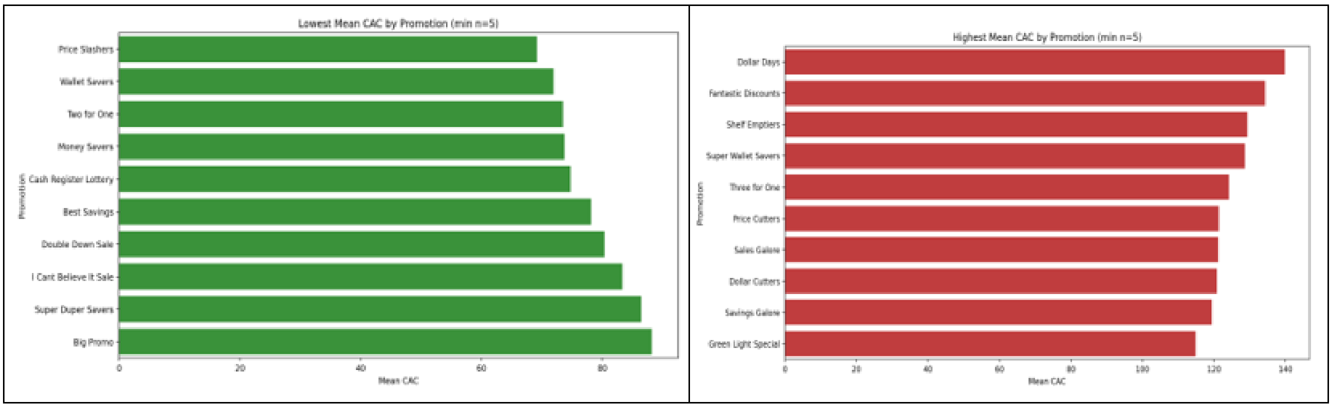

Figure 7 and 8: Highest and Lowest Mean CAC by Promotion

Media type used for advertising and promotions heavily impacts the customer acquisition cost. Certain media is more economical as it can target large groups for cheaper, however other types of media is necessary for convincing certain groups of people about a store's offering. In the case of the supermarket chain (Fig.9), TV and Daily Paper channels deliver the lowest customer acquisition costs on average, with mixed-media bundles that include TV (e.g., Daily Paper + Radio + TV) also performing efficiently; in contrast, in-store handouts (Cash Register Handout, Street Handout) and certain print bundles (e.g., Sunday Paper + Radio) are the most expensive. Mean and median CAC tell a consistent story, though higher spreads in some media types suggest outliers and variability. Practically, this points to prioritizing TV-led and Daily Paper placements for cost efficiency, while reassessing or tightening targeting for handout-based tactics that are driving up CAC.

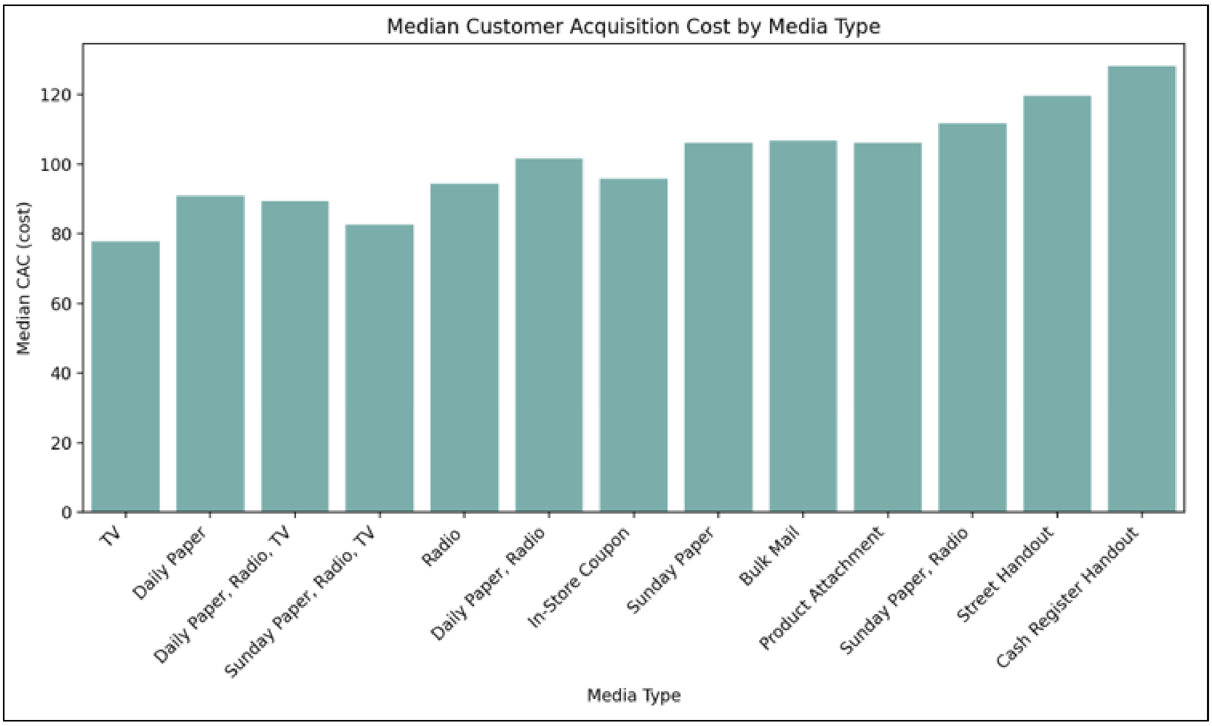

Figure 9: Median Customer Acquisition Cost by Media Type

At the store-format level, Deluxe Supermarkets consistently yielded the most efficient acquisition at $94.49 with the largest customer base, while small grocery stores suffered from the highest CAC at $112.13 despite limited reach (fig.10). Geographic patterns reinforced these dynamics, with Salem demonstrating the lowest CAC at $84.43, while Portland, though more expensive at $105.31, maintained profitability due to scale (fig. 11). At the country level, the U.S. achieved the lowest CAC at $98.40, compared with higher but still stable values in Canada and more variability in Mexico.

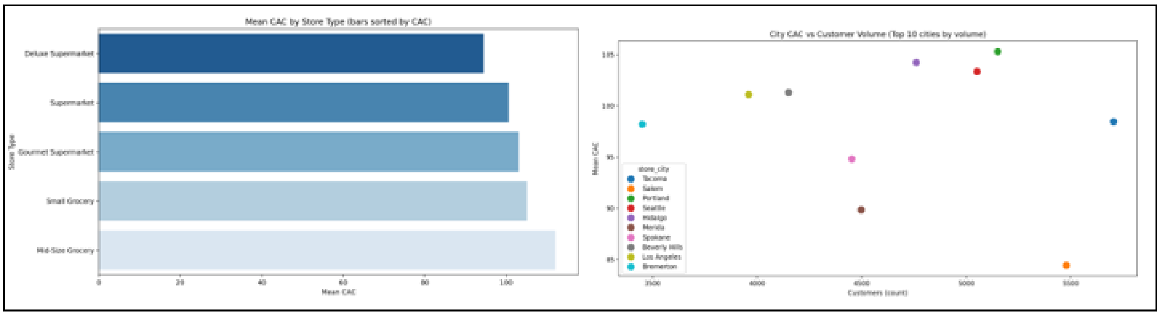

Figure 10 and 11: *CAC By Store Type and City*

To make sure results weren't biased by factors like seasonality or product mix, we compared stores of similar size within the same metro areas. Outlier trimming and percentile-based robustness checks were applied to prevent distortions from extreme observations. Amenity interactions were also evaluated, with combinations such as coffee bars and prepared foods showing incremental reductions in CAC beyond individual effects (fig. 12).

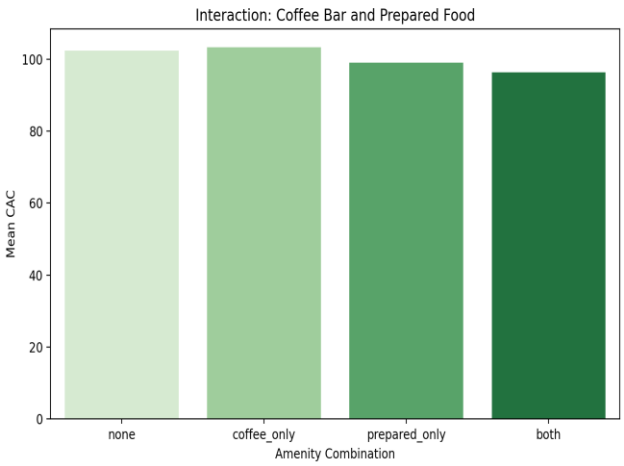

Figure 12: Interaction of Coffee bar and prepared Foods Amenities

Notably, video store amenities reduced CAC by nearly 7%, (fig. 13) and medium-sized stores consistently achieved the best balance between acquisition efficiency and operational performance.

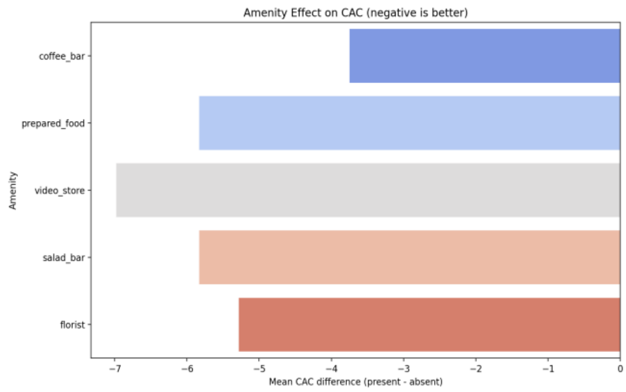

Figure 13: Amenity Effect on CAC

The number of amenities also leads to an eventual absolute reduction in customer acquisition cost on average, with stores having all 5 amenities seeing the lowest customer acquisition cost while those with no amenities have the maximum (fig. 14).

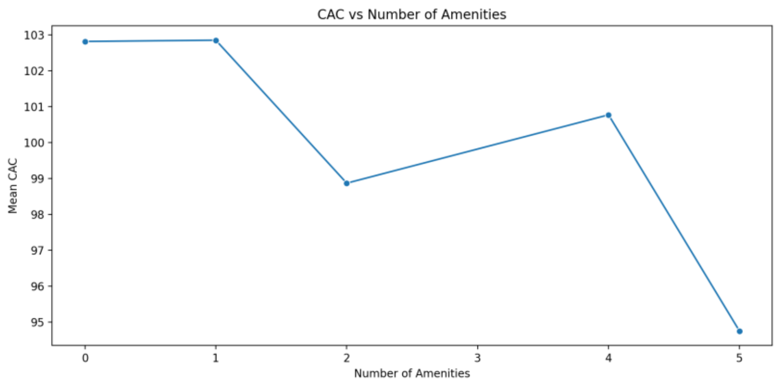

Figure 14: CAC vs Number of Ammenities

Distributional analysis of Deluxe Supermarket CAC demonstrated a near-normal bell curve (Fig. 16), with a mean of $94.49 and standard deviation of $30.36. Approximately 68% of customers fell within $64–125, and 95% within $34–155, indicating high predictability and enabling precise forecasting and budgeting of acquisition costs.

International Journal of Science and Research (IJSR)
ISSN: 2319-7064
Impact Factor 2024: 7.101

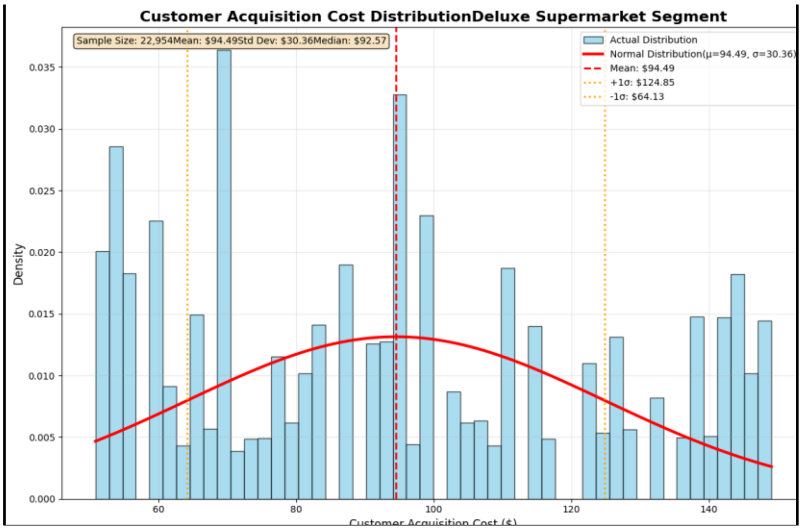

Figure 15: Customer Acquisition Cost Distribution

Synthesizing these results into business format strategies, small stores demonstrated the highest CAC efficiency and fastest payback, benefiting most from amenities such as coffee and prepared food and from media investments in newspapers and radio. Medium stores offered the strongest balance between cost efficiency and scalability, with selective amenities and mixed-media strategies proving optimal. Large stores, while less efficient on CAC, delivered the highest absolute EBITDA contributions when supported by complete amenity bundles and heavier TV spending. Financial modeling of revenue per square foot, gross margin rates, and operating expenses confirmed these trade-offs (fig. 16), suggesting a sequencing strategy: seed new markets with small formats for efficiency, anchor growth with medium stores for capital effectiveness, and selectively invest in large flagship stores for destination-driven revenue.

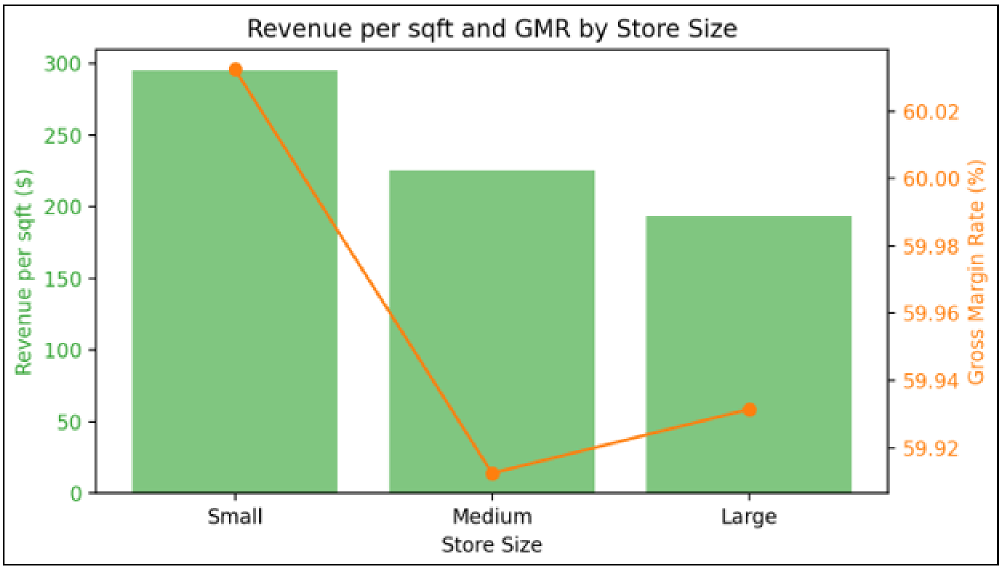

Figure 16: Revenue per Sqft and GMR by Store Size

Overall, the integrated analysis demonstrates that promotional design and media allocation are the most powerful levers for optimizing customer acquisition costs. Store size and amenity configurations play important supporting roles, with medium-sized stores offering the best capital efficiency, small stores providing superior acquisition efficiency, and large stores maximizing revenue density. Geographic targeting further enhances optimization, with U.S. operations offering the most favorable CAC conditions. By linking predictive modeling with engineered financial metrics, the study provides a robust, data-grounded framework for optimizing customer acquisition strategies across diverse store formats and markets.

Proposed business model for CAC optimisation

Our business plan, shown in Table 1 operationalizes the analysis through a modular decision architecture that links store format, location, amenities, and marketing strategy to financial outcomes. The model is built on four layers: (1) a data layer aggregating store attributes, market context, and campaign variables; (2) a feature layer computing revenue density, gross margin, CAC, and promotional/media lift by format and city; (3) a rules layer encoding which promotions and media channels minimize CAC or maximize lift in specific regions; and (4) an optimization layer dynamically allocating spend across channels and markets subject to CAC and ROI thresholds. This framework ensures that format-specific playbooks for small, medium, and large stores, as well as location plans, are grounded in the metrics surfaced by the analysis—namely, revenue per square foot, CAC ranges, and promotion/media lift patterns across markets.

Table 1: Grocery Store Business Model Comparison Table

Category	Small-Format Grocery (15–25k ft²)	Medium-Format Grocery (25–35k ft²)	Large-Format Grocery (≥35k ft²)
Annual Sales	$2.0–$3.4M	$5.7–$9.2M	$9.8–$14M+
Revenue Density	≈$135/ft²	≈$198/ft²	≈$279/ft²
EBITDA Margin	~8%	9–10%	11–12%
Capex & Payback	Low capex, quicker payback	Strongest payback efficiency	High capex ($5–12M), 6–7 yr payback
Key Amenities	Coffee bar, prepared-food counter	Coffee, prepared foods, salad bar, florist	Full suite: coffee, salad, prepared foods, florist, specialty counters
Best-Performing Categories	Frozen meals, breakfast items, ice cream	Fresh, pantry, frozen, dairy	Broad assortment with specialty
Operating Model	Self-checkout, simplified layouts, 20% opex	Scale-driven leverage, subregional siting	Regional nodes, freeway access, commissary support
Marketing Mix	60–70% daily paper + radio, tactical promos	50/30/20 split: paper/radio/TV	40/35/25 split: TV/paper/radio
Effective Promotions	Pick-Your-Savings, Save-It Sale, Weekend Markdown	Coupon Spectacular, Best Savings	Best Savings, Double-Down Sale
Personalization	Value: coupons; Commuters/students: ready-to-eat; Families: frozen bundles	Families: meal bundles; Young professionals: salad/prepared; Value-seekers: coupons	High-value: stock-up & occasion shoppers; specialty counters, holiday
Location Strategy	Dense neighborhoods, commuter corridors	Subregional hubs, first-ring suburbs	Only in regional nodes, destination sites
Clustering / Anchoring	4–6 stores per metro for efficiency	Anchor 4–5 small stores once $8M sales reached	Build only after ≥2 medium anchors exist

Dynamic Marketing Allocation and Personalization Across Formats

The optimization layer continuously reallocates spend based on observed CAC and sales lift by channel and promotion. Constraints include diminishing returns caps, minimum brand presence, and saturation limits by city. As seen in Table 1, small stores bias toward daily paper and radio; medium stores maintain a balanced trio and flex weekly toward the most efficient channel; large stores ride heavier TV during event peaks but rebalance when CAC spikes. Customer segmentation translates into actionable creative strategies: value-seekers receive coupon and price-led offers, convenience-seekers receive coffee/prepared-foods campaigns, families are targeted with frozen and bundle offers, and occasion buyers with florist and specialty promotions. Each segment is served through the channels where the analysis showed the strongest lift for that profile.

Comparative Insights and Rollout Logic

The integrated evidence shows that small stores are the most CAC-efficient and quickest to build, ideal for seeding new markets and generating early cash flow. Medium stores deliver the best capital-weighted returns and shortest payback at scale, making them the portfolio's growth engine and anchors for local clusters. Large stores generate the highest absolute EBITDA and strong brand halo effects but require superior sites and longer paybacks, limiting them to selective deployment. The recommended sequence is therefore: seed metros with clusters of small stores, add medium anchors once cumulative sales cross ≈$8 million, and introduce large flagships only in dense, high-income markets with mature logistics. This plan translates the model's quantitative findings—CAC ranges ($52–62 small; $51–60 medium; mid- to high-$50s large), revenue-per-square-foot benchmarks ($135, $198, $279), promotion and media efficiency rankings, and format-level financial models—into a coherent, rules-based expansion strategy optimized for both profitability and risk management.

3. Conclusion and Future Work

This research shows that predictive and prescriptive analytics can change how businesses handle Customer Acquisition Cost (CAC). By using a detailed dataset of store features, promotions, media channels, and customer demographics, the study illustrates how machine learning models can pinpoint the main factors that drive acquisition efficiency and turn those into usable strategies. The analysis found that optimizing promotional design and media mix are the most effective ways to lower CAC, while store format and amenity setup offer valuable advantages. In practical terms, small-format stores turned out to be the most cost-effective for acquiring customers. Medium formats provided a good mix of scalability and efficiency, while large formats yielded the highest revenue density and EBITDA when combined with complete amenity packages. These insights were organized into a business model framework that connects predictive insights with actionable marketing and expansion strategies. This gives managers and entrepreneurs a structured guide for decision-making.

What makes this study valuable is its practical application, not merely its predictive capacity. The framework demonstrates how businesses can better manage budgets, tailor promotional strategies, and plan store rollouts to achieve profitability and sustainable growth by integrating analytics with business model design. Accurately predicting CAC can help startups and small businesses avoid overspending, increase marketing returns, and accelerate growth in resource-constrained scenarios. The framework provides a means for larger organizations to optimize acquisition costs across various customer segments and markets.

There are some restrictions, though. The findings' applicability to other areas, sectors, and retail models is limited because the dataset was sourced from a single North American supermarket chain. Furthermore, there is some estimation bias because CAC was calculated using proxies based on store expenses and sales rather than real marketing expenditures. The model's capacity to adjust to abrupt changes in market conditions is hampered by the incomplete accounting of larger external factors like seasonality, competition, and macroeconomic shifts.

These results can be built upon in a number of important ways in future studies. A more comprehensive view of acquisition effectiveness and long-term profitability would be provided by including customer lifetime value (LTV) in addition to CAC. Developing real-time CAC tracking systems that integrate offline and digital marketing data could assist companies in continuously modifying their tactics rather than depending on static reports. Creating AI applications that are easy to use for small businesses is another possible avenue. With the help of these tools, business owners could enter campaign and customer data to receive customized promotional strategies, suggested media expenditures, and predictive CAC forecasts. With this strategy, businesses that can't currently afford sophisticated data science teams would have access to advanced analytics.

The real-world effects of this work are important. For startups, more accurate CAC predictions can protect against cash flow problems. For larger companies, it can drive efficiencies at scale. Collaborations between researchers, technology providers, and industry players could bring these models into real retail settings, confirming their usefulness and refining their predictive abilities through operational feedback. Ultimately, this study shows that predicting CAC is not just a theoretical task; it is a key business strategy. With the help of AI, it offers companies a route to better profitability, faster growth, and increased resilience in competitive markets.

Glossary/Appendix

1) **Data Sources & Preparation**
   a) **Primary dataset:** customers.xlsx – includes sales, costs, units, store size, promotions, media, and demographics.
   b) **Data cleaning:** Standardized column names, confirmed presence of key fields, and excluded rows with missing/invalid values (e.g., zero store sales or sqft)

2) **Key Metrics & Formulas**
   a) **Revenue per square foot (RPSF):**
   - Formula: (store_sales × 1,000,000) / store_sqft
   - Used to compare store types and locations.

   b) **Gross Margin Rate (GMR):**
   - Formula: (store_sales − store_cost) / store_sales
   - Used for comparing profitability by geography and format.

   c) **Customer Acquisition Cost (CAC proxy):**
   - Formula: (store_cost × 1,000,000) / max (unit_sales × 1,000, 1)
   - Used as a relative efficiency index since actual ad spend data was not available.

3) **Analysis Methods**
   a) **Aggregation rules:**
   - Averages (means) used for grouped comparisons.
   - Sums used for totals across time or regions.
   - Medians checked for stability but not primarily reported.

   b) **Outliers:** Not capped; only inspected.
   c) **Promotion effectiveness:** Compared average CAC proxy and sales lift (extra sales generated during a promotion or media campaign compared to a similar period without it) by promotion type.
   d) **Media effectiveness:** Compared CAC proxy, sales, and margins by media channel.
   e) **Store format & location recommendations:** Based on ranking of RPSF and GMR, with geographic overlays to identify stronger markets.

4) **Data Quality**
   a) Excluded rows with missing/invalid denominators (e.g., zero sqft, zero sales, or missing unit sales).
   b) Verified detection of core fields (sales, cost, sqft).
   c) Corrected initial visualization errors; final charts were based on clean aggregations.

5) **External Assumptions**
   a) CAC proxy is **relative only**, no actual marketing spend was imported.
   b) All financial data was in **millions**, scaled where needed.
   c) No inflation, rent, labor, or external benchmarks applied.
   d) Assumed standard retail relationships: higher RPSF and GMR = stronger unit economics.

References

[1] Ababyomi, A. A., Ogeawuchi, J. C., Onifade, A., & Agboola, O. A. (n.d.). Systematic Review of Marketing Attribution Techniques for Omnichannel Customer Acquisition Models. *International Journal of Advanced Multidisciplinary Research and Studies 3(6):1621-1633.*

[2] Ashik, K. (2023). *Customer Acquisition Cost Analysis using Python*. Kaggle. Retrieved 10 13, 2024, from https://www.kaggle.com/code/dkhalidashik/customer-acquisition-cost-analysis-using-python/code

[3] Long, B. (2025, January 24). *How to Calculate Unit Economics for Startups (LTV, CAC & More)*. Kruze Consulting. Retrieved September 13, 2025, from https://kruzeconsulting.com/blog/unit-economics/

[4] Nair, R. (2025, 7 10). *Why Is Marketing Attribution Broken?* Lifesight. https://lifesight.io/blog/marketing-attribution-problems/

[5] Rosler, D. (2024, January 2). *LTV:CAC - An Important (But Often Misunderstood) SaaS Metric*. Burkland Associates. Retrieved 2025, from https://burklandassociates.com/2024/01/02/ltvcac-an-important-but-often-misunderstood-saas-metric/

[6] Stewart, B. (2019, June 11). *EVA, Not EBITDA: A Better Measure of Investment Value*. The Harvard Law School Forum on Corporate Governance. Retrieved September 13, 2025, from https://corpgov.law.harvard.edu/2019/06/11/eva-not-ebitda-a-better-measure-of-investment-value/