Guide

6 Strategies for Building Personalized Customer Experiences

databricks

Contents

Introduction

In today's experience-driven world, the most beloved brands are the ones that know their customers. Customers are loyal to brands that recognize their needs and preferences — and tailor user journeys and engagements accordingly.

A study from McKinsey shows 76% of consumers are more likely to consider buying from a brand that personalizes the shopping and user experience to the wants and needs of the customer. And as organizations pursue omnichannel excellence, these same high expectations of online experiences also extend to brick-and-mortar locations — revealing for many merchants that personalized engagement is fundamental to attracting customers and expanding share of wallet.

But achieving a 360-degree view of your customers to serve personalized experiences requires integrating various types of data — including demographics, behavioral and transactional — to develop robust profiles. This guide focuses on six actionable strategic pillars for businesses to leverage automation, real-time data, AI-driven analysis and well-tuned ML models to architect and deliver customized customer experiences at every touch point.

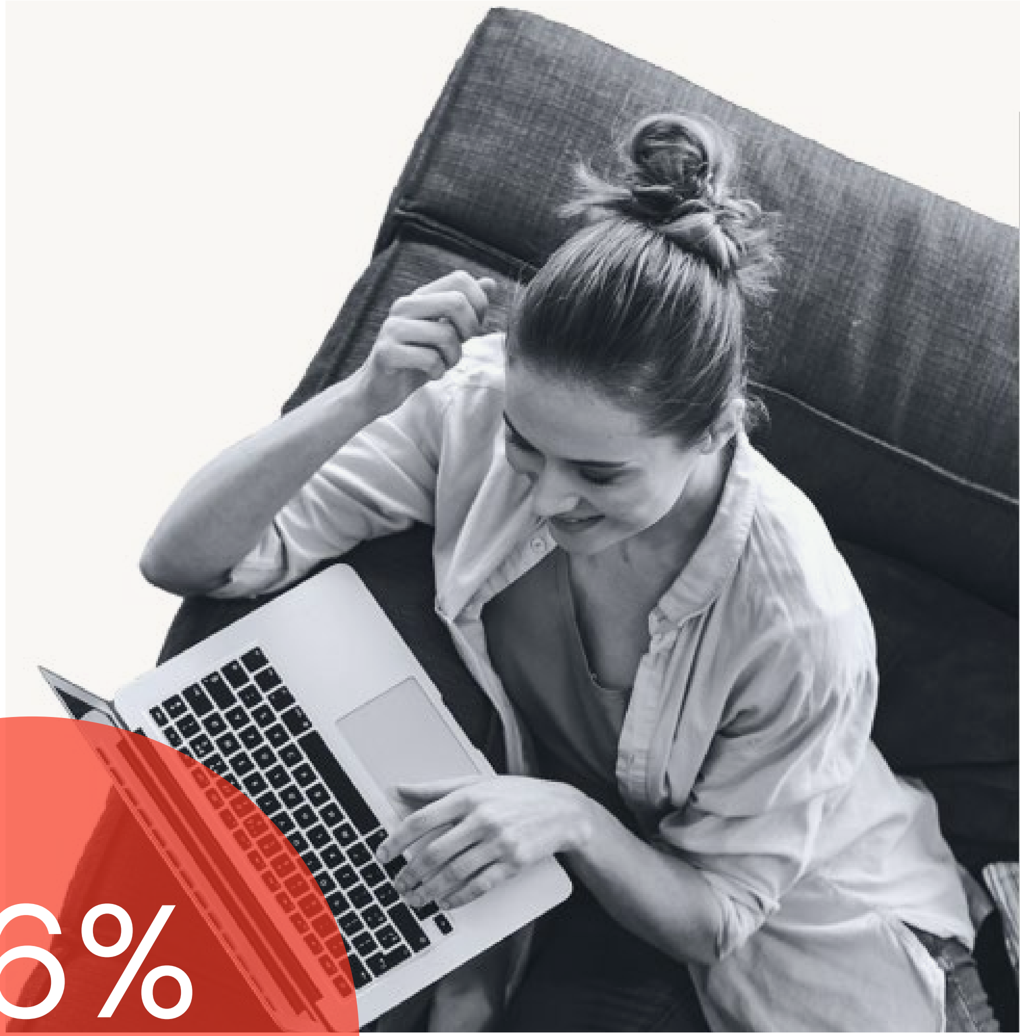

76%

of consumers are more likely to purchase due to personalization

McKinsey & Company

Building a Foundation for Personalization

Get a 360-degree view of the customer by leveraging ML-based entity resolution

To create truly personalized interactions, you need actionable insights about your customers. Start by establishing a common customer profile and accurately linking together customer records across disparate data sets.

Get a 360-degree view of your target customer by bringing together:
- Sales and traffic-driven first-party data
- Product ratings and surveys
- Customer surveys and support center calls
- Third-party data purchased from data aggregators and online trackers
- Zero-party data provided by customers themselves

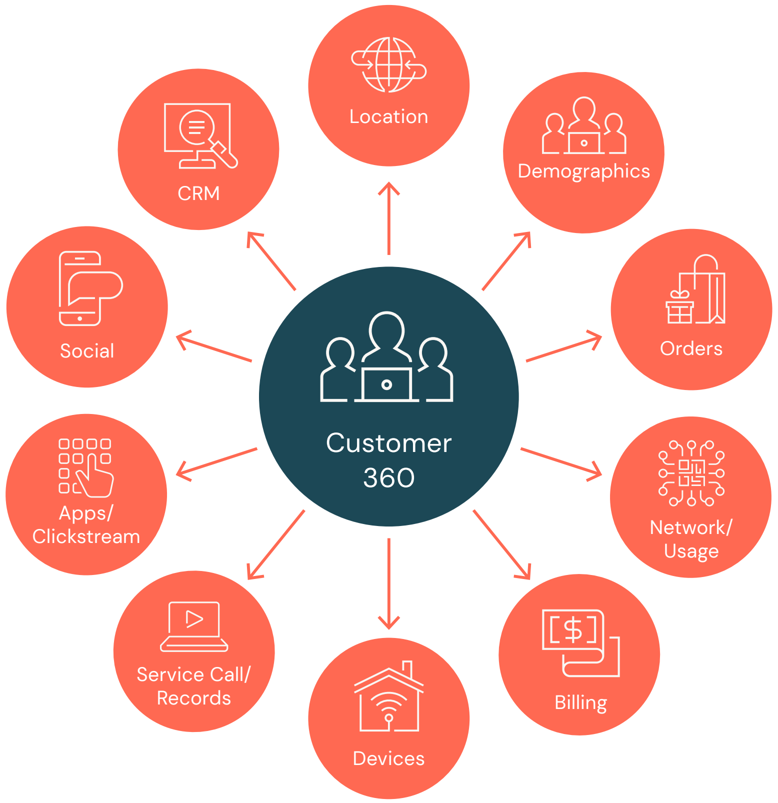

CASE STUDY

Personalizing experiences with data and ML

Grab is the largest online-to-offline platform in Southeast Asia and has generated over 6 billion transactions for transport, food and grocery delivery, and digital payments. Grab uses Databricks to create sophisticated customer segmentation and recommendation engines that can now ingest and optimize thousands of user-generated signals and data sources simultaneously, enhancing data integrity and security, and reducing weeks of work to only hours.

Get the full story →

> "The C360 platform empowered teams to create consumer features at scale, which in turn allows for these features to be extended to other markets and used by other teams. This helps to reduce the engineering overhead and costs exponentially."

NIKHIL DWARAKANATH
Head of Analytics, Grab

Given the different data sources and data types, automated matching can still be incredibly challenging due to inconsistent formats, misinterpretation of data, and entry errors across various systems. And even if inconsistent, all that data may be perfectly valid — but to accurately connect the millions of customer identities most retailers manage, businesses must lean on automation.

In a machine learning (ML) approach to entity resolution, text attributes like name, address and phone number are translated into numerical representations that can be used to quantify the degree of similarity between any two attribute values. But your ability to train such a model depends on your access to accurately labeled training data. It's a time-consuming exercise, but if done right, the model learns to reflect the judgments of the human reviewers.

Many organizations rely on libraries encapsulating this knowledge to build their applications and workflows. One such library is Zingg, an open source library bringing together ML-based approaches to intelligent candidate pair generation and pair-scoring. Oriented toward the construction of custom workflows, Zingg presents these capabilities within the context of commonly employed steps such as training data label assignment, model training, data set deduplication, and (cross-data set) record matching.

Built as a native Apache Spark™ application, Zingg scales well to apply these techniques to enterprise-sized data sets. Organizations can then use Zingg in combination with platforms such as Databricks Lakehouse to provide the back end to human-in-the-middle workflow applications that automate the bulk of the entity resolution work and present data experts with a more manageable set of edge case pairs to interpret.

As an active-learning solution, models can be retrained to take advantage of this additional human input to improve future predictions and further reduce the number of cases requiring expert review. Finally, these technologies can be assembled to enable their own enterprise-scaled customer entity resolution workflow applications.

Need help building your foundation for a 360-degree view of your customers?

Get pre-built code sample data and step-by-step instructions in a Databricks notebook in the **Customer Entity Resolution Solution Accelerator.**

- Translating text attributes (like name, address, phone number) into quantifiable numerical representations
- Training ML models to determine if these numerical labels form a match
- Scoring the confidence of each match

GET THE ACCELERATOR >

Estimating Customer Lifetime Value

Building brand loyalty to drive share of wallet with data

Once you've set up a 360-degree view of the customer, the next challenge is how to spend money to profitably grow the brand. The goal is to spend marketing dollars on activities that attract loyal customers and avoid spending on unprofitable customers or activities that damage the brand. Keep in mind, that making decisions solely based on ROI isn't the answer. This one-track approach could ultimately weaken your brand equity and make you more dependent on lowering your price through promotions as a way to generate sales.

KOLIBRI Games

CASE STUDY

Predicting and increasing customer lifetime value with ML

Kolibri Games, creators of Idle Miner Tycoon and Idle Factory Tycoon, attracts over 10 million monthly active users. With Databricks, they achieved a 30% increase in player LTV, improved data team productivity by 3x, and reduced ML model-to-production time by 40x.

Get the full story →

Within your existing customer base are people ranging from brand loyalists to brand transients. Brand loyalists are highly engaged with your brand, are willing to share their experience with others, and are the most likely to purchase again. Brand transients have no loyalty to your brand and shop based on price. Your focus should be on growing the group of brand loyalists while minimizing interactions with brand transients.

Identifying and engaging brand loyalists

Today's customer has overwhelmingly abundant options in products and services to choose from. That's why personalizing customer experiences is so important, as it increases revenue, marketing efficiency and customer retention.

Not every customer carries the same potential for profitability. Different customers derive different value from your products and services, which directly translates into differences in the overall amount of value a business can expect in return. Mutually beneficial relationships carefully align customer acquisition cost (CAC) and retention rates with the total revenue or customer lifetime value (CLV).

Calculating customers' lifetime intent

To assess the remaining lifetime in a customer relationship, businesses must carefully examine the transactional signals and other indicators from previous customer engagements and transactions.

For example, if a frequent customer slows down their buying habits — or simply doesn't make a purchase for an extended period of time — it may signal the upcoming end of the relationship. However, in the case of another customer who engages infrequently, the same extended absence may not signal anything notable. The infrequent buyer may continue to purchase even after a long pause in activity.

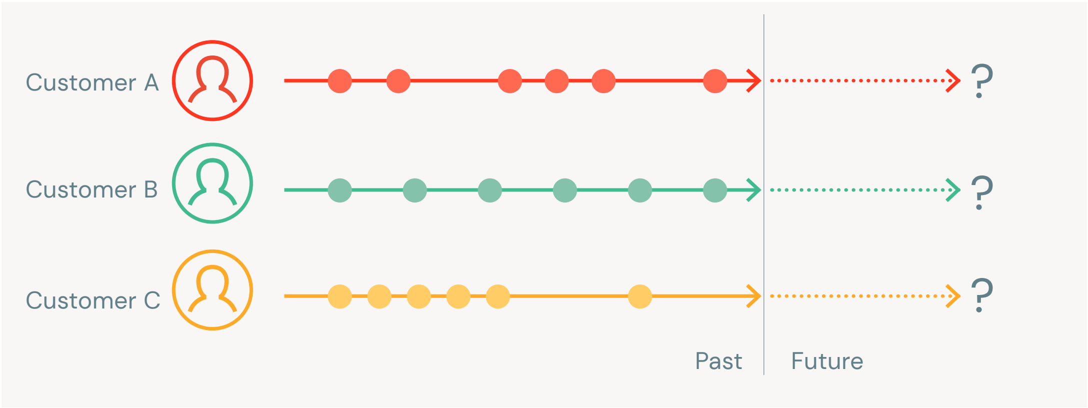

Different customers with the same number of transactions, but signaling different lifetime intent.

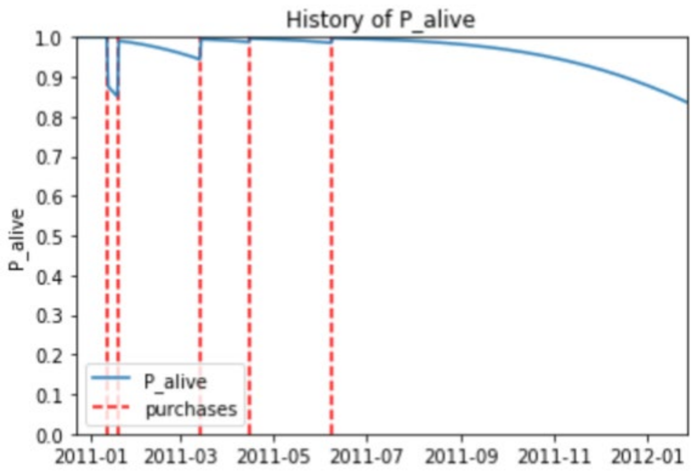

The probability of re-engagement (P_alive) relative to a customer's history of purchases.

Every customer relationship with a business has a lifespan. Understanding what point in the lifespan at a given time provides critical insight to inform marketing and sales tactics. By proactively discovering shifts in the relationship, you can adapt how to respond to each customer at the optimal time. For example, a certain signal might prompt a change in how to deliver products and services, which could help maximize revenue.

Transactional signals can be used to estimate the probability that a customer is active and likely to return in the future. Popularized as the Buy 'til You Die (BTYD) model, analysts can compare a customer's frequency and recency of engagement to similar patterns across their user population to accurately predict individual CLV.

The mathematics behind these predictive CLV models is complex, but the logic behind these critical models is accessible through a popular Python library named Lifetimes, which allows the input of simple summary metrics in order to derive customer-specific lifetime estimates.

CASE STUDY

PUBLICIS GROUPE

How personalized experiences keep customers coming back for more

Publicis Groupe empowers brands to transform retail experiences with digital technologies, but data challenges and team silos stood in the way of delivering the personalization that their customers required. See how they use Databricks to create a single customer view that allows them to drive customer loyalty and retention. As a result, they've seen a 45%–50% increase in customer campaign revenue.

Get the full story →

Delivering customer lifetime estimates to the business

Using the Lifetimes library to calculate customer-specific probabilities at speed and scale can be challenging — from processing large volumes of transaction data to deriving data curves and value distribution patterns and, eventually, to integration with business initiatives. But with the proper approach, you can resolve all of them.

These models depend on three key per customer metrics:

FREQUENCY
The number of times within a given time period in which a repeat transaction is observed

AGE
The length of time between the occurrence of an initial transaction to the end of a given time period

RECENCY
The "age" of a customer (how long they've engaged with a brand) at the time of their latest repeat transaction

Spark natively distributes this work across a multi-server environment, enabling consistent, accurate and efficient analysis. Spark's flexibility allows models to adapt in real time as new information is ingested, eliminating the bottlenecks that come with manual data mapping and profile building.

With per customer metrics calculated, the Lifetimes library can be used to train multiple BTYD models, such as Pareto/NBD and BG/NBD. Training models to predict engagements over time using proprietary data can take several months and thousands of training runs. Hyperopt, a specialized snippet library, helps businesses tap into the infrastructure behind their Spark environments and distribute the training outputs across models.

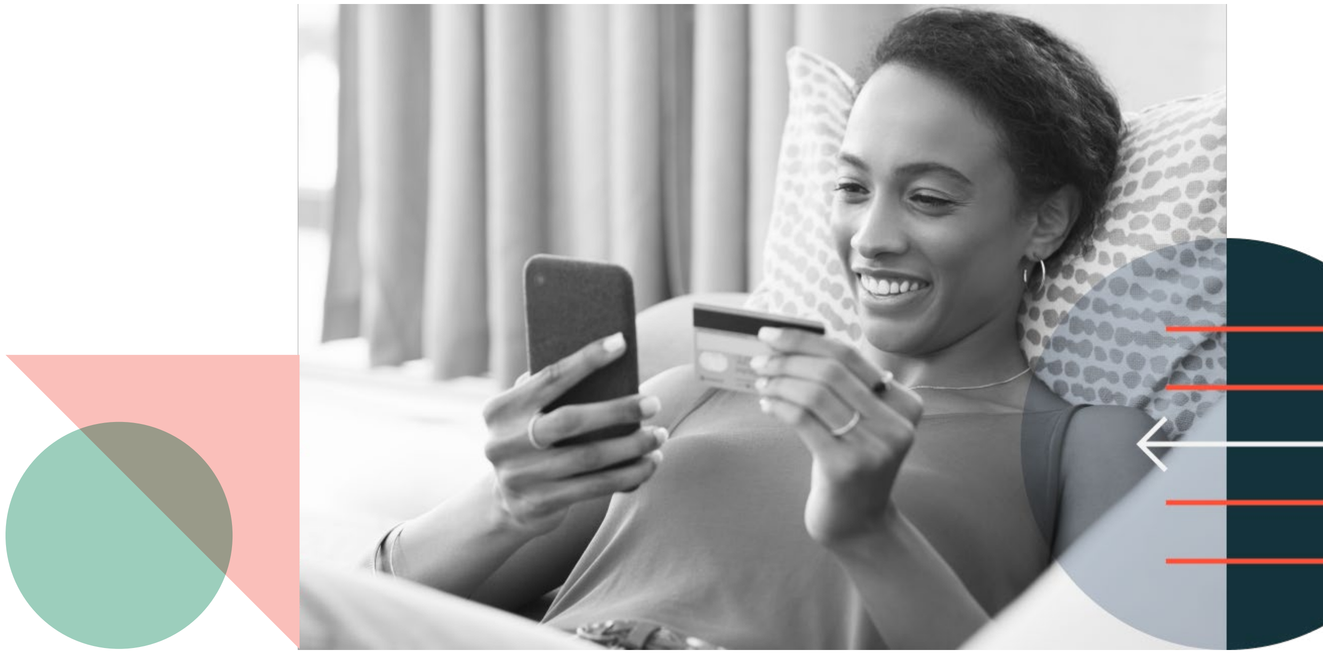

Solution deployment

Once properly trained, these models can determine the probability that a customer will re-engage, as well as the number of engagements a business can expect from that customer over time. But the real challenge is putting these predictive capabilities into the hands of those that determine customer engagement.

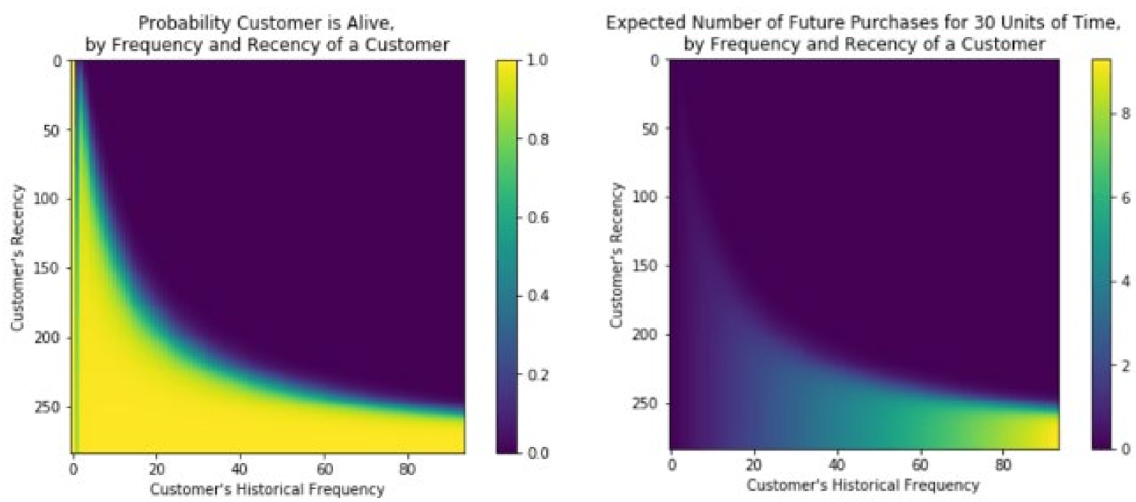

Matrices illustrating the probability a customer is alive (left) and the number of future purchases in a 30-day window given a customer's frequency and recency metrics (right).

Businesses need a way to develop and deploy solutions in a highly scalable environment with a limited upfront cost. Databricks Solution Accelerators leverage real-world sample data sets and pre-built code to show how raw data can be transformed into real solutions — including step-by-step instructions ready to go in a Databricks notebook.

Need help determining your customers' lifetime value?

Use the **Customer Lifetime Value Accelerator** to

- Ingest sample retail data
- Use pre-built code to develop visualizations and explore past purchase behavior
- Apply machine learning to predict the likelihood and nature of future purchases

GET THE ACCELERATOR >

Mitigating Customer Churn

Balancing acquisition and retention with personalized experiences

There are no guarantees of success. With a bevy of options at their disposal, customer churn is a reality that companies face and are focused on overcoming every day. One recent analysis of consumer-oriented subscription services estimated a segment average 7.2% monthly rate of churn. When narrowed to brands focused on consumer goods, that rate jumped to 10.0%. This figure translates to a lifetime of 10 months for the average subscription box service, leaving businesses of this kind with little time to recover acquisition costs and bring subscribers to net profitability.

Riot Games

CASE STUDY

Creating an optimal in-game experience for League of Legends

Riot Games is one of the top PC game developers in the world, with over 100 million monthly active users, 500 billion data points, and over 26 petabytes of data and counting. They turned to Databricks to build a more efficient and scalable way to leverage data and improve the overall gaming experience — ensuring customer engagement and reducing churn.

Get the full story →

Organizations must take an honest look at the cost of acquisition relative to a customer's lifetime value (LTV) earned. These figures need to be brought into a healthy balance and treated as a "chronic condition" to be managed.

Understanding attrition predictability through subscriptions: Examining retention-based acquisition variables

Public data for subscription services is extremely hard to come by. KKBox, a Taiwan-based music streaming service, recently released over two years of anonymized subscription data to examine customer churn. Through analyzing the data, we uncover customer dynamics familiar to any subscription provider.

Most subscribers join the KKBox service through a 30-day trial offer. Customers then appear to enlist in one-year subscriptions, which provide the service with a steady flow of revenue. Subscribers typically churn at the end of the 30-day trial and at regular one-year intervals.

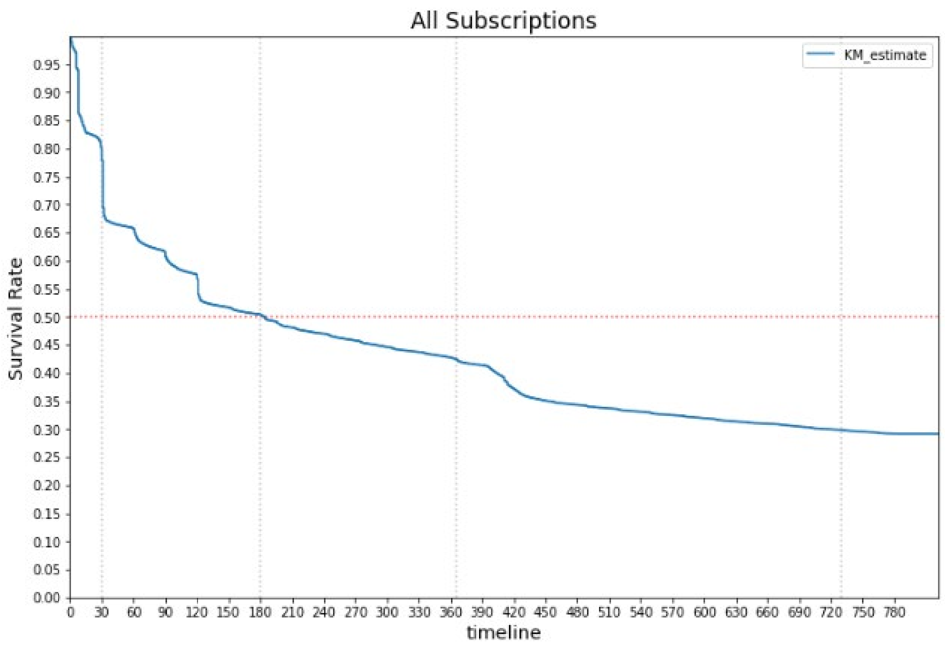

The Survival Rate reflects the proportion of the initial (Day 1) subscriber population that is retained over time, first at the roll-to-pay milestone, and then at the renewal milestone.

This pattern of high initial drop-off, followed by a period of slower but continuing drop-off cycles makes intuitive sense. Where it gets interesting is when the data changes. The patterns of customer churn become vastly different as time passes and new or changing elements are introduced (e.g., payment methods and options, membership tiers, etc.).

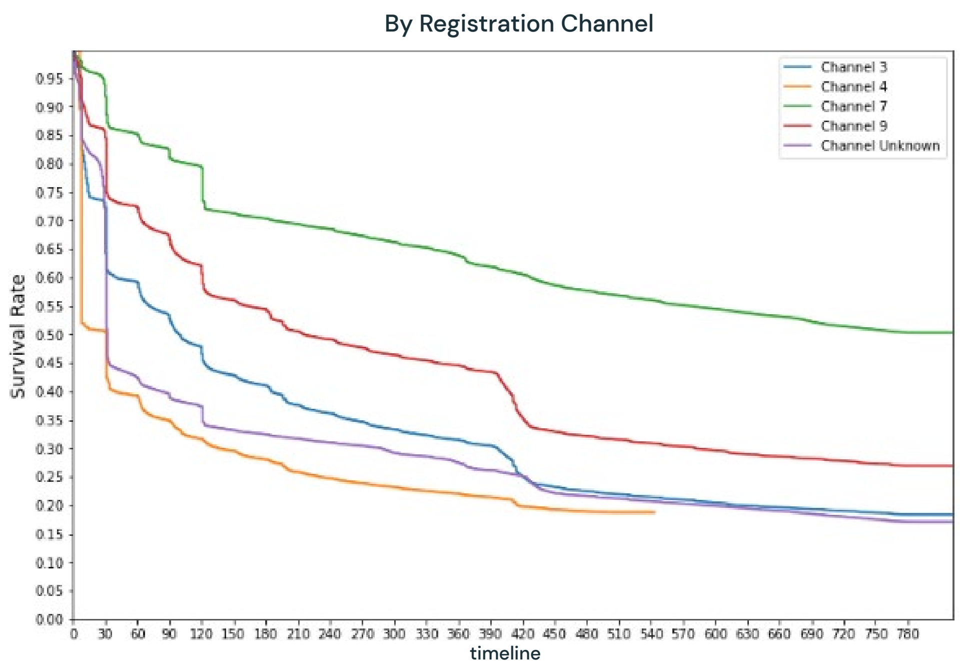

Customer attrition by subscription day on the KKBox streaming service for customers registering via different channels.

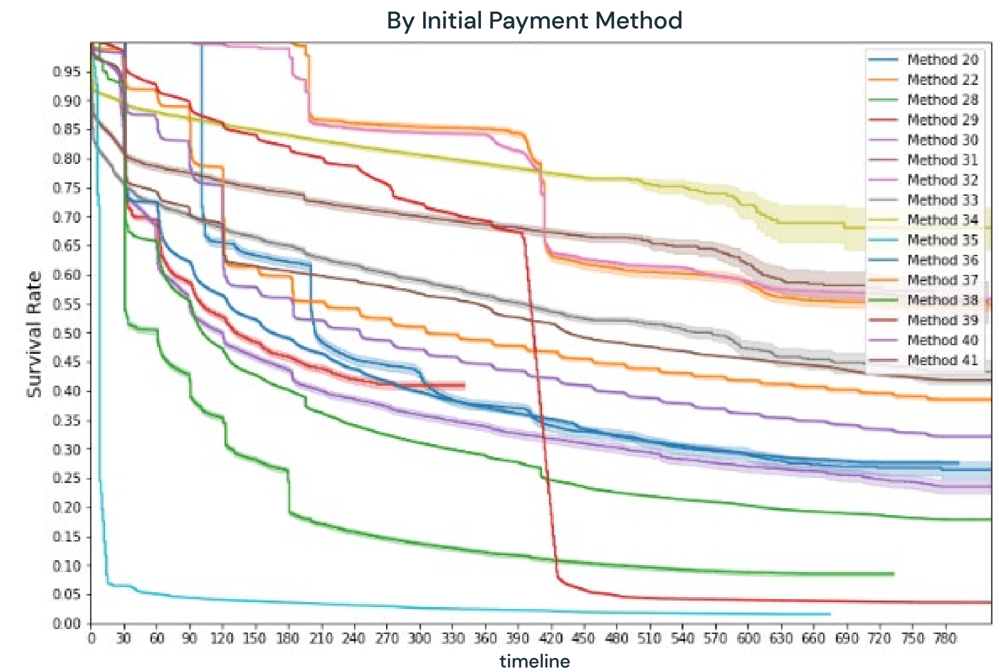

Customer attrition by subscription day on the KKBox streaming service for customers registering via different payment methods.

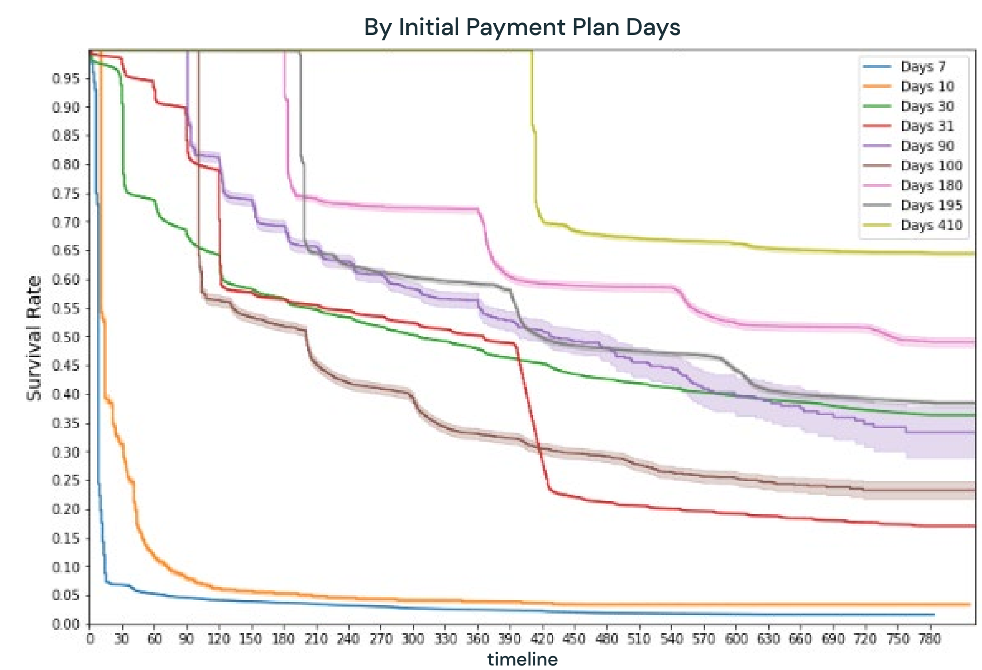

Customer attrition by subscription day on the KKBox streaming service for customers selecting different initial payment methods and terms/days.

Category	Feature	Factor
Registration Channel	channel_3	0.96
	channel_4	1.20
	channel_7	1.00
	channel_9	0.92
Initial Payment Method	method_20	5.15
	method_22	3.00
	method_28	5.31
	method_29	2.50
	method_30	2.27
	method_31	0.94
	method_32	2.89
	method_33	1.53
	method_34	0.57
	method_35	4.43
	method_36	2.30
	method_37	1.10
	method_38	3.32
	method_39	1.19
	method_40	1.32
	method_41	1.00

The channel and payment method multipliers combine to explain a customer's risk of attrition at various points in time. The higher the value, the higher the proportional risk of churn in the associated period.

These patterns seem to indicate that KKBox *could* potentially differentiate between customers based on their lifetime potential, using only the information available at subscriber acquisition. In the same way, non-subscription businesses could use similar data techniques to get an accurate illustration of the total lifetime value of a particular customer, even before collecting historical data.

This information can help businesses target certain shoppers with effective discounts or promotions as early as trial registration. Nevertheless, it's always important to consider more than individual data points.

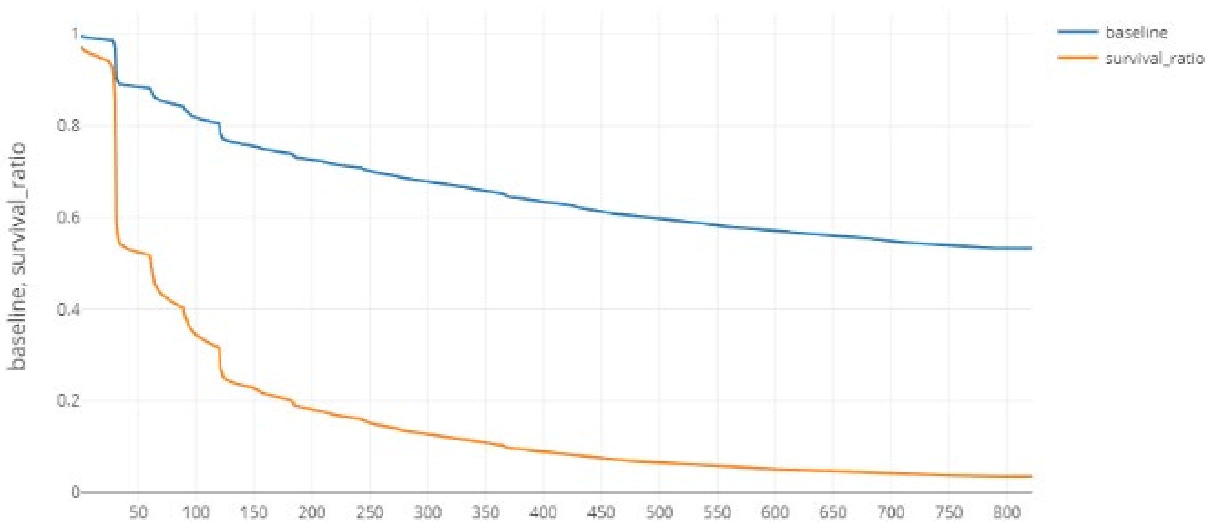

The baseline risk of customer attrition over a subscription lifespan.

Applying churn analytics to your data

This analysis is useful in two ways: **1)** to quantify the risk of customer churn and **2)** to paint a quantitative picture of the specific factors that explain that risk, giving analysts a clearer understanding of what to focus on, what to ignore and what to investigate further. The main challenge is organizing the input data.

The data required to examine customer attrition may be scattered across multiple systems, making an integrated analysis difficult. Data lakes support the creation of transparent, sustainable data processing pipelines that are flexible, scalable and highly cost-efficient. Remember that **churn is a chronic condition to be managed**, and attrition data should be periodically revisited to maintain alignment between acquisition and retention efforts.

Need help predicting customer churn?

Use the **Subscriber Churn Prediction Accelerator** to analyze behavioral data, identify subscribers with an increased risk of cancellation, and predict attrition. Machine learning lets you quantify a user's likelihood to churn, identifying factors that explain the risk.

GET THE ACCELERATOR >

Streamlining Customer Analysis and Targeting

Creating efficient and highly targeted customer experiences with behavioral data

Effective targeting comes down to one fundamental element: the cost of delivering a good or service relative to what a consumer is willing to pay.

In the earliest applications of segmentation, manufacturers recognized that specialized product lines targeting specific consumer groups could help brands stand out against competitors.

PANDORA

CASE STUDY

Finding that special something every time

Pandora is a jewelry company with global reach. They built their master consumer view (MCV) dashboard on the Databricks Lakehouse Platform, giving them the insights necessary to deliver highly targeted messaging and personalization — resulting in 80% growth in email marketing success, a 50% increase in click-to-open rate across 65 million emails, and 255M DKK (Danish Krone) in quarterly revenue.

Get the full story →

This mode of thinking extends beyond product development and into every customer-oriented business function, requiring specific means of ideation, production and delivery. The work put into segmentation doesn't need to be a gamble. Scrutinizing customers and testing responsiveness is an ongoing process. Organizations must analyze and adapt to shifting markets, changing consumer demand and evolving business objectives.

Bagelcode

CASE STUDY

Powering insight-driven dashboards to increase customer acquisition

Bagelcode is a global game company with more than 50 million global users. By using the Databricks Lakehouse Platform, they are now able to support more diversified indicators, such as a user's level of frequency and the amount of time they use a specific function for each game, enabling more well-informed responses. In addition, the company is mitigating customer churn by better predicting gamer behavior and providing personalized experiences at scale.

Get the full story →

> "Thanks to Databricks Lakehouse, we can support real-time business decision-making based on data analysis results that are automatically updated on an hourly and daily basis, even as data volumes have increased by nearly 1,000 times."
>
> JOOHYUN KIM
> Vice President, Data and AI, Bagelcode

A brand's goal with segmentation should be to define a shared customer perspective on customers, allowing the organization to engage users consistently and cohesively. But any adjustments to customer engagement require careful consideration of organizational change concerns.

reckitt

CASE STUDY

Responding to global demand shifts with ease

Reckitt produces some of the world's most recognizable and trusted consumer brands in hygiene, health and nutrition. With Databricks Lakehouse on Azure, they're able to meet the needs of billions of consumers worldwide by surfacing real-time, highly accurate, deep customer insights, leading to a better understanding of trends and demand, allowing them to provide best-in-class experiences in every market.

Get the full story →

A segmentation walk-through: Grocery chain promotions

A promotions management team for a large grocery chain is responsible for running a number of promotional campaigns, each of which is intended to drive greater overall sales. Today, these marketing campaigns include leaflets and coupons mailed to individual households, manufacturer coupon matching, in-store discounts and the stocking of various private-label alternatives to popular national brands.

Recognizing uneven response rates between households, the team is eager to determine if customers might be segmented based on their responsiveness to these promotions. They anticipate that such segmentation may allow the promotions management team to better target individual households, driving overall higher response rates for each promotional dollar spent.

Using historical data from point-of-sale systems along with campaign information from their promotions management systems, the team derives a number of features that capture the behavior of various households with regard to promotions. Applying standard data preparation techniques, the data is organized for analysis and using a variety of clustering algorithms, such as k-means and hierarchical clustering, the team settles on two potentially useful cluster designs.

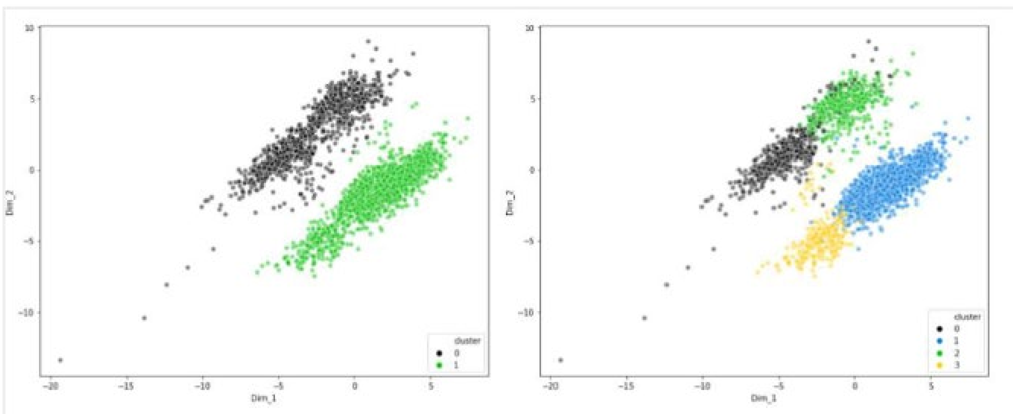

Overlapping segment designs separating households based on their responsiveness to various promotional offerings.

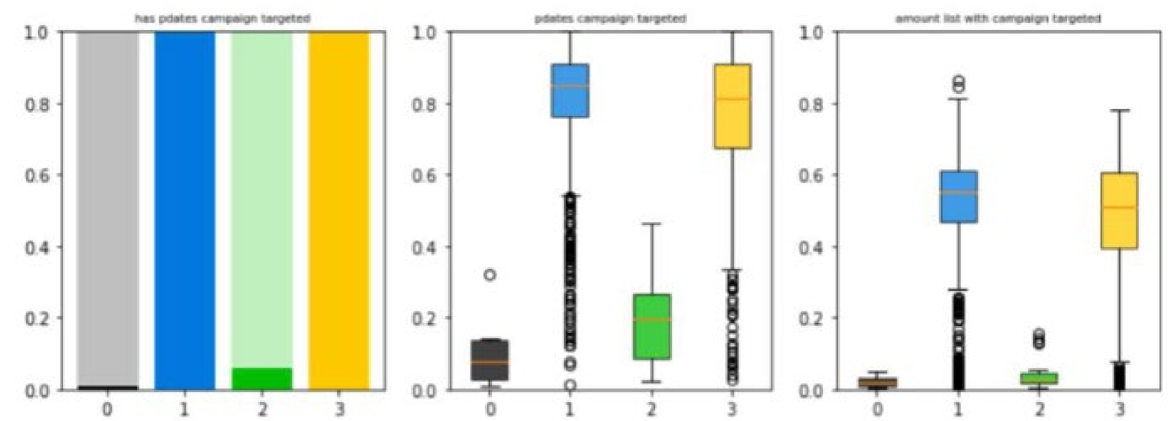

Profiling of clusters to identify differences in behavior across clusters.

Assessing results

With profiling, marketers can discern those customer households in the highlighted example fall into two groups: those who are responsive to coupons and mailed leaflets, and those who are not. Further divisions show differing degrees of responsiveness to other promotional offers.

Comparing households by demographic factors not used in developing the clusters themselves, some interesting patterns separating cluster members by age and other factors are identified. While this information may be useful in not only predicting cluster membership and designing more effective campaigns targeted to specific groups of households, the team recognizes the need to collect additional demographic data before putting too much emphasis on these results.

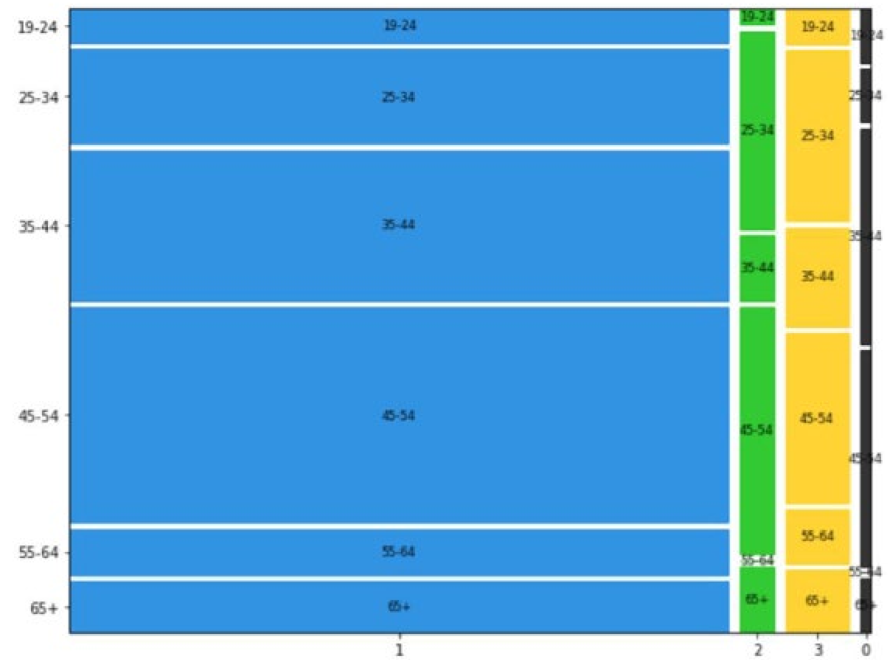

Age-based differences in cluster composition of behavior-based customer segments.

The results of the analysis now drive a dialog between the data scientists and the promotions management team. Based on initial findings, a revised analysis will be performed focused on what appear to be the most critical features differentiating households as a means to simplify the cluster design and evaluate overall cluster stability. Subsequent analyses will also examine the revenue generated by various households to understand how changes in promotional engagement may impact customer spending.

Using this information, the team believes they will have the ability to make a case for change to upper management. Should a change in promotions targeting be approved, the team makes plans to monitor household spending, promotions spend and campaign responsiveness rates using much of the same data used in this analysis. This will allow the team to assess the impact of these efforts and identify when the segmentation design needs to be revisited.

Need help segmenting your customers for more targeted marketing?

Use the **Customer Segmentation Accelerator** and drive better purchasing predictions based on behaviors. Through sales data, campaigns and promotions systems, you can build useful customer clusters to effectively target various households with different promos and offers.

GET THE ACCELERATOR >

Assessing Consumer Interest Data to Inform Engagement Strategies

Fine-tuning ML recommendations to boost conversions

Personalization is a journey. To operationalize personalized experiences, it's important to identify high-value audiences who have the highest likelihood of specific actions. Here's where **propensity scoring** comes in.

Specifically, this process allows companies to estimate customers' potential receptiveness to an offer or to content related to a subset of products, and determine which messaging to apply. Calculating propensity scores requires assessment of past interactions and data points (e.g., frequency of purchases, percentage of spend associated with a particular product category, days since last purchase and other historical data).

Databricks provides critical capabilities for propensity scoring (like the Feature Store, AutoML and MLflow) to help businesses answer three key considerations and develop a robust process:

1. How to maintain the significant number of features used to train propensity models
2. How to rapidly train models aligned with new campaigns
3. How to rapidly re-deploy models, retrained as customer patterns drift, into the scoring pipeline

Boosting model training efficiency

With the Databricks Feature Store, data scientists can easily reuse features created by others.

The feature store is a centralized repository that enables the persistence, discovery and sharing of features across various model training exercises. As features are captured, lineage and other metadata are captured. Standard security models ensure that only permitted users and processes may employ these features, enforcing the organization's data access policies on data science processes.

Extracting the complexities of ML

Databricks AutoML allows you to quickly generate models by leveraging industry best practices. As a glass box solution, AutoML first generates a collection of notebooks representing various aligned model variations. In addition to iteratively training models, AutoML allows you to access the notebooks associated with each model, creating an editable starting point for further exploration.

Streamlining the overall ML lifecycle

MLflow is an open source machine learning model repository, managed within the Databricks Lakehouse. This repository enables tracking and analysis of the various model iterations generated by both AutoML and custom training cycles alike.

When used in combination with the Databricks Feature Store, models persisted with MLflow can retain knowledge of the features used during training. As models are retrieved, this same information allows the model to retrieve relevant features from the Feature Store, greatly simplifying the scoring workflow and enabling rapid deployment.

How to build a propensity scoring workflow with Databricks

Using these features in combination, many organizations implement propensity scoring as part of a three-part workflow:

1. Data engineers work with data scientists to define features relevant to the propensity scoring exercise and persist these to the Feature Store. Daily or even real-time feature engineering processes are then defined to calculate up-to-date feature values as new data inputs arrive.

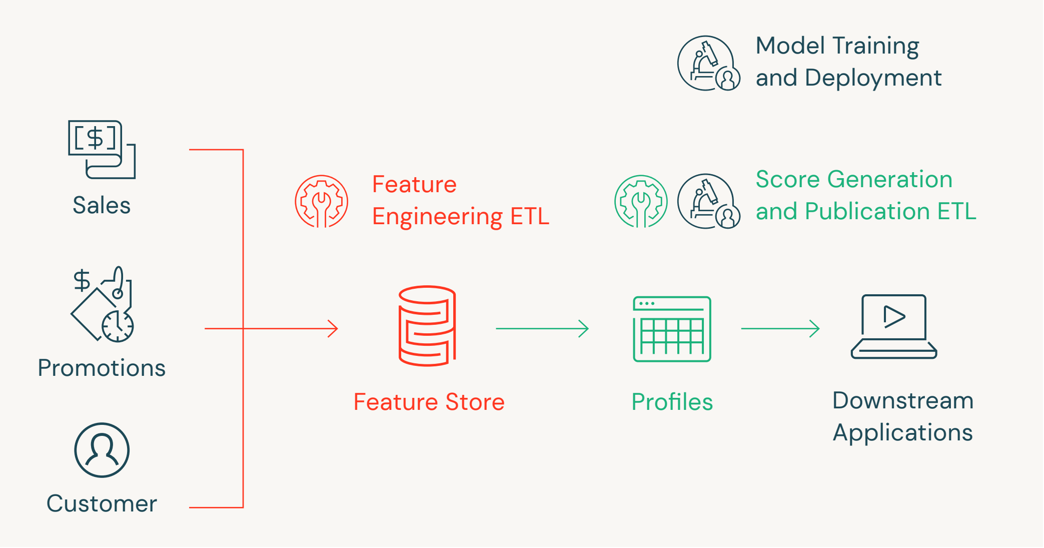

A three-part propensity scoring workflow.

2. As part of the inference workflow, customer identifiers are presented to previously trained models in order to generate propensity scores based on the latest features available. Feature Store information captured with the model allows data engineers to retrieve these features and easily generate the desired scores, which can then be used for analysis within Databricks Lakehouse or published to downstream marketing systems.

3. In the model-training workflow, data scientists periodically retrain the propensity score models to capture shifts in customer behaviors. As these models are persisted to MLfLow, change management processes are used to evaluate and elevate those models that meet organizational criteria-to-production status. In the next iteration of the inference workflow, the latest production version of each model is retrieved to generate customer scores.

Need help assessing interest from your target audience?

Use the **Propensity Scoring Accelerator** to estimate customers' potential receptiveness to an offer or to content related to a subset of products. Using these scores, marketers can determine which of the many messages at their disposal should be presented to a specific customer.

GET THE ACCELERATOR >

Delivering Personalized Customer Journeys

Strategies for crafting a real-time recommendation engine

As the economy continues to weather unpredictable disruptions, shortages and demand, delivering personalized customer experiences at speed and scale will require adaptability on the ground and within a company's operational tech stack.

CASE STUDY

Creating a unified view across 200+ brands

As a driving force for economic growth in the Middle East, Al-Futtaim impacts the lives of millions of people across the region through the distribution and operations of global brands like Toyota, IKEA, Ace Hardware and Marks & Spencer.

Al-Futtaim's focus is to harness their data to improve all areas of the business, from streamlining the supply chain to optimizing marketing strategies. But with the brands capturing such a wide variety of data, Al-Futtaim's legacy systems struggled to provide a single view into the customer due to data silos and the inability to scale efficiently to meet analytical needs.

With the Databricks Lakehouse, Al-Futtaim has transformed their data strategy and operations, allowing them to create a "golden customer record" that improves all decision-making from forecasting demand to powering their global loyalty program.

Get the full story →

> "Databricks Lakehouse allows every division in our organization — from automotive to retail — to gain a unified view of our customer across businesses. With these insights, we can optimize everything from forecasting and supply chain, to powering our loyalty program through personalized marketing campaigns, cross-sell strategies and offers."

DMITRIY DOVGAN
Head of Data Science, Al-Futtaim Group

As COVID-19 forced a shift in consumer focus toward value, availability, quality, safety and community, brands most attuned to changing needs and sentiments saw customers switch from rivals to their brand. While some segments gained business and many lost, organizations that had already begun the journey toward improved customer experience saw better outcomes, closely mirroring patterns observed in the 2007–2008 recession.

[1] Comparison of total returns to shareholders for publicly traded companies ranking in the top 10 or bottom 10 of Forrester's Customer Experience Performance Index in 2007-09.

Source: Forrester Customer Experience Performance Index (2007-09); press search

CX leaders outperform laggards, even in a down market, in this visualization of the Forrester Customer Experience Performance Index as provided by McKinsey & Company.

The personalization of customer experiences will remain a key focus for B2C and B2B organizations. Increasingly, market analysts are recognizing customer experience as a disruptive force enabling upstart organizations to upend long-established players.

Focus on the customer journey

Personalization starts with a careful exploration of the customer journey. The digitization of each stage provides the customer with flexibility in terms of how they will engage and provides the organization with the ability to assess the health of their model.

flaconi

CASE STUDY

Personalizing the beauty product shopping experience

Flaconi wanted to leverage data and AI to become the No. 1 online beauty product destination in Europe. However, they struggled with massive volumes of streaming data and with infrastructure complexity that was resource-intensive and costly to scale. See how they used Databricks to increase time-to-market by 200x, reduce staff costs by 40% and increase net order income.

Get the full story →

Careful consideration of how customers interact with various assets — and how these interactions may be interpreted as expressions of preference — can unlock a wide range of data that enables personalization.

CASE STUDY

Connecting shoppers to savings with data-driven personalization

Flipp is an online marketplace that aggregates weekly shopping circulars, so consumers get deals and discounts without clipping coupons. Siloed customer data sources once made getting insights difficult. Now with Databricks, Flipp's data teams can access and democratize data, helping them do their jobs more effectively while bringing better deals to users, more meaningful insights to partners, and a 10% jump in foot traffic to brick-and-mortar retailers.

Get the full story →

The engines we use to serve content based on customer preferences are known as recommenders. With some recommenders, a heavy focus on the shared preferences of similar customers helps define what recommendations will actually make an impact. With others, it can be more useful to focus on the properties of the content itself (e.g., product descriptions).

The complexity of these engines requires that they be deployed thoughtfully, using limited pilots and customer response assessments. And in those assessments, it's important to keep in mind that there is no expectation of perfection — only incremental improvement over the prior solution.

Need help generating personalized recommendations?

Use the **Recommendation Engines Accelerator** to estimate customers' potential receptiveness to an offer or to content related to a subset of products. Using these scores, marketers can determine which of the many messages at their disposal should be presented to a specific customer.

GET THE ACCELERATOR >

Building a Direct Path to Winning the Minds and Wallets of Your Customers

Providing deep, effective personalized experiences to customers depends on a brand's ability to intelligently leverage consumer and market data from a wide variety of sources to fuel faster, smarter decisions — without sacrificing accuracy for speed. The Databricks Lakehouse Platform is purpose-built for exactly that, offering a scalable data architecture that unifies all your data, analytics and AI to deliver unforgettable customer experiences.

Created on open source and open standards, Databricks offers a robust and cost-effective platform for brands to collaborate with partners, clients, manufacturers and distributors to unleash more innovation and efficiencies at every touch point. Businesses can rapidly ingest available data in real time, at scale, and create accessible, data-driven insights that enable actionable strategies across the value chain.

Databricks is a multicloud platform, designed for quick enterprise development. Teams using the Lakehouse can more effectively reveal the 360-degree view into their company's operational health and the evolving needs of their customers — all while empowering teams to easily unify data efforts, perform fine-grained analyses and streamline cross-functional data operations using a single, sophisticated solution.

Learn more about Databricks Lakehouse for industries like Retail & Consumer Goods, Media & Entertainment and more at **databricks.com/solutions**

About Databricks

Databricks is the data and AI company. More than 7,000 organizations worldwide — including Comcast, Condé Nast, H&M and over 50% of the Fortune 500 — rely on the Databricks Lakehouse Platform to unify their data, analytics and AI. Databricks is headquartered in San Francisco, with offices around the globe. Founded by the original creators of Apache Spark™, Delta Lake and MLflow, Databricks is on a mission to help data teams solve the world's toughest problems. To learn more, follow Databricks on Twitter, LinkedIn and Facebook.

START YOUR FREE TRIAL

Contact us for a personalized demo
databricks.com/contact

databricks