Acquia
EXPERIENCE DIGITAL FREEDOM

MACHINE LEARNING MODELS IN ACTION

Making AI Easy for Marketers

TABLE OF CONTENTS

SECTION 01

RISE OF THE MACHINES

Need to process data at a scale and speed impossible for human beings? That's what machine learning is for.

The amount of customer data available to marketers has increased exponentially, especially as so much commerce now takes place on multiple channels. At the same time, the tools available to analyze and operationalize this data have become incredibly sophisticated. As a result, marketers today can make extremely accurate predictions about customer behavior, both individually and in segmented groups.

One of the most powerful tools in a marketer's arsenal – one that gives brands significant competitive advantage in customer acquisition, engagement and retention – is machine learning.

Machine learning, which some equate with artificial intelligence and others see as a subcategory of it, refers to the ability of computers to "learn" from the analysis of massive datasets. While machine learning has applications in multiple fields, in marketing it is most frequently used to uncover patterns in customer data. These patterns can be used to explain observed customer behaviors as well as predict future behavior.

The extreme power of machine learning models lies in the fact that they can process data at a scale and speed impossible for human beings. Machine learning tools can reveal commonalities among thousands or even millions of customers across all interaction points. These insights can drive increasingly relevant recommendations, improve the likelihood that specific marketing strategies will succeed and increase everything from marketing ROI to revenue growth.

EMPOWERING MARKETERS WITH MACHINE LEARNING MODELS

Over time, data scientists have evolved a wide range of machine learning models optimized for specific use cases. To harness their power, marketers need to familiarize themselves with these models and, ideally, have a way to quickly and easily put them to use. In our data-driven economy, brands that empower their marketers with configurable, user-friendly machine learning tools will be the winners.

The Acquia Customer Data Platform (CDP) makes predictive analytics and machine learning technology available to the average marketer through intuitive machine learning models that don't require technical expertise to utilize. This e-book explores Acquia's comprehensive set of machine learning models, the specific applications for which these models were developed and the tangible benefits they provide.

These machine learning models can influence and improve marketing performance in three fundamental ways. They can predict customer behavior. They can provide a more refined and granular understanding of customer personas. And they can enable one-to-one personalization.

Most importantly, they can do all of the above at an unparalleled scale and with a level of customization that enables adaptation to any marketing scenario.

SECTION 02

PREDICTIONS

Predicting Behavior to Improve Marketing Performance

Being able to predict how customers or potential customers will react to offers, promotions and other marketing tactics is incredibly powerful. Acquia's predictive machine learning models cover behaviors across the entire customer lifecycle.

1. LIKELIHOOD TO BUY

Ideally, marketing campaigns would only target those customers most likely to make a purchase. The Likelihood to Buy model aims to predict near-future repeat purchasing behavior based on past transactions as well as email and browsing behavior.

This model allows you to aim your marketing activity at those with a high purchase probability. This helps you to improve conversion rates and tailor offers to the recipient. From an analytics standpoint, this model can also help you understand what variables in your marketing mix actually drive purchase behavior.

Arcelik is a multinational appliance and electronics manufacturer with multiple brands that operates in over 30 countries. Arcelik used Acquia CDP to turn customer support calls into a marketing opportunity.

During the assembly/repair process, customers were asked "Would you consider buying an electronic product within the next 6 months?" An affirmative answer automatically triggered a daily automated marketing campaign called QuickWin. By using Acquia CDP's Likelihood to Buy machine learning model, **Arcelik saw a 6x increase in conversion rates from this campaign.**

2. LIKELIHOOD TO PAY FULL PRICE

Many brands fall into the bad habit of over-discounting or offering blanket discounts to all customers. The Likelihood to Pay Full Price model predicts the degree to which a customer is likely to purchase a product without a discount.

This model can be used for pricing optimization, allowing you to avoid cannibalizing margins by offering discounts to those who would buy without them. It can also be used to optimize revenue by recommending full-price items to those willing to pay and discounted items to those who are less likely to. You can also create targeted segments of customers likely to pay full price.

Shoe manufacturer and retailer Clarks leveraged Acquia CDP to gather customer data from their retail, e-commerce, email marketing and analytics systems and identify which specific customers were incentivized by discounts and which were not. Rather than blasting all customers with discount offers, the company tailored emails and campaigns to customers likely to pay full price. Using this approach, a set of campaigns projected to earn $500,000 ended up generating **$1.4 million in revenue**, helping the marketing team exceed its year-end goal by **more than $200,000** and putting them on track to double revenue the following year.

3. PREDICTIVE LIFETIME VALUE

Knowing how much a given customer is likely to spend over time is crucial for identifying your most valuable customers and building segments around them. Once you've created an accurate profile of your most valuable customers, you can analyze new customer data and identify those who most resemble customers with a high predicted lifetime value.

The Predictive Lifetime Value model predicts the expected revenue or margin associated with a particular customer in the next 12 months. This model allows you to identify your VIP customers and create targeted campaigns – special offers, special treatment, gifts – aimed at increasing their loyalty and optimizing their spend. It also allows you to segment direct mail campaigns, limiting them to these customers specifically to ensure higher ROI.

Francesca's is a women's boutique clothing chain with hundreds of locations and a significant e-commerce operation. As the company engaged customers across multiple channels, it gathered a plethora of data: an email address in one channel, a phone number collected in another and payment information on the e-commerce site. With data in various silos, one customer could have multiple records featuring different emails and phone numbers scattered across their systems.

Francesca's needed to aggregate this data to create a single, unified view of the customer at a granular level. Acquia CDP made it possible to integrate these disparate data points and break down their data silos. As a result, Francesca's could effectively **calculate the customer lifetime value for specific customers.**

This exercise provided a number of "aha" moments – such as the extent to which the lifetime value of a customer whose first purchase was a dress differed from that of a customer whose first purchase focused on accessories – and drove decisions about the categories the brand chose to leverage during customer acquisition campaigns.

4. LIKELIHOOD TO CHURN

Companies inevitably spend more money acquiring new customers than they do coaxing existing customers to buy more. For this reason, identifying customers at risk of churn and reengaging them should be a top priority.

The Likelihood to Churn model identifies at-risk customers, particularly those who are becoming less likely to buy over time. This model can be used to create win-back campaigns focused specifically on customers you are most likely to lose. This model can be combined with the Predictive Lifetime Value model to further segment churners.

Headwear retailer Lids leveraged Acquia CDP to run a win-back campaign focused on re-engaging at-risk customers. Lids's marketers identified 1 million target customers and offered them a time-sensitive, in-store loyalty reward. **The campaign ended up driving $250,000 in revenue.**

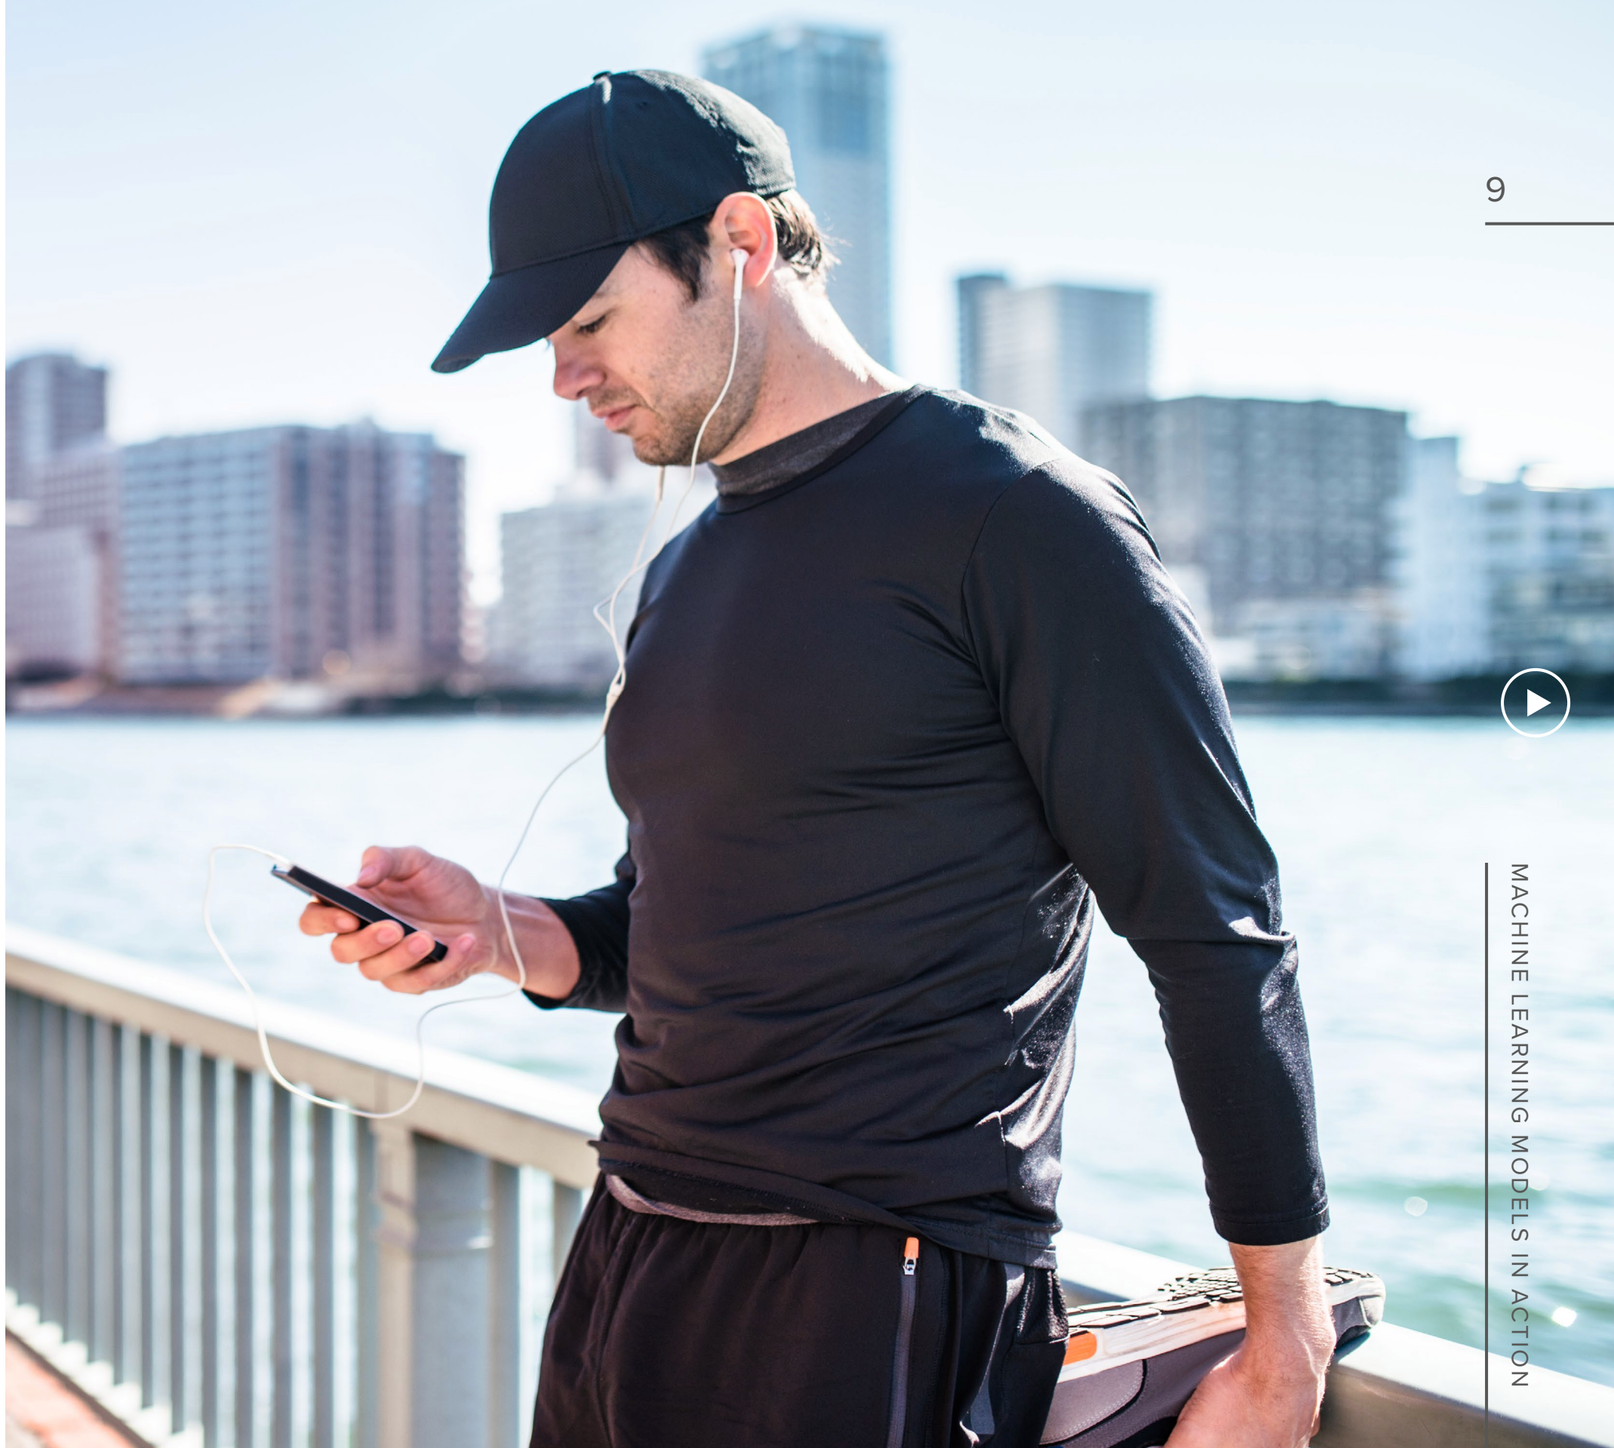

5. LIKELIHOOD TO ENGAGE

Before a person becomes a customer, they need to engage with your brand. Ongoing engagement can also distinguish one-time customers from lifetime customers. Understanding how likely someone is to engage with your marketing is key to ensuring that you are not only targeting the right people, but also targeting them in the right channel with the right message.

The Likelihood to Engage model predicts the odds a customer will open an email, click on a link, subscribe to a newsletter, join a mailing list or otherwise opt to engage with the brand. Marketers can use this model to improve the performance of marketing campaigns.

A home improvement and garden retailer wanted to increase engagement and conversions through more relevant email campaigns. Using Acquia CDP, they determined customers who had abandoned their carts had a higher likelihood to engage with emails if the messaging and creative was personalized according to the product or category they were interested in. To be systematic about it, they only tested the relevance of these targeted emails to the recipients – no offers or discounts were included.

This approach increased open rates 108% and saw a 275% increase in conversions.

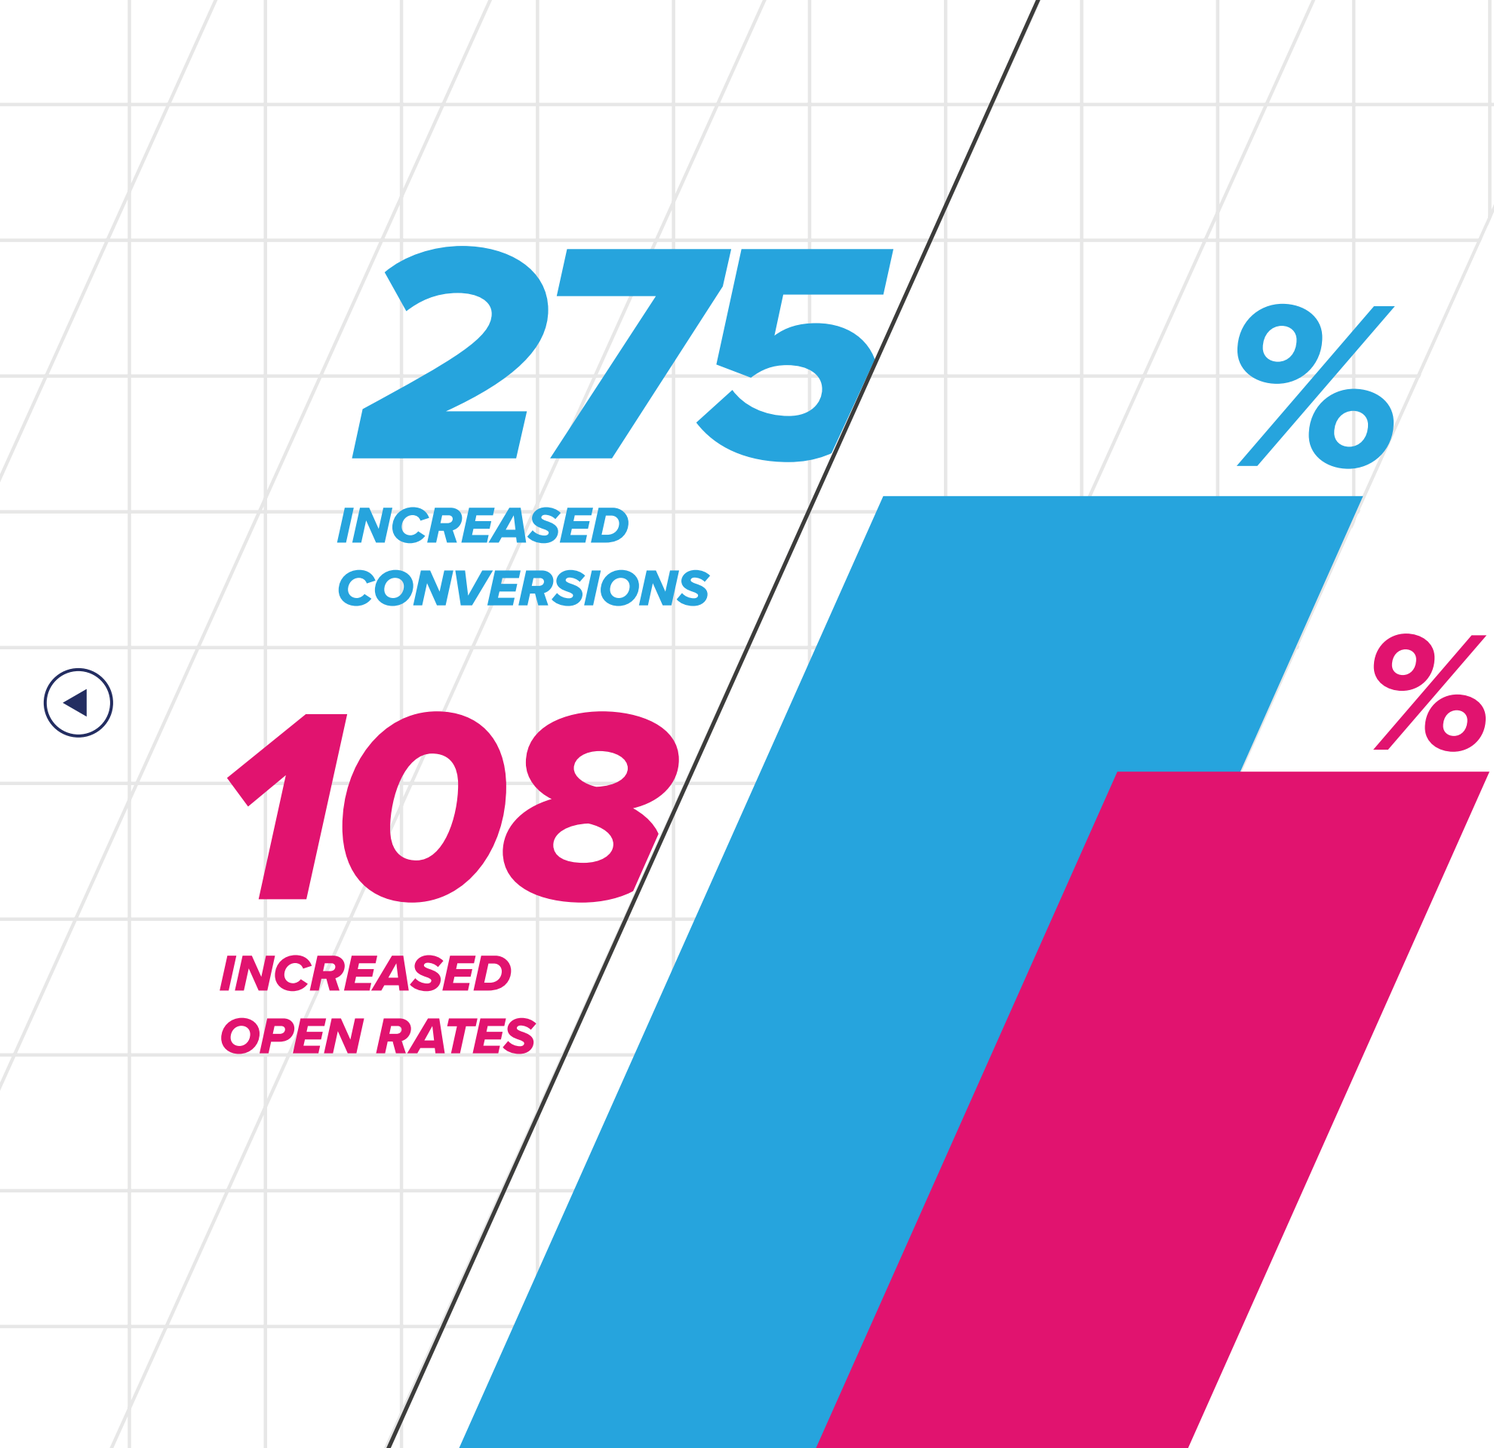

SECTION 03

PERSONAS

Advanced Segmentation through Customer Clustering

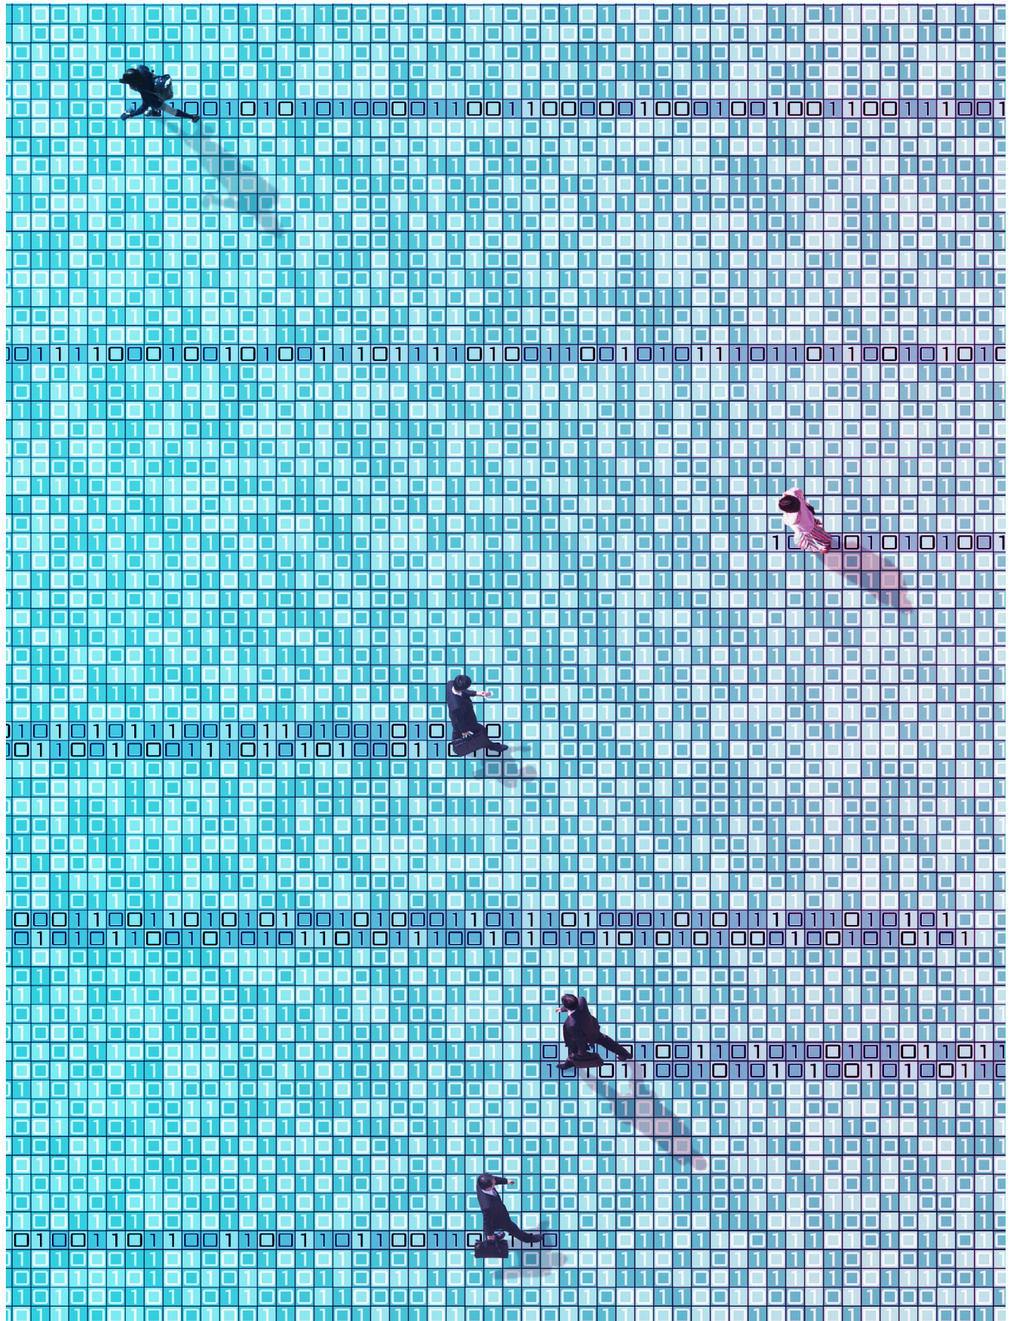

All brands segment their customers. Unfortunately, this segmentation is rarely as useful as it could be because the segments created are simply too broad (“homeowners”) or too generic (“people who have made a purchase in the last six months”). By uncovering meaningful patterns in customer data, machine learning can use a wide variety of behavioral and demographic data to create much more granular customer segments or clusters.

Acquia’s machine learning models focused on customer segmentation and clustering bring more accuracy and nuance to your marketing by helping you group your customers and prospects based on a range of dimensions. To ensure a comprehensive view of your customer clusters, Acquia CDP takes advantage of fuzzy clustering.

FUZZY CLUSTERING

Machine learning and the predictions it enables are based on probabilities. That is, when we say that a given customer fits within a particular cluster, we are really saying that this customer will probably exhibit the traits we associate with that cluster.

Since we’re dealing with probabilities, it’s possible that a customer may fall into more than one cluster — one with a high probability, but several others with a low probability. At the same time, it means that each cluster itself can be further segmented into high probability and low probability groups.

Fuzzy clustering is not a model in and of itself, but a way to make other machine learning models better. It allows marketers to explore the many different categories a customer belongs to. By identifying the top three clusters to which a customer may belong while also specifying the probability that they belong in a particular cluster, fuzzy clustering provides a much more comprehensive view of the customer and we use it in all our clustering models.

6. PRODUCT-BASED CLUSTERING

One simple way to segment customers is to group them based on the products they have purchased. The Product-Based Clustering model groups customers based on products they prefer or types of products they tend to buy together.

Product-based clustering (sometimes called “category-based clustering”) can help you separate customers who only buy one type of product, sweaters, for example, from those who buy sweaters and other related products, such as sportswear, outerwear, swimwear and so on.

Clothing retailer J.Crew leveraged Acquia CDP to send personalized emails to customers who had browsed or purchased cashmere within the previous 365 days and, at the same, sent emails to its “business-as-usual” (BAU) audience. The cashmere audience saw a double-digit lift over the BAU audience in average order value (AOV), conversion rates, open rates and click rates. **While the cashmere audience was only 10% of the full circulation, it drove almost 50% of the total demand.**

A similar model, brand-based clustering, tells you what specific brands people like and tend to buy. When a brand releases new products you can target those who will be most interested. Brand-based clustering can also determine what other brands people are most likely to be interested in.

For example, customers who buy Converse sneakers may also like Adidas, but might not be interested in Bally shoes. This is useful information when deciding which product offers or email content to send. It also provides valuable insight into your overall product mix, revealing which products and product groups drive the most valuable customer behaviors.

7. SEASONAL CLUSTERING

Just as we find customer cohorts coalescing around products, we also find cohorts that exhibit seasonal buying behavior. The Seasonal Clustering model helps you identify these cohorts. They can include parents who purchase athletic gear for children at the beginning of different sports seasons or customers who only visit a store during the holidays.

Clustering customers based on the seasonality of their purchasing behavior allows you to target them at appropriate times as well as suppress them from marketing campaigns when they are less likely to purchase. Of course, it also allows you to target them during off times with special offers. Seasonal clustering, by providing insight into the ebb and flow of buying activity, helps with marketing and resource planning throughout the year.

Believing that previous gift buyers could be enticed to purchase more gifts, premium chocolate vendor Godiva created media and email campaigns retargeting buyers who had shopped during at least one of the previous gifting seasons — Mother's Day, the holiday season or Valentine's Day. The following Mother's Day, the brand saw its campaigns yield a **return on advertising spend (ROAS) increase of 621% and 42% more clicks.**

8. BEHAVIORAL CLUSTERING

While product-based and seasonal-based clusters do reflect certain behaviors, there is a wide range of other behaviors that can be used to create meaningful segments for marketers. The Behavioral Clustering model groups customers into homogeneous clusters based on buying behavior, including recency of purchase, frequency of purchase, average order value and predicted (or actual) lifetime value.

This clustering allows marketers to build VIP programs around customers with a high average order value, for example, or reach out to “discount junkies” with aggressive offers. Marketers can also better structure programs by defining personas that cover a broad swathe of customers. The number of dimensions marketers must consider when defining messaging can be significantly reduced, allowing them to focus on five or six personas, rather than millions of individuals.

Headquartered in Turkey, Aydinli is a retailer of children and adult clothing. It is the licensed brand distributor in Asia, the Middle East and south Europe for brands including Pierre Cardin, Cacharel and U.S. Polo Assn. Given the competitiveness of their market, the company wanted a system that would allow them to quickly and accurately create audiences for targeted campaigns.

To improve campaign effectiveness, they used Acquia CDP to create behavior-based clusters, allowing them to target or eliminate segments, such as high-returners, digital-only buyers and omnichannel customers. This drove better conversion rates and helped them avoid revenue cannibalization.

The result? An ROI of more than 3,500%.

SECTION 04

PERSONALIZATION

Making Messages More Relevant

Optimized recommendations are a powerful and effective form of personalization.

Survey after survey emphasizes how much customers desire personalization. While it may be nice to personalize communications using a customer's first name, it has a much greater impact when personalization takes the form of respecting their preferences and making relevant recommendations.

The machine learning models we have developed for personalization focus primarily on recommendations, both in terms of product recommendations for the customer as well as tactical recommendations for the marketer.

9. NEXT-BEST PRODUCT

There are a few different ways to approach next-best product recommendations. The key is to use your data to ensure that the recommendations you make are relevant. This means not looking at recommended purchases in terms of the customer's past transactions but rather in terms of similar purchases by customers who occupy related clusters or segments.

The Next-Best Product model covers three use cases:

- **Upselling at the time of purchase.** In this scenario, you recommend that the customer buy a better version or brand of their intended purchase. "Supersizing" your McDonald's order is one example of this but it could also take the form of recommending a watch of the same brand but with more features, for example. This model helps you identify the upsell offer with the highest chance of success.
- **Cross-selling at time of purchase or in a follow-up email.** These recommendations focus on items that are often purchased along with the primary product. You can also offer a modest discount if the customer buys the entire cross-sell bundle. This model uses aggregated buying behavior to identify the bundles most commonly sold together.
- **Future purchase.** This model uses purchasing information from across the business to identify typical patterns of future purchases. For example, a home improvement store found that people who build decks tend to be in the market for a grill shortly thereafter. They devised a recommendation program to capitalize on this insight.

10. NEXT-BEST CHANNEL

Because customers interact with brands on so many different channels, marketers need to be sure they're engaging with customers on the channels they prefer. Of course, they also need to be sure to select the channels where they are most likely to get a response.

The Next-Best Channel model considers the customer and similar customers in their cluster(s) and recommends the next-best channel on which to engage them. This model balances out the customer's stated preferences against the customer's actual observed behavior.

An athletic apparel retailer with hundreds of stores globally had over a million customers shopping both on- and offline. They wanted to leverage the massive amount of customer data they had collected to orchestrate customer experiences across physical and digital channels. To do that, they needed to unify all of this customer intelligence.

With Acquia CDP, the company could finally connect the dots between myriad digital channels and in-person customer touchpoints, allowing them to create and execute tailored, in-person and online campaigns.

For example, they would customize online and offline marketing campaigns and offers based on a customer's preferences and behaviors (whether a customer preferred running to indoor cycling or tended to buy more online than in store, for example).

By creating a seamless cross-channel experience, the company **increased site visits by nearly 50%, increased attendance at local events by 25% and achieved a 10-15% increase in baseline revenue from digital marketing campaigns.**

11. SEND TIME OPTIMIZATION

Email response rates can be significantly improved if customers receive emails when they are most likely to read and respond. The Send Time Optimization model determines the best time of day to send emails to particular customers. This data can also be aggregated to schedule marketing emails to entire clusters of customers who exhibit similar response behaviors.

A home decor brand consistently maximized the impact of sale events with an email push on the sale's final day. They would first schedule an email blast in the morning reminding customers that the sale was almost over. Towards the end of the day, in the last hour of the sale, they would send another batch of emails, but only to customers who had previously browsed sale products without buying.

Because it didn't inundate customers with irrelevant offers, this targeted approach helped keep unsubscribe rates low and improve overall clicks and conversions.

SECTION 05

LEVERAGING ML MODELS WITH ACQUIA CDP

Since its creation, Acquia CDP has offered organizations predictive analytics capabilities. Acquia CDP gives marketers the models they need to use machine learning and seamlessly optimize marketing efforts without the need for deep technical expertise.

LEVERAGE ANY DATA SOURCE AT ANY SCALE

Acquia CDP's Machine Learning framework performs **1.4 billion predictions a day** and can process any data ingested by the CDP. Many vendors offer ML models that only work with a predefined set of data sources and attributes. There are also limits with other solutions when it comes to the amount of data that can be processed with their ML models.

Acquia CDP's ML models, however, don't have these limitations and can work with data sets of any size.

CONFIGURE MODELS TO FIT YOUR NEEDS

Acquia CDP embraces configurability in machine learning so that brands can customize models to suit their specific needs. The static machine learning models provided by other vendors don't allow for customized use cases.

Many vendors will claim their models are configurable, but often, "configurable" turns out to be a development project in Python, R or Java.

Acquia's framework is built with a configuration-first approach, making it possible for marketers to configure ML models without the need for extensive support from developers.

TRUST AND UNDERSTAND THE MODELS THAT YOU USE

From our experience, clients will only adopt what they trust, understand and can easily test. We build our products with this in mind.

Marketers need to be able to understand and trust the predictions generated by machine learning rather than blindly rely on predetermined scoring systems assigned by black-box models. By giving marketing teams direct access to all of their customer data via a CDP and the ability to design what tests they want to run through configurable machine learning, teams will have a deeper understanding of what each prediction means and be able to interpret results to inform strategic actions.

When marketers have the power of trusted, configurable machine learning models at their fingertips, there is really no limit on how much they can improve the customer experience and achieve predictable, ongoing marketing success.

CDPs offer a very efficient solution to leverage the massive amount of unstructured data brands have accumulated. This data has been increasingly difficult to manage for brands: it comes from a lot of different sources — sometimes in real-time, sometimes in batch and it rarely shares the same structure.

Modern martech systems need to be built with AI in order to create contextual meaning from this massive amount of data. Customer data platforms with built-in machine learning capabilities like the Acquia CDP are perfectly suited to address these complex challenges."

ARTHUR GALIBERT, LEAD MARKETING CONSULTANT AT SQLI SWITZERLAND

MOVE FROM INSIGHT TO ACTION TODAY

Increase the power of your data with machine learning

LEARN MORE ▸

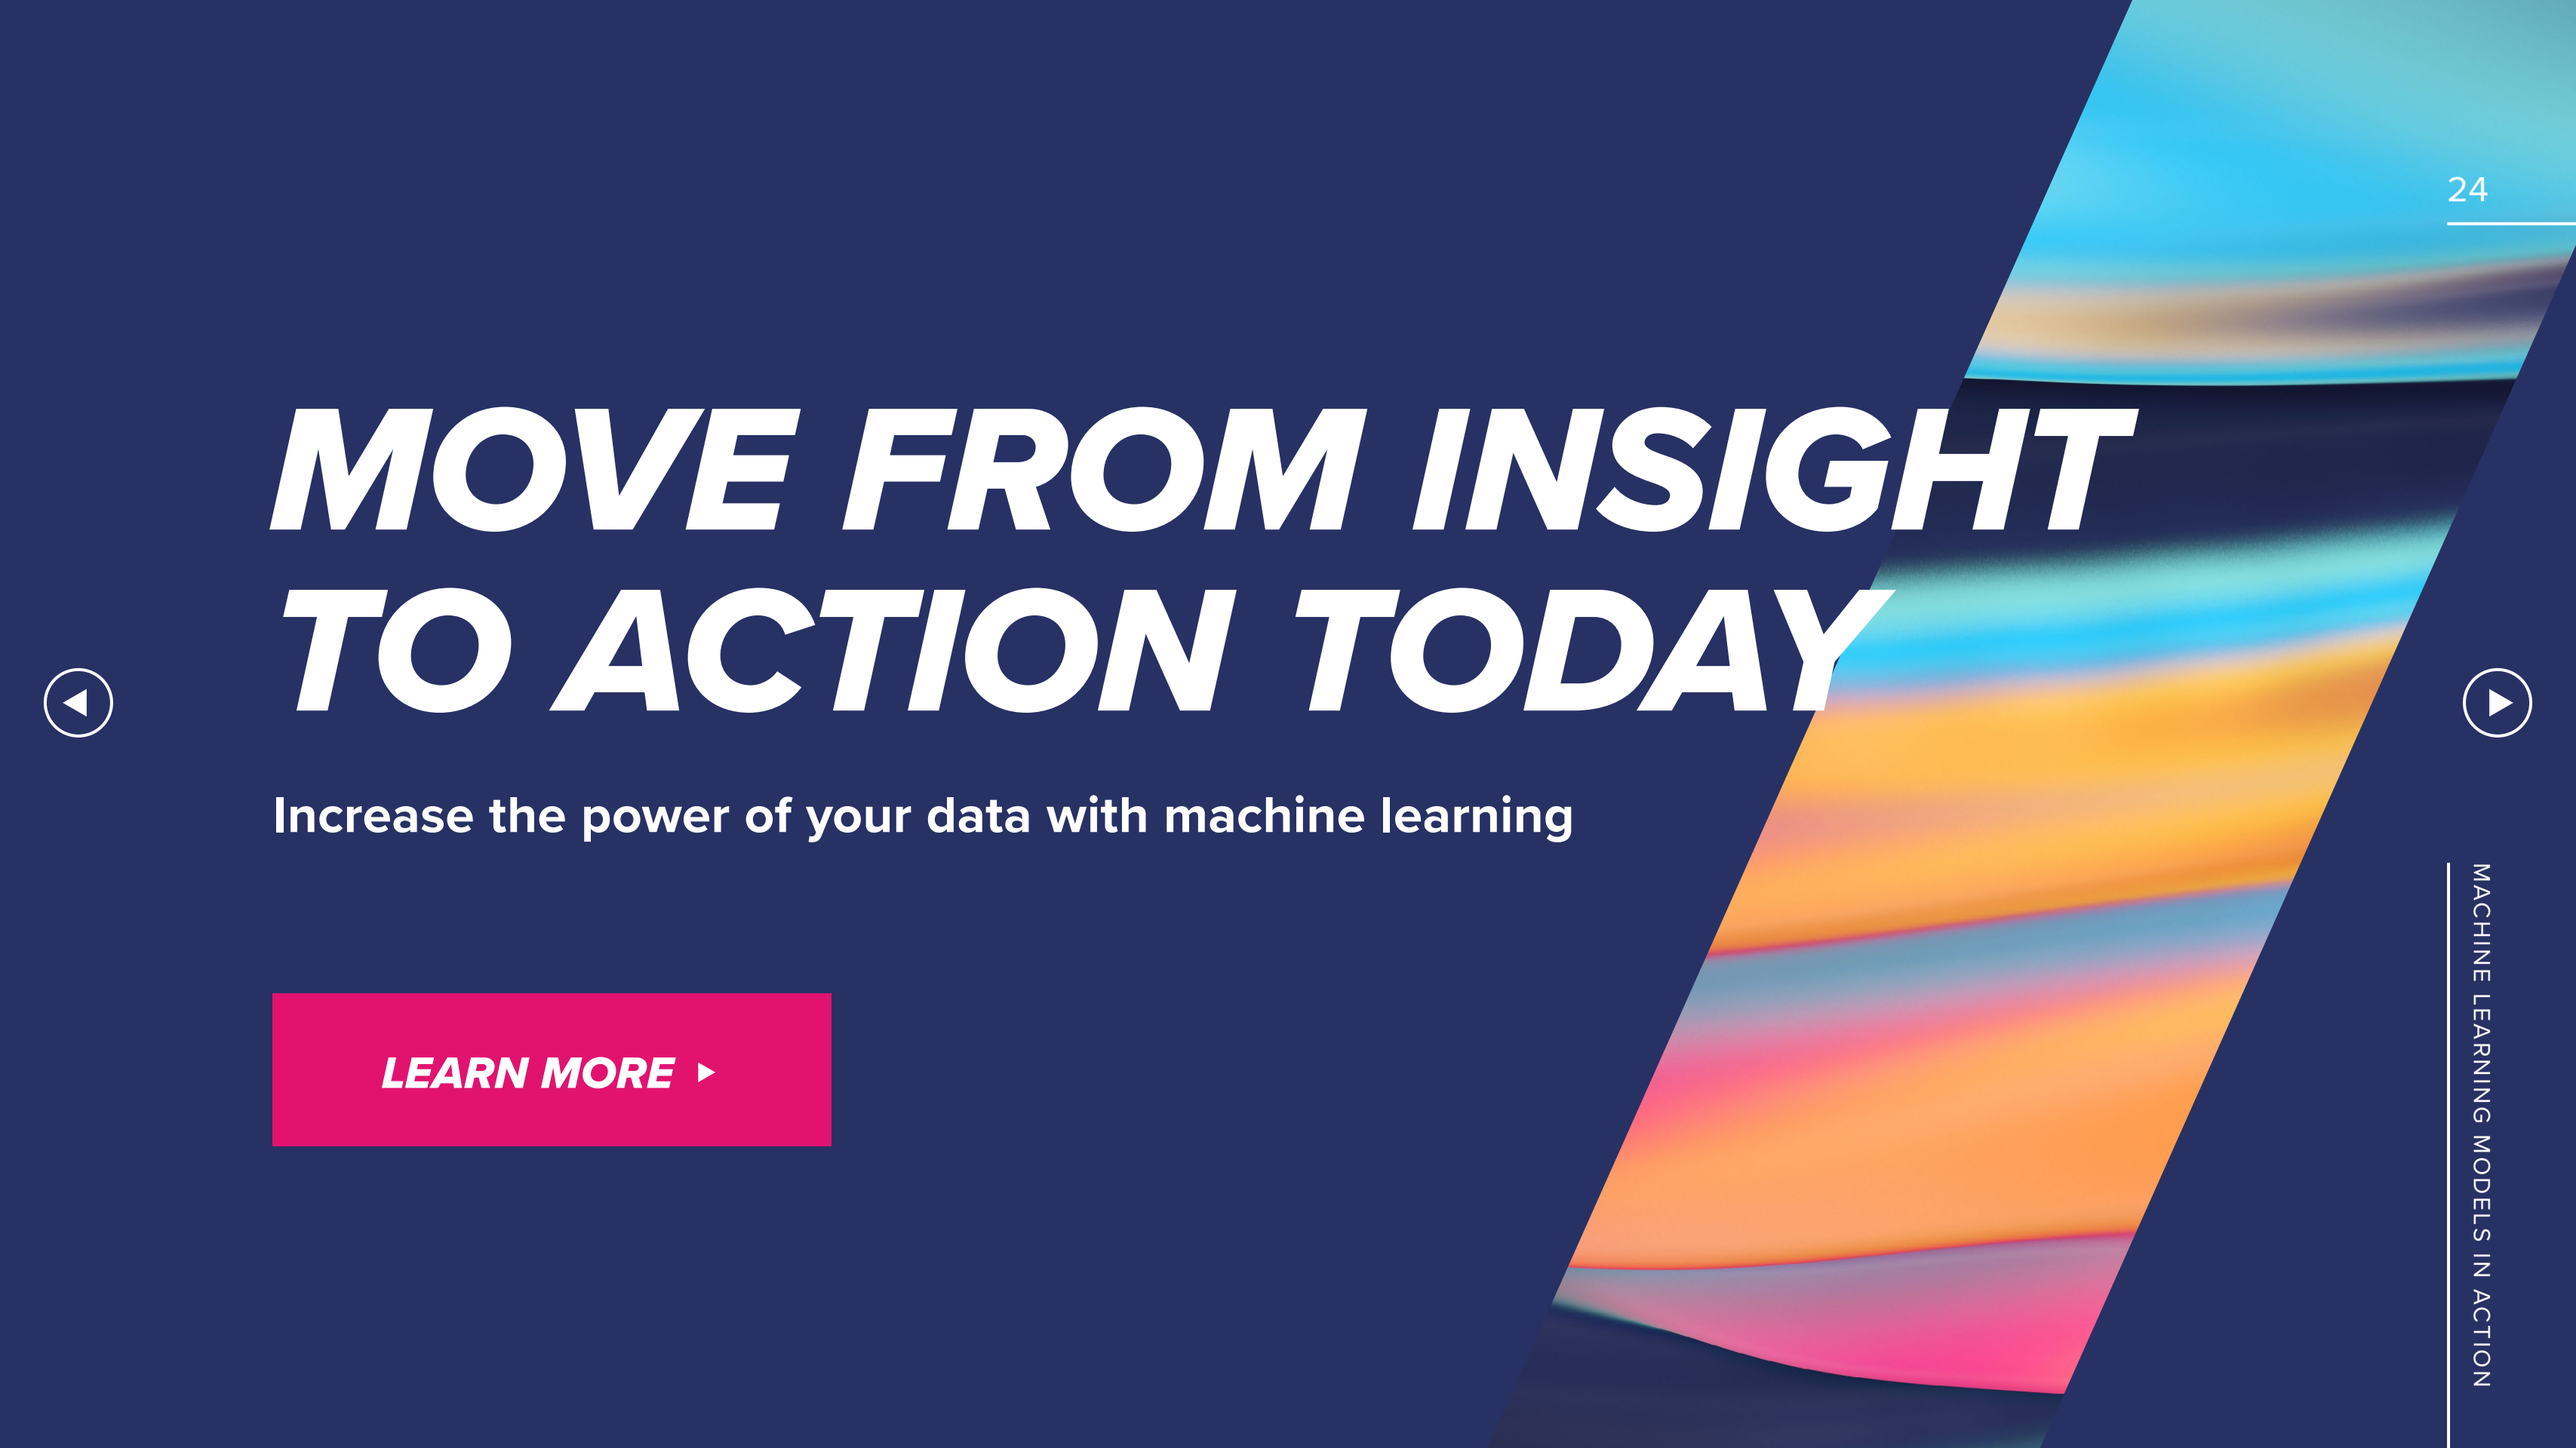

Acquia

ACQUIA.COM

ABOUT ACQUIA

Acquia is the open digital experience platform that enables organizations to build, host, analyze and communicate with their customers at scale through websites and digital applications. As the trusted open source leader, we use adaptive intelligence to produce better business outcomes for CX leaders.