

The background is a dark blue gradient with a large, stylized 'A' shape formed by two overlapping triangles. The left triangle is dark blue, and the right triangle is a lighter blue with a yellow-to-orange gradient at its base. The right side of the image features a bokeh effect of out-of-focus light points in various shades of blue.

Acquia
EXPERIENCE DIGITAL FREEDOM

ARTIFICIAL INTELLIGENCE AND CUSTOMER ANALYTICS

A Guide for Marketers

CONTENTS

SECTION 1:

**What Can Machine Learning Do
for Marketers?**

SECTION 2:

**How Can Marketers Apply
Machine Learning?**

SECTION 3:

**Harnessing the Power of
Machine Learning**

SECTION 4

**Choosing the Right Customer
Data Solution**

Introduction

Over the past 20 years, marketing has become increasingly data driven. One reason for this is the simple fact that the sheer quantity of customer data marketers can access has rapidly increased, especially as more and more commerce, both in the consumer and B2B world, takes place across digital channels.

At the same time, the tools available to analyze data – tools capable of transforming data into an immensely valuable, even strategic asset – have become increasingly sophisticated. **This is particularly true of those tools that leverage artificial intelligence (AI) in the form of machine learning.** These tools can make extremely accurate predictions about customer behavior, both at the level of customer groups as well as at the level of individual customers (through one-to-one personalization, for example).

Machine learning can make valuable predictions and recognize significant patterns because it works on massive datasets. Indeed, the scale of data machine learning can process, not to mention the number of variables it can consider, goes well beyond what a human can do. A small shopkeeper might remember the preferences of her longtime customers, but only machine learning tools can remember those preferences and make recommendations or find commonalities between thousands — or even millions — of customers.

In this e-book, **we will provide an overview of three high-level problems machine learning can address**, examples of specific applications of machine learning to improve marketing performance, and some recommendations on how marketing departments must change in order to take advantage of machine learning, particularly when it comes to data management and the marketing tools that support it.

Is There a Difference between AI and Machine Learning?

Artificial intelligence, in the broadest sense, refers to a human-like ability to perceive, understand and act on information when exhibited by a machine.

Machine learning (ML), generally considered a subset of artificial intelligence, is the ability for a machine to process data and derive new information from it, usually at a scale or speed that would be difficult (or even impossible) for humans to match.

For the purposes of this e-book, we will use the terms interchangeably. **Specifically, we will focus on machine learning's capacity for pattern recognition and prediction as the component of artificial intelligence that interests us most.**

SECTION 1

**What Can Machine
Learning Do
for Marketers?**



What Can Machine Learning Do for Marketers?

Machine learning can influence and improve marketing performance in three fundamental ways. It can predict customer behavior. It can provide a more refined and granular understanding of customer personas. And it can power one-to-one personalization. **Most importantly, it can do all of the above at an unparalleled scale.**

Predictions

Likelihood to engage

Likelihood to buy

Likelihood to convert

Personas

Product cluster

Behavioral cluster

Seasonal cluster

1-to-1 Personalization

Next Best Product

Next Best Action

Send-Time Optimization

Predictions

Being able to predict how customers or potential customers will react to offers, promotions and other marketing tactics is incredibly powerful. Machine learning provides you with exactly this capability.

What are the types of outcomes that machine learning can predict? Well, for starters, it can predict the likelihood that a customer will buy a particular product or set of products. Beyond this, and based on the data available, machine learning can predict behaviors such as:

- Likelihood to engage
- Likelihood to convert
- Likelihood to respond to a discount offer
- Likelihood to make a repeat purchase
- Likelihood to make a return
- Likelihood to churn

Understanding how customers will respond to different marketing tactics increases the effectiveness of your marketing efforts. For example, if you can predict that a customer will buy whether they receive a discount or not, then you can become more targeted in your discounting strategy. And by preventing over-discounting, you can automatically improve margins. This is just a small example of the incredible power of the predictions machine learning can generate.

Personas

Segmentation is hardly a new concept when it comes to marketing. However, marketers tend to create customer personas either based on assumptions or, when they have the data, based on what customers have purchased. Unfortunately, neither approach accurately reflects the makeup of your customer base.

Machine learning goes beyond historical purchase data to uncover commonalities between customers and customer attributes that allow you to group them into multiple clusters. Examples of common clusters include:

Product-Based Clusters group customers based on products they prefer or types of products they tend to buy together. For example, people who buy skis may also buy boots and other accessories.

Behavior-Based Clusters focus on how customers behave while making a purchase. For example, one can cluster customers based on responsiveness to discounts, frequency of purchases, average purchase amount and so on.

Seasonal Clusters consider customers based on whether their purchasing behavior is seasonal. For example, parents who make purchases at the beginning of soccer season, or people who only visit your store when the holidays roll around.

Machine learning can cluster customers based on a wide range of variables – from average order size, order variety and discount sensitivity, to total number of orders placed, order seasonality and number of returns. In doing so, it can also identify overlooked or unexpected customer personas, letting you better refine your marketing tactics.

1-to-1 Personalization

Machine learning makes it possible to better understand your customers in the aggregate. But it also allows you to understand them better at the individual level. In this way machine learning powers greater personalization.

The personalization enabled by machine learning can take many forms. It can inform product and content recommendations. It can provide customer service reps with next-best actions. And it can help you alter your marketing tactics by aligning them with individual customer preferences.

Whether you use machine learning for predictions, personalization or something else, the beauty is that, over time, the performance of your models will improve. It is called machine “learning,” after all! This means that, as you feed the results of your marketing activities back into your machine learning algorithms, they get better. It’s a truly virtuous circle.

SECTION 2


How Can Marketers
Apply Machine Learning?



How Can Marketers Apply Machine Learning?

There are countless ways that marketers can apply the power of machine learning. In fact, Omer Artun, Acquia's chief science officer, wrote an entire book about the paradigm shift initiated by this technology, **Predictive Marketing** (Wiley, 2015). We encourage anyone interested in taking a deep dive into this transformational approach to check it out.

For our purposes, let's consider three specific use cases that demonstrate the far-reaching impact of machine learning on marketing.



The Power of Clustering: Cross-Selling, Promotion and Customer Acquisition

Product-based clustering can improve the effectiveness of cross-selling.

Consider this example:

A big sports retailer thought it would make sense to target moms who had bought soccer equipment for their children with promotions for yoga accessories. The assumption the company made was that “soccer moms” were also “yoga moms.” As it turned out, they were mistaken. Indeed, when they looked at the data, predictive clustering showed no real overlap between these two groups.

Predictive clustering did reveal, however, that soccer moms were also baseball moms. The company shifted their cross-selling efforts accordingly to better reflect the actual behavior of their customers.

Similarly, a leading vitamin and wellness brand found that a certain joint supplement had become a big seller. They assumed this reflected the purchasing behavior of their senior customers. They even built campaigns around this assumption. As it turned out, the leading buyers of these supplements turned out to be bodybuilders. The company and its agency reworked their campaigns to more accurately target the real customers of these products.

Just as clustering can help with targeting existing customers, it can also be used to identify prospective customers. This tactic can be particularly effective when buying advertising on Facebook, Twitter or through Google Display Network. The idea is simple. With 100 or more customer records you can train the machine learning algorithm to identify “lookalike” customers and use these insights for ad targeting. You can also optimize your models for specific similarities between targets and existing customers.


For example, a company could create a custom audience consisting of customers with a high likelihood to buy and share that dataset with Facebook. Facebook would then run lookalike modeling on that audience and use that data to target new prospects. **The result: New customers that resemble some of the best customers you already have.**

The Power of Prediction: Discounting, Sales Enablement and Likelihood to Buy

Every marketer knows that discounting can influence purchasing behavior. Machine learning allows marketers to be far more targeted when it comes to discounting and prevent over-discounting.


One online pet pharmacy, like many retailers, followed a set calendar of discounts and promotions for all customers. **Using predictive analytics, however, they determined that while discounts were effective for some customer clusters, they were totally unnecessary for certain customers and totally ineffective for others.**

By ranking customers based on likelihood to buy, the pharmacy was able to improve its discounting strategies. They offered minimal discounts to customers already very likely to buy and larger discounts to those less likely. This approach drove a 38% increase in quarterly sales, a 24% increase in profits, and a 14% increase in customer retention.



Of course, this same method can be applied to optimize operational efficiency in sales and marketing. For example, many software vendors rely on free trials to create leads and drive purchasing. Having your sales team follow up with everyone who has started a free trial doesn't make sense for the simple reason that these potential customers will likely fall into three groups: the simply curious (about 70%); the serious evaluators (20%); and those on the fence (10%).

By applying a likelihood-to-buy model to people who have initiated a free trial, **you can provide salespeople with the insights needed to focus on the 30% where contact will make a difference**, allowing them to ignore the 70% where they would simply waste their time.



The Power of Personalization: Recommendations and Relevance

When it comes to personalization, marketers tend to think in terms of addressing customers by name in communications and the like. While that can be important, and certainly have a negative impact when it goes wrong, it is not the most effective way to personalize. **In fact, marketers are better served when they think about personalization in terms of relevance.**

Machine learning fuels greater personalization by, for example, increasing the relevance of recommendations. The most common form of personalized recommendations in this context are product recommendations made at the time of purchase, generally taking the form of “Customers who bought X also bought Y.”

This model can also be applied to “next-best-purchase” recommendations. For example, a home improvement store found that customers who had built a deck tended to be in the market for a new grill shortly afterward. This insight led them to create a marketing campaign focused on the deck builders and promoting their grill options.

Of course, this type of personalization does not have to be about increasing purchase amounts or driving future purchases. Consider the recommendations made by the online fare comparison engine, Kayak. Kayak uses a price prediction tool that advises customers whether they should buy or wait based on potential price changes. **This type of recommendation is not focused on increasing spend at the moment of purchase but, instead, on improving the customer experience and driving repeat usage of the service.**

Marketers are better served when they think about personalization in terms of relevance.

Countless Applications

The few examples just discussed demonstrate that machine learning can be applied to most key marketing activities. It can be used to drive purchasing behavior by influencing promotions, offers and upsell/cross-sell initiatives. It can be used to acquire new customers and re-engage old customers. And it can be used to improve customer experience across every step of the customer journey.

Given the power of machine learning, what do marketers need to do to harness this power? We'll look at that in the next section.

SECTION 3

**Harnessing the Power
of Machine Learning**



Harnessing the Power of Machine Learning

As the use cases just described suggest, machine learning can have a profound impact on marketing performance and marketing effectiveness. As you also may have guessed, it can have a profound effect on every aspect of marketing operations. Consider the case of using machine learning to predict customer churn by scoring customers based on their flight risk.

Traditionally, if a customer has stopped visiting the company website or opening the company's emails, the company will initiate a reactivation campaign. These campaigns may consist of a series of emails with various offers, but, generally speaking, they will be "one-size-fits-all." With a scoring system in place, however, marketers can begin diversifying their tactics.

First of all, on the predictive front, if a marketer knows that someone is at a high risk for churn, they can proactively address it before the customer has actually churned.

The question is: Should you? For example, let's say you have identified a customer as a "serial returner." If it looks like you might lose that customer, maybe your attitude should be, "Good riddance!"



At the same time, it could be the case that a customer only buys around the holidays. Going overboard trying to reactivate them outside of this season won't make much sense. Instead, you should focus on staying in touch and priming the pump for when the holiday rolls around again.

Finally, if a customer has been loyal and is now at risk for churn, this calls for a much more aggressive response and a tailored campaign focused on recovering their loyalty.

Truly diversifying your tactics based on identified customer clusters and behaviors means that the days of "one-size-fits-all" marketing are coming to an end. What will take its place is predictive marketing. This will require a host of changes in the marketing department. These changes will take the form of new marketing tactics, processes and metrics. But they will also take the form of new approaches and new roles associated with data management and analytics.

And these approaches and roles will demand new tools.

SECTION 4

Choosing the Right
Customer Data Solution



Choosing the Right Customer Data Solution

Machine learning, as should be clear by now, runs on data. This means that harnessing the power of machine learning calls for marketers to get their data house in order. The most effective way to do that is through a customer data platform (CDP). The right CDP will help you aggregate the data you have, ensure its integrity by “cleaning” and deduping it, and continually supplement that data to create ever-more complete customer profiles.

Effective data management, however, is simply table stakes. What is needed are capabilities, ideally built into the CDP, that will allow you to analyze and manipulate customer data in a way that yields the many benefits just described.

When assessing solutions based on their predictive machine learning capabilities, you need to think in terms of three high-level categories:

Quality

How accurate or actionable are the predictions that the solution produces?

Scalability

How much data can the system handle? More importantly, can it scale to accommodate the data requirements of large enterprises?

Adaptability

Can the system adapt to new data types and new marketing problems? Can it handle unpredictable situations that may arise from new technologies, new business initiatives and other unforeseen circumstances?

Each of these attributes is critical. Without quality data, your predictions are meaningless. If the system can't work with large datasets, then it can't provide the customer insights you need. And if it can't adapt, then its value is painfully limited.

The good news is, if you find the right solution, the impact it will have on your marketing is immeasurable.

Conclusion

The era of machine learning and the predictive marketing it enables has arrived. Indeed, machine learning has already proven its value for a number of marketing organizations.

Machine learning solutions will only continue to become more powerful. Marketing leaders need to think about what this means for their organizations. First and foremost, they need to consider what they could accomplish with the right machine learning tools. As we have tried to show, the possibilities are endless.

On the other hand, marketing leaders also need to consider the cost if their competitors take advantage of these tools first. In some cases, of course, this has already happened.

The good news is, these tools are available now. Imaginative and forward-thinking marketers will make the most of them!

Learn how [Acquia CDP](#) can help you harness the power of machine learning to build better experiences for your customers.

ABOUT ACQUIA

Acquia is the open source digital experience company. We provide the world's most ambitious brands with technology that allows them to embrace innovation and create customer moments that matter. At Acquia, we believe in the power of community — giving our customers the freedom to build tomorrow on their terms.



[acquia.com](https://www.acquia.com)

Acquia