



Make the Leap from A/B to AI

The Three Types of Machine Learning & Their Uses in Marketing

A Marketer's Guide to selecting the Right Machine Learning

George Khachatryan

Co-Founder & CEO, OfferFit

Visctor Kostyuk

Co-Founder & CTO, OfferFit

Scott Brinker

Editor, Chiefmartec.com

OfferFit

August 2022

Contents

Foreword by Scott Brinker, Editor at chiefmartec.com	3
What are Artificial Intelligence & Machine Learning?	4
The Three Types of Machine Learning	6
Type 1: Supervised Learning	7
Type 2: Unsupervised Learning	9
Type 3: Reinforcement learning	10
Which Type of Machine Learning is Right For You?	12

Foreword

You've heard about **Artificial Intelligence (AI)** and **Machine Learning (ML)** — transforming marketing for years. But the explanation of what that actually means is too often just high-level, hand-wavy promises.

AI isn't magic, it's machinery. In fact, ML uses relatively straightforward machinery that simply excels at working with large amounts of data and computing in the cloud.

By understanding the three types of ML — supervised learning, unsupervised learning, and reinforcement learning — and how each can be applied to marketing, you'll be better prepared to identify opportunities to harness ML in your work and to evaluate vendor promises of how they can help.

George Khachatryan and Victor Kostyuk, authors of this paper and co-founders of OfferFit, have been successfully employing ML in enterprise analytics for years. They launched OfferFit to apply ML to scale marketing experimentation and conversion optimization, tapping the tremendous volume of data that marketing organizations have within their reach to take A/B testing to a whole new level.

In the next few pages, you'll learn about the specific ML techniques OfferFit uses, as examples of how the different types of ML solve real-world, marketing problems. Supervised learning to make predictions, such as churn propensity. Unsupervised learning to reveal new customer segments. Reinforcement learning to automate experiments at speed and scale. It's powerful machinery. And there's no one better than George and Victor to explain how it works and how it can enhance your business's marketing performance.

Here's to putting marketing technology to good use!

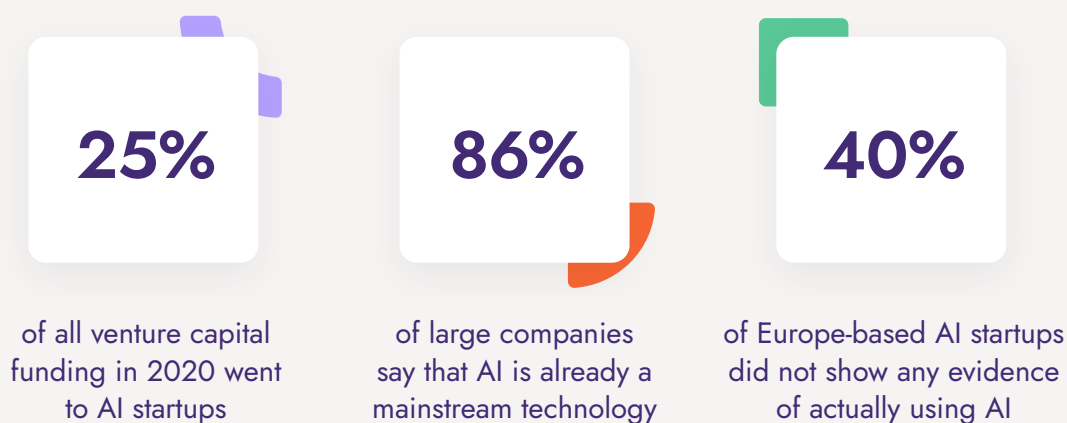
Scott Brinker

Editor, chiefmartech.com

What are Artificial Intelligence & Machine Learning?

Artificial intelligence (AI) is hot. Approximately 25% of all venture capital funding in 2020 went to AI startups,¹ a staggering \$75 billion in total.² Enterprises are adopting AI at a fast pace; according to the Harvard Business Review, 86% of large companies say that AI is already a mainstream technology for them.³

Despite this rapid adoption of AI – or perhaps because of it – there remains a great deal of confusion about what exactly constitutes AI and machine learning in marketing. Most business decision-makers – including marketers, general managers, and investors – do not have computer science degrees, making it hard for them to evaluate the underlying technology used by AI products. The result is a murky industry landscape, with seemingly every company claiming to be using “AI.” A recent analysis by London-based venture capital firm MMC found that 40% of Europe-based AI startups did not use any AI technology.⁴



So what exactly is AI and **Machine Learning (ML)**? AI has always been a vague concept with no universally-accepted definition. The term “artificial intelligence” was first coined in 1956 at a famous workshop held at Dartmouth College. Computer scientist John McCarthy, the organizer of the workshop, defined AI as “making a machine behave in ways that would be called intelligent if a human were so behaving.” At the time, this would be taken to include such tasks as recognizing handwritten digits, playing chess, and making accurate predictions about future events (e.g., customer churn). But this definition was problematic since it relied on public perceptions of what human-like intelligence meant. McCarthy himself later lamented that “as soon as it works, no one calls it AI anymore.”

1 Cunchbase, Global VC Report 2020: Funding And Exits Blow Past 2019 Despite Pandemic Headwinds

2 VentureBeat, VCs invested over \$75B in AI startups in 2020

3 Harvard Business Review, AI Adoption Skyrocketed Over the Last 18 Months

4 Verge, Forty percent of ‘AI startups’ in Europe don’t actually use AI, claims report


What are Artificial Intelligence & Machine Learning?

Machine learning, while synonymous with AI in most people's minds, is a term with a more precise meaning. Machine learning deals with algorithms that make predictions or inferences based on experience (in the form of data). For example, some machine learning models "learn" from past data and then make predictions for new situations.

One critique leveled against machine learning is that it is "just statistics." This criticism is both apt and misguided: apt because machine learning indeed grew out of statistics and can be seen as a natural extension of it, but misguided because it radically understates the extraordinary new capabilities that machine learning has given us.

Machine learning can indeed be seen as a part of statistics, but this is not a disparagement of machine learning, but rather a testament to the power of statistics.

Classical statistics began around the early 1900s. Its primary uses included making predictions based on existing data, inferring numerical values from a limited sample, and assessing the strength of evidence from quantitative experiments. The mathematical methods used by classical statistics were, in general **linear**, – that is, it was assumed that relationships between variables x and y took the form $y = x \times a + b$, for a and b constants. This assumption was made not for lack of mathematical sophistication, but because of computational constraints: it is impossible to do the computations needed for **nonlinear** statistics with paper and pencil, slide rules, or 20th-century pocket calculators.



Machine learning can be seen as an extension of the methods of classical (linear) statistics to nonlinear situations

The 1980s, 1990s, and 2000s saw an explosion of available computing power, which gave impetus for the development of nonlinear statistics. Different communities of researchers tackled the problem, and each gave their own name to the new field – algorithm researchers developed "machine learning," database experts worked on "data mining" and "big data," and statisticians pursued "statistical learning." Each came from a slightly different angle but was solving an overlapping problem: extending the tools of statistics to the new world of enormous datasets and virtually limitless computing power.

The move from linear to nonlinear models was not just an incremental improvement; it enabled the creation of algorithms that solved problems that were inaccessible using the old techniques. Today, machine learning models can recognize speech and images, write human-like text, play board games at superhuman skill levels, and predict consumer behaviors with unprecedented accuracy.

The Three Types of Machine Learning

For marketers to understand the specific uses of machine learning, it is helpful to think in terms of a customary division of machine learning into three types:

Supervised Learning

Predicting an outcome based on historical data

Examples

- Churn prediction
- Customer Lifetime Value prediction

Unsupervised Learning

Identifying patterns in data

Examples

- Cluster analysis
- Customer segmentation

Reinforcement Learning

Experimentally discovering optimal actions

Examples

- Campaign optimization
- Personalization



“For each type of problem in our business that can be solved with machine learning, are we fully leveraging the available tools?”

These are not just different techniques – the three types of machine learning solve fundamentally different problems. In other words, they are used to answer different questions.

This distinction is helpful for marketers since it helps them go beyond asking the simplistic question of “Are we using machine learning?” A better question is, “For each type of problem in our business that can be solved with machine learning, are we fully leveraging the available tools?”

The rest of this article will discuss the three types of machine learning and how each is applied in marketing.

TYPE 1: Supervised Learning

Supervised learning involves predicting an outcome based on historical data. You start with a collection of data points that include “labels” representing the outcome of each.

For example:

- 100,000 customers with different features for each customer (e.g., how long they have been with each company) and labels (outcomes) (e.g., whether each customer churned in the subsequent 12 months)
- 50,000 photographs of pets: With labels (outcomes) (e.g., whether the photograph is of a dog or cat)

This historical data is called the “training set.” Using this dataset as the guide, the machine learning model is built to predict the labels as accurately as possible. This is why this type of machine learning is called supervised learning – the “labels” on the historical dataset are providing supervision, like a teacher who gives a student feedback based on their answers to questions. The goal of the machine learning model is to provide accurate predictions of labels beyond the training set – i.e., to correctly predict labels for new, unlabeled data.

As an example, consider this dataset on customer churn:

	Name	Years as Customer	Annual Spend	Customer Service Calls	Churned in 12 Months
Training Data	Mickey Mouse	17	\$53	2	No
	Chapulín Colorado	23	\$112	17	Yes
	—	—	—	—	—
	Cheburashka	3	\$12	0	Yes
New Data	Dora the Explorer	8	\$83	3	?

Here we have historical data for some customers (in a supervised learning scenario, we would need thousands or millions of data points to create an accurate model). Each customer comes with a set of “**features**”: how long they have been a customer, how much they spend with us per year, and how many times they have called customer service. In real-life applications, we would usually have many more features for each customer, typically dozens or even hundreds. For our training set, we also have labels showing whether each customer churned in the subsequent 12 months.

When we are given a new datapoint (in this example, Dora the Explorer), the model will predict what the label would be – in this case, whether the customer is likely to churn in the next 12 months.

When marketers wish to make predictions for the future based on past experiences, supervised learning can help. Examples include:

- Predicting customer churn (as in our example)
- Predicting whether customers will be receptive to upsell or cross-sell offers
- Predicting Customer Lifetime Value (CLV) for various marketing tactics (which is useful to decide how to allocate budget across different paid channels)

To implement such predictive models, marketers can work with a company’s data scientists to build models in-house, hire consultants to create models, or purchase software products that include this predictive capability. Some examples of software companies that provide this are **VOZIQ** (for churn prediction) and **Retina.ai** (for Customer Lifetime Value prediction).

TYPE 2:

Unsupervised Learning

Unsupervised learning involves identifying patterns in data. Here, there are no “labels” you are aiming to predict, which is why this type of machine learning is called “unsupervised” – there are no labels providing supervision.

One type of unsupervised learning is **cluster analysis**: identifying the natural groups into which a dataset can be divided. For example, you might start with a dataset of species of living organisms and use a cluster analysis model to determine what natural groups they fall into; if the model works well, you’d expect that it might divide them into groups along the same lines as biologists have defined – e.g., the kingdoms of animals, plants, fungi, protists, and monera.



Cluster Analysis

Cluster analysis algorithms identify the natural groups into which a dataset can be divide

In marketing, a natural application of cluster analysis is to create fact-based customer segmentations. For example, consultancies like the **Boston Consulting Group** and **McKinsey & Company** help companies define segments by combining this type of model with qualitative research: you can identify the natural segments in the customer set discovered by a machine learning model, and then interview customers in each segment to understand what common underlying qualities really define this group of customers. A second, even more famous, application of unsupervised learning in marketing is **lookalikes**. For a large group of customers (e.g., the full population) and a small subgroup, you can use unsupervised learning to identify which customers in the full population look similar to the small subgroup. Platforms like **Facebook** offer this type of machine learning to allow advertisers to build lookalike audiences to target individuals similar to their existing high-value customers.

TYPE 3:

Reinforcement Learning

Reinforcement learning (which is OfferFit's focus area) automates the process of experimentation. Reinforcement learning models (which are also called "agents") select "actions" and then receive "rewards" from the environment. The models proactively experiment to discover which actions yield the highest reward.

This is somewhat akin to supervised learning: the rewards in reinforcement learning serve a similar purpose to the labels in supervised learning. However, in reinforcement learning, there is a different division between labeled historical data and unlabeled new data – the agent might start from scratch with no reward and then learn over time as it gets rewarded for new actions. Supervised learning models make predictions, while reinforcement learning models make decisions.



A central concept in reinforcement learning is the "**exploration-exploitation tradeoff.**" Every time an agent needs to select an action, it has a choice: go with the reward-maximizing option based on the agent's current knowledge, or experiment to attempt to discover even better actions. Agents who always side with the current best-looking option are called "greedy," and they usually end up underperforming over time because they don't explore and learn fast enough. On the other hand, agents that predominantly explore are also suboptimal, since they miss opportunities to gather reward. They're like the cousin everyone has who keeps getting master's degrees but never holds down a real job for longer than a year or two. The models that maximize long-term reward are the ones that find the perfect balance between exploring and exploiting; they continuously learn in the most sampleefficient way possible while concurrently putting their knowledge into practice.

The most famous example of reinforcement learning is AlphaZero, a model created by DeepMind (a company now owned by Google). In 2017, AlphaZero learned to play Chess and Go from scratch by repeatedly playing against itself. By experimentally discovering the best plays, it quickly reached unprecedented levels of skill in both games, defeating the best human players as well as all previous computer models.

TYPE 3: Reinforcement Learning

Reinforcement learning has reached maturity more recently than supervised and unsupervised learning; the last several years have seen active research in academia in this field, and now cutting-edge reinforcement learning models are finally starting to see commercial applications.

Examples of applications of reinforcement learning in marketing include:

- Automatically optimizing digital ad campaigns
- Personalizing email creative (including subject lines)
- Selecting the best channel for each individual customer
- Personalizing lead-nurturing sequences

Essentially, any time a marketer might consider A/B testing, reinforcement learning is a more powerful alternative. This is why OfferFit's motto is "from A/B to AI" – we believe that reinforcement learning will replace A/B testing in marketing.

The simplest type of reinforcement learning model – and the kind most familiar to marketers – is a **multi-armed bandit**. This type of model chooses among several options (such as different subject lines for an email) and automatically adjusts its choices based on real-time results (e.g., resulting email open or click rates). A multiarmed bandit is like a more dynamic version of an A/B test, and it is very simple to implement.

However, multi-armed bandits are very simplistic: they do not take into account differences between customers or other aspects of the environment. For example, a multi-armed bandit that selects the best email sending time will not adjust its recommendations for different customer profiles or days of the week (e.g., Monday vs. Saturday). For this reason, more advanced types of reinforcement learning (which OfferFit utilizes), such as **contextual bandits** or **deep-Q networks**, are better suited for most marketing applications.

Different marketing applications call for different types of machine learning. So should you use supervised, unsupervised, or reinforcement learning?

Which type of Machine Learning is right for you?

The three types of machine learning are not alternatives to each other – i.e., it doesn't make sense to ask the question, "Which type is generally better to use?" Instead, they exist to solve completely different problems.

Are you trying to predict something?

Example

Churn prediction

Use

Supervised learning

Are you trying to identify patterns or similarities?

Example

Lookalike identification

Use

Unsupervised learning

Are you trying to automate the process of experimentation?

Example

Campaign optimization or personalization

Use

Reinforcement learning

This means marketers seeking to reap the full benefits of machine learning should separately ask about each type of machine learning: Are we using it to the fullest to optimize our business performance?

Today, many marketers at enterprises are already using supervised learning to create propensity scores for things like churn, repurchase, and cross-sell. Some marketers are also using the most basic type of reinforcement learning – multi-armed bandits – to optimize campaigns. However, this leaves a lot of opportunities on the table. By adding unsupervised learning and moving to more advanced types of reinforcement learning, marketers can significantly improve their understanding of the customer and the performance of campaigns.

Want to learn more?

Visit us at offerfit.ai or email us at hello@offerfit.ai!

About OfferFit

OfferFit accelerates the creation of knowledge. Trial and error has always been the core of human progress. At OfferFit, we automate experimentation using reinforcement learning, a type of self-learning AI, to make knowledge creation faster than ever before.

A/B testing can be effective for lifecycle marketing, but it's slow and doesn't scale. Lifecycle marketers use OfferFit to radically accelerate experimentation. Marketers choose options for multiple dimensions such as messaging, creative, incentive, channel, and timing. Then OfferFit experiments to discover the best-performing recommendations for each customer.

OfferFit's Automated Experimentation Platform lets marketers unlock the value in their first-party data and maximize whichever KPIs are most important to their business.



Make the Leap from A/B to AI

offerfit.ai
hello@offerfit.ai

About the Authors

George Khachatryan
Co-Founder & CEO

Before founding **OfferFit**, George was an Associate Partner at McKinsey & Company, where he led transformations of some of the world's largest companies. Previously, George co-founded a technology company whose products are used today by over 1 million users. He holds a PhD in Mathematics from Cornell University.

Victor Kostyuk
Co-Founder & CTO

Before founding OfferFit, Victor was a Lead Data Scientist at the Boston Consulting Group. He is an expert in self-learning AI, with experience serving Fortune 500 companies on offer personalization. He holds a Ph.D. in mathematics from Cornell University.