SEARCH AND DIGITAL FROM VERTICAL LEAP EST 2001

A marketer's guide to machine learning and AI

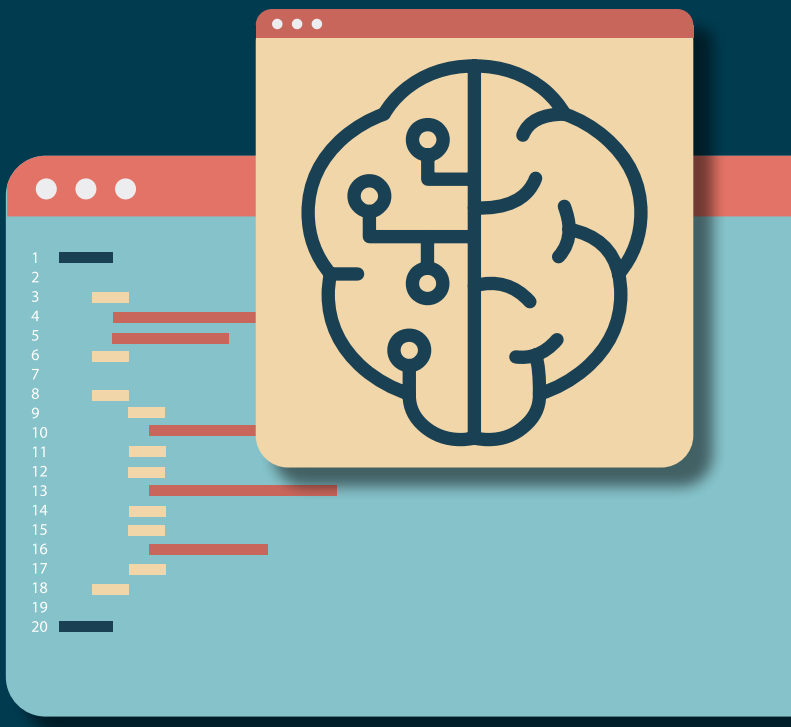

Vertical Leap, 6 Boathouse,
College Road, HM Naval Base,
Portsmouth, PO1 3LR

Visit **vertical-leap.uk**

Call us on **023 9283 0281**

Email **info@vertical-leap.uk**

Contents

Machine learning is something every marketer needs to come to terms with – and the sooner the better. It's set to be one of the biggest technology shifts the industry has seen since the dawn of the internet and it's going to change everything you do as a marketer.

Up to 75% of executives plan to implement AI within three years, missing out today could mean losing a competitive edge in the future. Over the next five years, marketing leaders expect AI to transform and improve these business areas:

AI could lead to significant transformations within your marketing department. However, embracing AI-driven marketing is a tough call for most companies; only 26% of business leaders have confidence in their organisation's ability to create a distinct, result-driven AI business strategy.

In this guide, we will

- address the strategy for unlocking the power of big data using machine learning;
- examine artifical intelligence, machine learning and deep learning;
- discuss common implementation mistakes;
- present a roadmap for defining and developing your AI-driven marketing strategy.

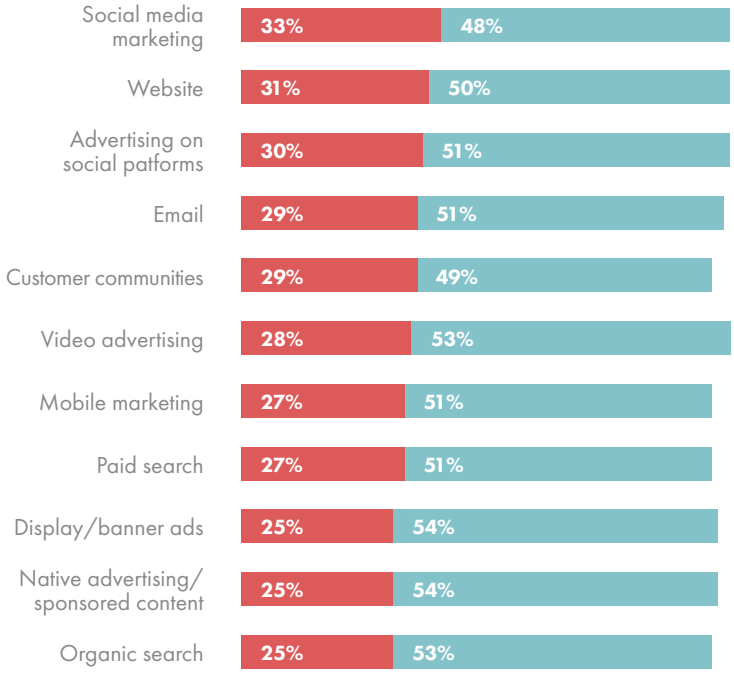

Extent to which marketers coordinate messages in the following channels with other channels

Messages evolve across channels based on customer actions

Identical messages are broadcast across channels

Source: Salesforce

What are machine learning and AI?

The terms machine learning and artificial intelligence are often used interchangeably, but they're not quite the same thing.

Artificial intelligence (AI) is human-like intelligence showcased by machines.

An "intelligent machine" is deemed to be one that can mimic certain cognitive functions typically inherent in human minds, such as problem-solving or learning.

Artificial intelligence, in the sense of rules-based programmes capable of basic "intelligence", has been around for decades. As information became more digitised and readily available via the internet, "mining" knowledge became simpler. Now you can "feed" data to the smart algorithms and have routine tasks such as image or voice recognition performed by machines.

Machine learning is one of the methods that achieves artificial intelligence. It enables computers to learn directly from experience/examples presented in the form of data, instead of merely following the programmed instructions.

Over time, as the system "learns", it becomes capable of optimising its performance and improving by its own accord.

An "intelligent machine" is deemed to be one that can mimic certain cognitive functions such as problem-solving or learning.

For example, while it's easy for humans to differentiate a dog from a cat, it's not been so simple for a machine. Now though, modern machines are gradually getting better – in 2015 a neural network managed to pass a benchmark test in image recognition with 95% accuracy, exceeding human performance for the first time.

Progress can be largely attributed to deep learning – a subfield technique for enacting machine learning. Here's how Qualcomm visualises the three technologies:

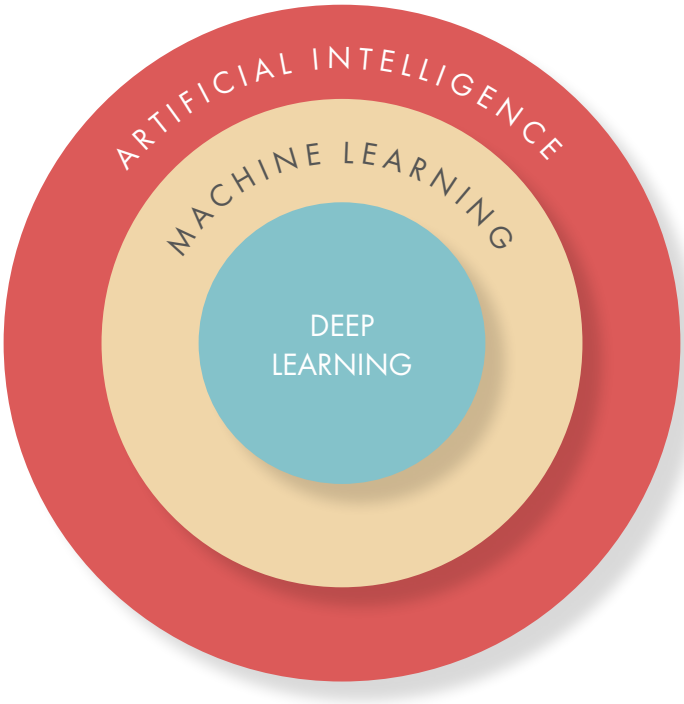

Source: Qualcomm

Today's most advanced AI products are powered by **deep learning**, and, more specifically, **deep neural networks (DNNs)**. In a nutshell, the code structures of DNNs are grouped in layers, modelled loosely on the human brain. They are designed to recognise patterns and learn the patterns of patterns.

Based on actual neural networks, DNNs include multiple hidden layers between the input (your data) and output (desired result).

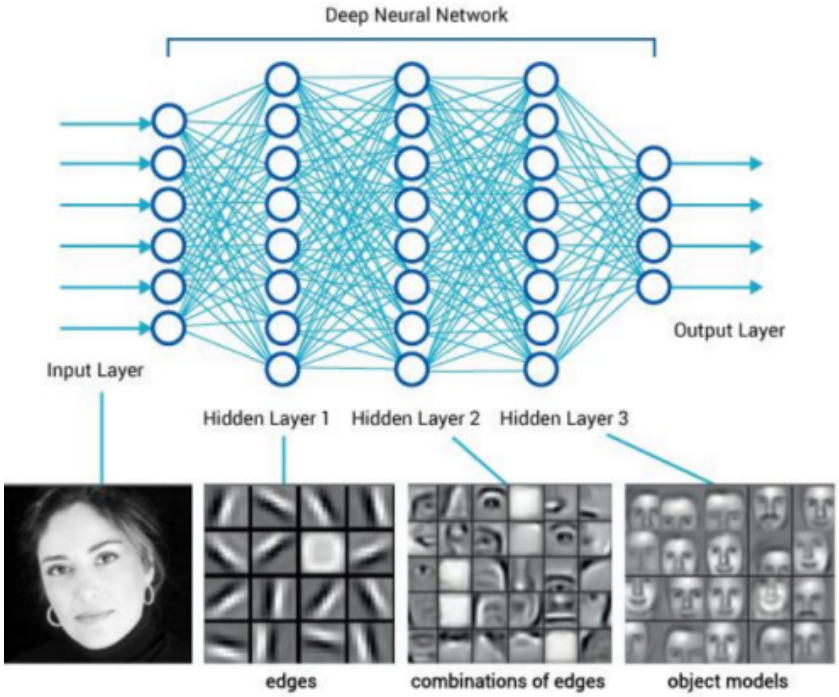

Source: Gitbook

Each layer helps process the data. The first layer looks for the most obvious repetitive patterns – black circles or smooth lines. Then, the first layer outputs the "knowledge" to a new layer, capable of analysing more complicated patterns – combining the smooth lines into some parts of the picture. Then a more advanced level goes on top of that, and so on until we receive the desired output.

It's important to note that the network itself decides what each layer will be analysing. The network designer merely indicates how many layers there are, how they are connected and whether any additional processing will be done to improve the results.

So, deep learning enables more advanced machine learning and, ultimately, better artificial intelligence.

Where machine learning is heading

Machine learning is already powering our search engines, translation software, speech recognition and even the discovery of new pharmaceutical drugs. We now have "smart" platforms that continue to automate, upscale and refine tasks to a level that humans can't keep up with.

Machine learning isn't just crunching more numbers in less time; it's making smarter decisions and even learning from its own mistakes.

From a marketing perspective, advancements in machine learning stand for multiple improvements:

- **Truly data-driven decision-making.** Today, 85% of executives admit being only partially effective at meeting their goals for data and analytics initiatives. Smart algorithms can significantly improve data analysis procedures and ultimately function as autonomous entities without the need for constant human supervision.
- **Intelligent agents driving customers' purchase behaviour.** Most customers remain unimpressed by the efforts of companies, attempts at personalisation. Machines can improve data collection and analysis, online behaviours, preferences, transactions and even emotions. These insights could be applied to influence consumers' opinions and decisions online.
- **Improved product development.** A future where computer programs are building other computer programs is now a genuine possibility. Websites that design themselves, marketing strategies that create themselves and ads that target themselves could be the future we're facing. The best part is that business owners can precisely estimate and compare the costs of creating certain elements of their product. For instance, the new computer-aided design system, Dreamcatcher, launched by Autodesk, allows comparing the cost, function, and material of different virtual prototypes before manufacturing. The AI-powered algorithm "remembers" the impact of each element and can offer advice on product optimisation.

The impact of AI and machine learning on marketing so far

The most obvious impact of machine learning for marketing to date is how the tech giants are applying it to enhance their own products. Various Google apps are now partially powered by AI and the same goes for Facebook, Apple and Amazon.

Up to 88% of customers expect to receive a personalised experience from a brand, yet, only 45% of them think marketers do personalisation well. With the help of machine learning, you can second-guess your customers' moves and make personalised suggestions based on past data.

Netflix recently discovered that its customer typically gives up searching for a film in 90 seconds. It deployed an advanced recommendation and search engine to deliver relevant suggestions within a few seconds. Netflix believes this has saved more than $1 billion of lost revenue in cancelled subscriptions.

But, personalisation is only one side of the machine learning revolution.

Up to 88% of customers expect to receive a personalised experience from a brand, yet, only 45% of them think marketers do personalisation well.

The more exciting impact for marketers is how this technology is empowering us to do more with data, on a much wider scale. The SEO team at Vertical Leap is using machine learning to dig deeper, achieve more and operate at a scale that would not be possible with humans alone.

Big data doesn't need large teams of analysts anymore. If you can build the right algorithm, you can process as much data as you need to.

Which means big data is no longer reserved for big brands.

Not all marketers are going to create their own custom algorithms, of course. This year, up to 80% of companies plan to rely on IaaS (insights-as-a-service) providers for analytical capabilities. Right now, the selection of third-party machine learning marketing tools isn't huge, but they're on the way.

How will machine learning help you as a marketer?

Machine learning is changing the way we approach data management in marketing.

Unlock the power of big data

Before machine learning, programmatic computing was limited to set rules and software engineers had to code every function, variable and parameter. This is slow and vulnerable to human fallibility.

With machine learning, you feed masses of relevant data to algorithms and leave them the modelling, based on past "knowledge". With a minimal risk of human bias or error, machine learning algorithms can accurately predict what a potential customer is about to do next and reccomending the best action.

For example, flagging a user with a 90% chance of churning, pinpointing that users exhibiting similar behaviour are most likely converted by remarketing campaign A, and so on.

Intelligent automation

Machine learning enables a more intelligent automation. Regular automation simply completes predefined tasks without our input.

Machine learning takes this one step further by improving its own workflow, learning from past experiences and becoming more efficient over time.

MonkeyLearn offers a machine learning add-on for Google Sheets that analyses spreadsheets without users needing to write a single line of code. Proof that machine learning can be simple, lightweight and still hugely powerful. Machine learning doesn't really care how much data you feed it; it's more interested in the quality of your data.

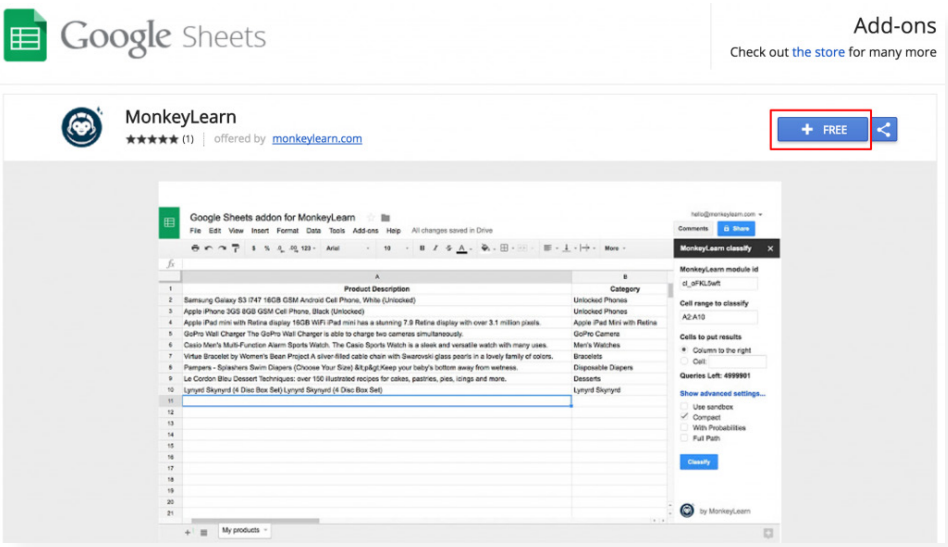

Advanced data management

The biggest challenge marketers have faced over the past decade has been organising the growing haul of data available.

While 98% of business leaders encourage employees to make data-driven decisions, only 23% of them claimed that their company was extremely successful in leveraging big data. Indeed, managing hoards of unstructured data requires large teams of specialists and a lot of time.

Now, machine learning is the next technology to offer a working solution to this problem.

It's not as if you can feed a bunch of unstructured data to your algorithm and expect it to do everything for you. But we're already seeing the emergence of AI marketing tools designed to process chunks of data from specific sources like Google Analytics.

Creating an algorithm that knows how to process data from a single source is one thing. Combining five, ten or one hundred algorithms to process data from any number of different sources turns seemingly unstructured data into something incredibly powerful.

Acquisition channels

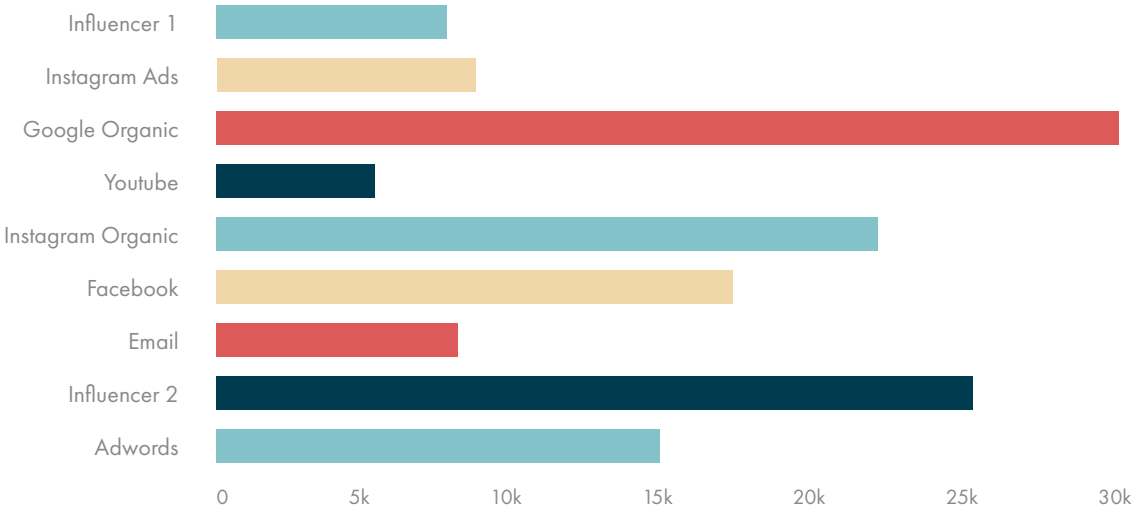

Email visitors who landed on /about/ are 97% more likely to convert.

Over multiple visits, users from Facebook are 84% more likely to convert when returning. Retarget Facebook users, or nurture them using email campaigns.

PaveAI processes Google Analytics data from up to 300,00 users per month and turns it into actionable insights.

Informed decision-making and forecasts

We've talked about automation and machine learning to handle large volumes of data. We developed Apollo Insights; our machine learning platform, to perform all those tasks, combining data from multiple sources. Apollo audits every page of a website, checking every backlink and checking every query against every page, every week, AND filters data to spot new opportunities for growth. But machine learning has more to offer than a more advanced form of data tracking and analysis.

One of the most exciting developments for marketers has been **predictive analytics**. By analysing the behaviour of website visitors, platforms like SAS are able to calculate the value of different leads and predict their next steps. This can warn you about leads likely to drop off, highlight cross-selling opportunities and pinpoint UX issues that are getting in the way of sales.

Let's take Uber as an example. The taxi service has mastered the art of predicting customer demand over Christmas and other holiday periods, allocating and optimising resources with the help of neural networks. The forecast's errors were less than 11% during the most unpredictable holiday – Christmas, and the average error rate was less than 10% for the other 11 holidays analysed.

This has allowed Uber to deliver excellent customer experience by allocating just the right amount of resources, and to save money by avoiding over-delivery.

Machine learning reduces human bias. As marketers, we're used to exploiting the cognitive biases of consumers to sell products, but we also fall into the same traps when it comes to our own decisions.

A classic experiment conducted at Stanford revealed that we tend to massively overestimate whether another person knows what we know, which often makes our communication ineffective. One group of participants was asked to choose a popular song like "Happy Birthday" and tap the song beat to the second group. "The tappers" predicted that 50% of the second group would be able to guess the song. In reality, only 2.5% named the tune correctly.

As a marketer, you may be prone to the same bias – overestimating your prospect's ability to understand your message.

AI insights

Decision-making isn't the only thing vulnerable to human weakness.

Our ability to spot things is often compromised by biases, distractions, fatigue or a lack of knowledge. Sometimes, you simply spend too long looking at the same set of numbers to see what they're telling you.

An AI-powered algorithm doesn't have these same weaknesses. It's capable of spotting patterns that are too complex for the most proficient data scientists to see. Machine learning's only handicap in this regard is the quality of data it has to work with.

Highly-targeted ad campaigns

We're already seeing the impact machine learning can have on advertising campaigns with Facebook's lookalike audiences and similar audiences for AdWords. The two advertising giants look at the behaviour of your existing customers and then find other users displaying the same habits, allowing you to target them directly.

There are a couple of important things happening here. First, AI-powered algorithms are tracking user behaviour and attributing probable actions. They are also connecting you to entirely new leads by matching similar behaviour patterns in other users, allowing you to cut through the noise and target potential customers right away.

Taking this process forward, we're now seeing ad and website personalisation finally become a reality, thanks to machine learning. Platforms like Optimizely X mean you can provide a personalised web experience based on the ads users click, the search terms they use and other previous actions.

With Optimizely X you can target users with custom experiences based on their previous actions.

Machine learning is making it possible for us to adapt ads and web pages based on user intent signals, as they happen.

Google identifies these signals as micro-moments – the moment when a consumer grabs his phone to satisfy a certain impulse – to check traffic, hail a cab, or receive information. On average, that happens around 150 times per day. Successful marketers are already tracking pivotal micro-moments and leveraging them to improve their customers' journey.

Red Roof Inn received data that some 90,000 passengers in the U.S. are stranded every day due to flight cancellations. Their team developed an algorithm to track real-time flight delays across different locations and trigger search ads for Red Roof Inns near the periled airports.

Targeting an "I-need-a-hotel-ASAP" micro-moment has resulted in a 60% increase in booking across non-branded search campaigns. Impressive, right?

IDC predicts that ad personalisation through machine learning will boom by 2020, with most businesses investing in the following areas:

Machine learning delivers applications today, with more to come by 2020

What communications applications do you think machine learning technologies can (1) deliver today or (2) will in addition be able to deliver in 2020?

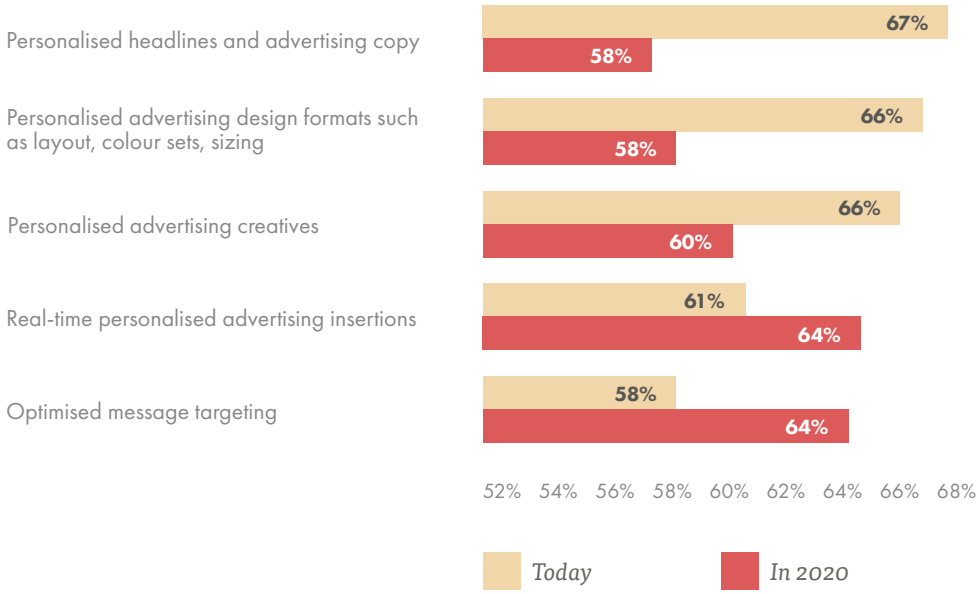

n=379-381. Note: percentages do not add up to 100% as respondents could respond both to (1) today and (2) in 2020.

Source: IDC

Common machine learning mistakes to avoid

With machine learning being a relatively new technology there are plenty of mistakes to be made – especially if you're developing your own algorithms. Mistakes are all part of the learning process, but here we want to give you a heads up on some common problems to avoid.

Starting without good data

This is the most important rule for any project involving data processing. You can't get valuable insights from the wrong set of information; you must start with good data or everything you do with machine learning will be a waste of time.

So, what do we mean by poor data?

- Noisy data - including volumes of conflicting or uncertain info.
- Dirty data - containing incorrect, inconsistent or missing values.
- Sparse data - data with too many values at zero or no value at all.
- Inadequate data - incomplete or biased datasets.

First, you'll want to create a system for collecting the right data, evaluating it and then cleaning it, when necessary. Then, you'll need to think about data preparation and integration between your machine learning algorithm and other platforms.

Data also needs to be formatted in a way that machine learning can work with.

You must start with good data or everything you do with machine learning will be a waste of time.

Implementing machine learning without a strategy

Agencies want to say they're using machine learning, business owners love the sound of using it and the tech giants are doing their bit to spread the excitement.

All of this is fine until marketers start implementing machine learning without a solid strategy in place. Don't jump on the bandwagon. Before you get started with machine learning, determine your business objectives.

Pinpoint areas where your current data handling process is coming up short and then devise a strategy to improve results with machine learning.

Data leakage

In a machine learning sense, data leakage is what happens when external data makes its way into the dataset you're using to train your algorithm, distorting the outcome.

In many cases, this happens when the intended result of a calculation ends up in the training data. For example, if you're building an algorithm that predicts the price of cars and the actual prices inadvertently end up in the training data, you've got yourself some data leakage. This might sound easy to avoid, but data leakage is surprisingly difficult to prevent.

It goes without saying that thorough testing and debugging is vital to creating reliable machine learning algorithms. Once again, having data scientists on board is the only way to go.

Sampling bias

Sampling bias happens when you train your algorithms using biased or non-random datasets, which skew the resulting hypothesis.

Going back to cars for a moment, a crude example would be training an algorithm with a sample of images of new luxury cars and then testing it against a selection of cars in a second-hand dealership. The algorithm will fail to distinguish between a new and a used car or to provide any classification at all.

In theory, sampling bias is easier to deal with than data leakage, but humans have a reputation for being biased samplers. Don't fall into the same trap.

How to get started with machine learning

1 Get your data scientists on board

Whether you're relying on third-party tools, custom machine learning algorithms or a mixture of both, you're going to need a good team of data scientists. Scouting talent isn't easy in this emerging field, but it's an absolute priority. If you can't build your own team, outsource to an agency that's already using machine learning to enhance its marketing services.

If your data handling consists of little more than Google Analytics, AdWords, and a CRM system, an agency is probably the place to start.

2 Run a data audit

Next, you need to identify where machine learning can improve your data management. This starts with a data audit to assess your brand's data maturity, readiness to adopt machine learning and opportunities to use the technology.

You'll want to look at the following areas in your audit:

- Data governance
- Data warehousing
- Data scientific research capabilities
- Data hygiene
- Data sources, uses, and volumes

It's fine if you don't have much in the way of data governance or sophisticated research capabilities. As we said earlier, machine learning isn't reserved for major brands collecting hoards of data on a daily basis. Even a relatively simple machine learning algorithm that visualises your Google Analytics and AdWords data more effectively can make all the difference.

It doesn't matter how big or small your data problems are. The key is to identify them and then explore how machine learning can help you overcome them.

3 Create your machine learning strategy

Gartner suggests starting with defining the end objectives first, such as:

- Develop an integration with tools X, Y, Z to gather up-to-date data on rankings,
- Identify the key metrics worth tracking for your company,
- Gather historical data on the correlation between rankings and number of clicks, conversions per each tracked page and average sale number.

The key here is to build a problem or task taxonomy outlining exactly what tasks you wish to solve with machine learning-powered algorithms and design your data science pipeline around them.

4 Prepare yourself for machine learning

Chances are, you'll have to make certain changes to introduce machine learning into your current marketing setup. You'll probably need to rethink your approach to handling and storing data. There's also a good chance you'll need to clean up your existing datasets.

You may find some of the tools you use simply weren't built to scale up to the volumes of data you'll be dealing with in the foreseeable future.

So, do you go for an agency that specialises in machine learning marketing solutions or build your own system in-house?

Building a custom algorithm in-house assumes allocating the budgets for a development and data science team; establishing a realistic timeline (six months plus) and ensuring there's proper supervision through the development stage. In most cases, you will first have to settle for an MVP – a simplistic algorithm doing one thing well e.g. ranking leads in your CRM based on a set of indicated parameters – before you move on to adding additional capabilities and integrations.

Opting for an existing solution means that you receive immediate access to a vast array of features and capabilities, and benefit from the future updates as long as your subscription lasts. That's exactly what we offer to our Apollo Insights clients – fast access to AI innovations in search marketing without the need to expand your IT department.

About Vertical Leap

Vertical Leap is the most effective search agency in the UK. We use data science and machine learning combined with experience and expertise to dig deeper, achieve more and operate at a scale that is at least 4x the capability of our nearest competitor.

Our Apollo Insights technology ensures that no stone is ever left unturned when it comes to your search marketing campaign. It collects everything that could possibly be known about your website, your wider online presence, your competitors and your overall market. It combines, verifies and analyses all this data to provide both prescriptive and predictive insights that identify all the opportunities and threats that you are facing. It is infused with artificial intelligence that provides our search specialists with a level of analysis that would simply not be possible manually; more than humanly possible.

It means that we spend more time implementing a finely targeted set of activities that are quantitatively proven to get the greatest impact. We spend more time on the details that will yield the greatest results. We spend more time on what matters most – your business.

For over 15 years, we have been leading the way in search marketing, pioneering innovative techniques that keep our customers ahead of their competitors. Our customers range from SMEs to enterprise-level, including the likes of P&O, Harvester, Ordnance Survey, Foyles and University of East Anglia.

We are part of the Sideshow group; an independent group of 170 staff with an annual turnover of £15m and offices in Portsmouth, London and Bournemouth. We are a RAR top 100 agency and a Premier Google Partner.

Contact us

If you want to understand the benefits that machine learning can make to your search campaigns, get in contact and we'd be happy to show you.

email:
info@vertical-leap.uk

or call:
023 9283 0281

Terminology

Artificial intelligence:

Human-like intelligence showcased by machines with the help of technology.

Big data:

Large data sets that require dedicated storage, processing and application systems.

Data governance:

An organisation's approach to handling, processing, protecting and using data.

Data hygiene:

Data is considered clean when it is relatively error-free and complete.

Data quality:

Quality data is relevant, clean and sufficiently structured for the required application.

Data warehousing:

The aggregation of structured data sources so they can be compared and analysed.

Deep neural network:

A set of algorithms designed to simulate the brain's ability to recognise patterns. DNNs usually involve multiple layers with varying tasks, passing the outcome from one layer to the next to complete a more complex set of micro-decisions.

Machine learning:

One of the methods to achieve artificial intelligence, which enables computers to learn directly from experience/examples presented in the form of data, instead of merely following the programmed instructions.

RankBrain:

A machine learning-powered algorithm that helps Google better understand complex search queries and determine the most important ranking factors for different search types.

WEBSITE

vertical-leap.uk

Vertical Leap, 6 Boathouse,
College Road, HM Naval Base,
Portsmouth, PO1 3LR