



The Business Applications of Machine Learning

How businesses can leverage machine learning to make data-driven decisions and gain a competitive edge.

12/6/2018

Author: Jack Perales

Intro

In recent years, the term machine learning has risen to prominence, now residing amongst phrases such as serverless architecture, agile development, and continuous deployment in the great pantheon of technology buzzwords. Similar to its predecessors, the term “machine learning” is apt enough on its own to convey a basic meaning, yet is also vague enough to suggest additional nuance. In truth, “machine learning” is somewhat of a misnomer. It does not, as the name would imply, have anything to do with providing machines the ability to learn autonomously; that’s more within the realm of artificial intelligence (AI). Rather, the discipline surrounds the strategies programmers have developed to train computers to draw accurate conclusions from data through a rudimentary learning process of trial and error.

That being said, how does it work, and why are businesses so eager to leverage it?

An Overview of Machine Learning

The idea of machine learning derives itself from how humans learn and modify their behavior based on experiences. For example, children learn exactly what a “cat” is by observing what people do consider cats (house cats, bobcats, big cats) and what people don’t consider cats (dogs, ferrets, raccoons). Through this, they learn to classify a mammal with a similar set of features as a cat. While computers can’t understand relative size, body shape, and color the same way humans do, they can understand numbers. Being that most images found online are simply arrays of pixels, each of which has numerical value, computers process images and come to understand what mathematical features an image of a cat has that other images don’t.

How exactly does a computer accomplish this? It usually involves an algorithm, a large amount of data, and multiple learning periods. For example, say we’re trying to classify images as containing a cat or a dog. There are many suitable machine learning algorithms for this classification, but we can use a neural network for our example here. In general, a neural network can use a series of matrix transformations (remember that images are simply matrices of pixel values) to turn a large X by Y matrix into a much smaller 1x2 matrix with two values. The first value is the designated probability the image is a cat, the other the probability the image is a dog. It will use this value to probabilistically

determine which category the image falls into.

Just as a young child will zealously label anything as a “cat”, our algorithm will be horribly wrong to start. However, much like children, it can learn with the right kind of reinforcement. Using a technique called back propagation, the values within the neural network, called “weights”, will be systematically adjusted to favor producing a higher “cat value” rather than a “dog value” when given an image with a similar numerical profile. After being trained like this on up to thousands of images, the weights will be finely tuned to have a higher success rate of calculating an accurate result. The neural network, even well trained, will likely not do well in edge cases, such as when part of the animal is shown or where the animal blends into the background—unless it’s explicitly trained for it. However, for the most part, the algorithm will have an accurate understanding of the different categories.

How Does This Relate to Businesses?

While I would not fault anyone for viewing the cat example as more of a useful teaching analogy than a serious application of machine learning, the Zoological Society of London has derived a lot of use from the image-parsing capabilities of computers. Partnering with Google Cloud, they’ve developed programs to identify and classify animals in the thousands of images they acquire through Camera-traps, which allows them to gather crucial data while sparing their employees from what would otherwise be weeks of repetitive, dull work.

Machine learning shines in cases where the rules for doing repetitive tasks that cannot be distilled into simple “if-then-else” cases, but can be done based off information that exists in past data, usually using either a classification or regression (prediction) algorithm. For example, while there’s no hard rules for whether or not to buy a stock, a regression algorithm can use recent trend data to determine which prices are ideal for buying and selling. Additionally, while it is hard to be certain what products to recommend to a customer based solely on their purchases, a classification algorithm can group them with similar customers to get more detailed insights on their potential interests.

While humans could certainly do these tasks, they are often either monotonous or require objective consideration of so much accumulated data that from which it is nearly impossible to derive a sound conclusion. By implementing machine learning, companies can benefit from faster turnarounds, new insights from

objective data analysis, and decreased resource utilization.

How Does This Relate to Businesses?

While I would not fault anyone for viewing the cat example as more of a useful teaching analogy than a serious application of machine learning, the Zoological Society of London has derived a lot of use from the image-parsing capabilities of computers. Partnering with Google Cloud, they've developed programs to identify and classify animals in the thousands of images they acquire through Camera-traps, which allows them to gather crucial data while sparing their employees from what would otherwise be weeks of repetitive, dull work.

Machine learning shines in cases where the rules for doing repetitive tasks that cannot be distilled into simple “if-then-else” cases, but can be done based off information that exists in past data, usually using either a classification or regression (prediction) algorithm. For example, while there's no hard rules for whether or not to buy a stock, a regression algorithm can use recent trend data to determine which prices are ideal for buying and selling. Additionally, while it is hard to be certain what products to recommend to a customer based solely on their purchases, a classification algorithm can group them with similar customers to get more detailed insights on their potential interests.

While humans could certainly do these tasks, they are often either monotonous or require objective consideration of so much accumulated data that from which it is nearly impossible to derive a sound conclusion. By implementing machine learning, companies can benefit from faster turnarounds, new insights from objective data analysis, and decreased resource utilization.

By implementing machine learning, companies can benefit from faster turnarounds, new insights from objective data analysis, and decreased resource utilization.

However, with such a broad use of applications and a wide array of potential benefits, it can often be challenging to convince key decision makers of the exact value it might provide to one's business. Fortunately, the topic of machine learning is not impossibly broad. While machine learning is regularly being utilized in unique and novel ways, there are definitely broader categories of problems to which machine learning algorithms are frequently applied.

How Do Businesses Use Machine Learning?

Recommendations

The complex recommendation systems that greatly support e-commerce sites like Amazon or online libraries like Netflix are made possible largely by machine learning. By choosing a suitable categorization algorithm and passing available customer data through it, computers can utilize the similarities and differences between user data profiles to group them together. Once grouped, recommendations can be pulled on the assumption that similar users will prefer similar products. Implementations of such systems can encourage desired behaviors by providing end-users with useful information with minimal effort and a high degree of accuracy.

Automatic Categorization

Beyond photos of animals, computers can often be taught to categorize other kinds of data based on its intrinsic features. By analyzing a subset of pre-labeled examples, computers can determine the qualities that differentiate one category of data from another. A common example is how computers can analyze word usage in spam vs. normal emails to determine which to send to your inbox and which to relegate elsewhere. While these cases can also require human review where the algorithm's accuracy is suboptimal, they generally reduce the workload of reviewers and can dramatically increase responsiveness and overall coverage, especially of certain categorizations which require specific action.

While machine learning is regularly being utilized in unique and novel ways, there are definitely broader categories of problems to which machine learning algorithms are frequently applied.

Process Refinement

Similar to how machine learning algorithms can use their own success and failures to self-improve, they can also analyze the success and failures of other processes and provide similar suggestions for process modification. Rather than iterating upon weights, these algorithms reveal commonalities between success and failure cases which provide an objective view of what is effective and what is not.

A common example are search engines, which often use user activity following a search to determine whether or not their query processor produced an accurate result. In this case, the algorithm will highlight the features of the query processing software that work, which features do not, and an overview of the specific cases in which they fail. Such data will eventually aid analysts in finding the optimal combination of known traits. For problems and processes that are especially complex, machine learning can help remove some of the baseline complexity and dramatically narrow the scope of a problem to a few key features.

Predictions

Certain algorithms can use past data to model the relationship between a number of features called parameters and leverage it to predict output values for an entirely unique combination of parameters. Think of it like deriving an equation, but the variable relationships can be anywhere from simple and linear to complex and multi-dimensional. Such algorithms fall within the scope of regression analysis and are utilized in a number of industries to provide fast predictions with a high degree of accuracy. Like most algorithms, it will do poorly in edge cases where the result does not logically follow from the data trends, but in this case its speed often outweighs its imperfections. Such predictions shine when applied to tasks like stock trading and fraud detection where speed is often far more important than the losses that can be incurred by an inaccurate prediction.

Pattern Discovery

Most of our previous examples focused on fitting data into predefined categories, and such strategies are called supervised learning models. While they can certainly be used to solve a broad variety of problems, they often fall short in cases where data is available but groupings for such data have yet to be determined. Here, unsupervised learning models step in to fill the gaps. These algorithms utilize clustering or association techniques to group data based on some or all of their features, deriving categories driven by data rather than higher-level concepts. Through this, such analyses can reveal new or surprising information about the data being processed. Businesses often utilize this strength of unsupervised learning to refine their understanding of their customer base, allowing them to make business decisions better attuned to the realities of their target market.

Overall

Even with such a diverse array of applications, machine learning is still very much in its early stages of growth, with some reports expecting its usage to increase massively within the next two years alone. With such high rates of expected growth, early adoption of machine learning could provide businesses with a technology advantage key to effectively utilizing an ever-increasing amount of available information. In a climate where making data-driven decisions is becoming an increasingly key component in retaining a competitive edge, adopting the technology that will enable such fast, accurate judgement on large and complex datasets becomes all the more imperative.

About Levvel

Levvel helps clients transform their business with strategic consulting and technical execution services. We work with your IT organization, product groups, and innovation teams to design and deliver on your technical priorities.

Our App Dev team is made up of technology agnostic enthusiasts with a wide array of knowledge in popular modern and legacy languages such as Java, .NET, Ruby on Rails, Node.js & JavaScript, Python, and PHP.

We firmly believe that mentoring can be integrated with delivery. Our main focus is on saving our partners as much as possible on the lifetime-total-cost of ownership and maintainability of their systems. For more information, contact us at hello@levvel.io.

