The Ultimate Business Owners' Guide to Machine Learning

"Nowadays barely a single geek talk, tech conference agenda or software product release can do without mentioning Machine Learning. Like "data is the new oil", Machine Learning is like the chemical industry that utilizes the raw material to change the way we live. But the excitement surrounding the topic is evolving disproportionately faster than the level of understanding, which we see as a problem potentially harmful for all levels: business, consumers, and scientists.

In this guide, we made an attempt to familiarize decision-makers, regardless of their background and industry, with theoretical and practical basics of Machine Learning. We hope that the knowledge one can acquire by reading this guide will accelerate technological change at companies and organizations and, at the same time, create an unbiased picture of the technology itself."

Yuri Svirid, PhD

CEO, Silk Data

Foreword

Machine Learning (ML) is a buzzword nowadays that is about to reshape every sphere of our life soon. Well, it already does, and the pace of adoption is only going to accelerate as cost of production and marketing rises and companies find it more challenging than ever to stay on top of competition.

Thus, Netflix save about $1 billion annually by using machine learning to make personalized recommendations while GE claims implementing machine learning software saved them $80 million. Whether in medicine, insurance, transportation or agriculture, every progressive business now realizes the importance of automation and smartification of their processes.

However, there's so much hype around the topic these days, that it is complicated for business owners, investors, and industry professionals to get a clear undistorted picture of the technology and its capabilities.

Some people approach machine learning as a magic wand that can solve their problems with little or no effort, other believe its practical value is overestimated. With this guide, we try to create a reliable and unbiased image of machine learning and provide people of different backgrounds and from various industries with comprehensible knowledge on the subject.

You will find this guide useful if you're:

• A business owner or CEO It is essential that business leaders and entrepreneurs not only keep up with their industry trends but can leverage advanced methods and technology to move beyond the traditional practices and stay on top of the competition. Machine Learning is indeed the great source of getting a competitive advantage, which makes the knowledge comprised in this eBook indispensable for a business owner.

• A product leader Machine Learning is extremely powerful when it comes to providing enhanced functionality and personalized user experience for digital products. If you are a product manager, you might have already considered exploiting machine learning capabilities – and this guide will help you explore a more practical side of implementing the technology.

• A change leader There are no managers who wouldn't want their work to be done better, faster, and more efficiently. Machine Learning can automate routine tasks at scale and at a much higher speed than legacy software. If you'd like to look a little deeper behind the articles on Digital Trends, this guide will give you a perfect introduction into the basics of machine learning.

This is the short version of our Machine Learning guide

We've removed a significant part of theory in this version of Machine Learning guide, to focus on what's important for business decision-making process. If you are looking for more solid knowledge and specific examples, please get our executive guide here.

What is Machine Learning?

Machine Learning is a quite multidimensional concept that uses approaches from statistics, computer science, and data processing.

In simple terms, **Machine Learning**-based product is a computer program (a combination of algorithms and data) designed to get correct answers, generate accurate predictions and make its own sensible decisions without being explicitly coded. Unlike classical programs, it mainly relies on **data**, not algorithms, to progressively improve itself upon a given task, learning from experience. Let's take an example of a predictive text function. Instead of giving the program explicit instructions (which could comprise millions of alternatives for a single step) you can show the examples of thousands of previously sent messages and the machine will use this data to build reasonable predictions of the words and phrases you're going to use next.

In far 1952, IBM's Arthur Samuel created one of the first machine learning programs that played checkers at the human level and got better with every new game (according to Wikipedia, it had been running on IBM's first commercial computer!)[1]. 54 years later, Google's Alpha Go beat a professional human player at Go, arguably the hardest board game ever created. One can be either excited or scared of the speed at which the progress is being made here. But despite all concerns around whether machines will be able to outsmart humans anytime soon, the primary goal of machine learning is to make people's life better and their work easier.

Over the last 70 years, Machine Learning experienced several waves of hype, with different buzzwords being used to promote new achievements. You might have heard things like Data Mining, Deep Learning, Synthetic Intelligence, Cognitive Computing, Affective Computing, DeepTech, Data Science and so forth. Despite some technical difference, all above-mentioned terms can be explained in a single way: an approach to making the computer solve tasks where exact rules cannot be easily set. In the rest of this book, we will normally refer to Machine Learning (ML) and Artificial Intelligence (AI), without paying much attention to difference between the terms.

---

[1] See https://en.wikipedia.org/wiki/Arthur_Samuel

In simple words, the typical application case for ML can be formulated as follows:

Machine Learning Rule of Thumb

> "If a typical person can do a mental task with less than one second of thought, we can probably automate it using AI either now or in the near future."
>
> Andrew Ng, Stanford University

The Typical Tasks Solved by Machine Learning

The intelligent systems built on machine learning algorithms have the capability to learn from experience or historical data which make them a powerful tool for solving a wide range of tasks. Currently, ML has been used in multiple areas, from social media (photo tagging) to genetics (identifying disease genes in a human DNA).

Here are some common applications. They demonstrate that the opportunities of practical use of ML are endless and not limited to any specific area.

Input	Response	Application
Picture	Are there cats?	Photo tagging
Loan application	Will they repay the loan?	Loan approvals
Ad plus user information	Will user click on ad?	Targeted online ads
Audio clip	Transcript of audio clip	Speech recognition
English sentence	Chinese language	Language translation
Digital text	Text summary	Text summarization
Users' transaction history	Customers grouped into segments	Customer segmentation
A user's saved tracks	Other tracks he is likely to appreciate	Recommendation systems
Utility bill	Extracted personal information	Named entity extraction
Text-based document	Similar documents in the collection	Similarity search

Introducing AI on the company level can help cut down on costs and increase performance manifold in the long run, first and foremost through automating repetitive tasks and reducing human errors. While technology giants like Google, Amazon, Facebook, Microsoft, Apple and other AI frontrunners are increasing their efforts in integrating Machine Learning, a wide adoption of ML and AI is still at its very roots.

But let’s take some cases (there are millions of them, but we take just a few as examples) where implementing AI and Machine Learning yielded significant results for companies.

AI and Machine Learning Success Stories

1. Using AI to reduce fashion chain customer friction

In an effort to better stock individual stores with what they customers wanted, H&M was using big data and Artificial Intelligence (AI) to analyze returns, receipts and loyalty card data and tailor the inventory offering for each store. In order to get customized merchandise to each store and to respond to consumers' demands for a hassle-free shopping experience, H&M invested in automated warehouses that ultimately resulted in next-day delivery for 90% of the European market over 12 months.

3. Online retail: Boosting conversion rates by 32% with AI-based recommendation engine

A large online retailer wanted to improve conversion rates and increase average order value by offering more personalized product listings on its website. Initially, historic transactional data across all markets were analyzed, then the AI-powered tool was built to discover cross-sell patterns (which products customers frequently bought together). Using these insights and user behavior data the software was able to provide custom recommendations for website visitors, displaying the products that a consumer was most likely to purchase. The online retailer managed to achieve a **32% conversion rate uplift**, as well as a **23% uplift in revenue**.

4. Applying AI in legal practices

AI has already established itself as a very useful technology in legal practices. Last year LegalTech companies hit $1 bln in combined investments mainly thanks to AI and machine learning being used to speed up and improve processes. Thus, a Canadian startup asserts that its software can perform a more accurate contract due diligence than lawyers by searching, highlighting, and extracting relevant content for analysis. A staff that need to perform multiple reviews of the content can search for the extracted information with links to the original source using the software. It is claimed the system can complete the task up to 40 percent faster when used for a new set of documents, and up to 90 percent faster when trained on real-time data.

ML vs Classical Software

Machine Learning as Data Analysis

The question many technical folks ask: How is Machine Learning different from a legacy software? In classical software engineering, there is always an expert – normally a person – that sets precise rules when writing a code. The software uses data to generate answers a user wants to obtain.

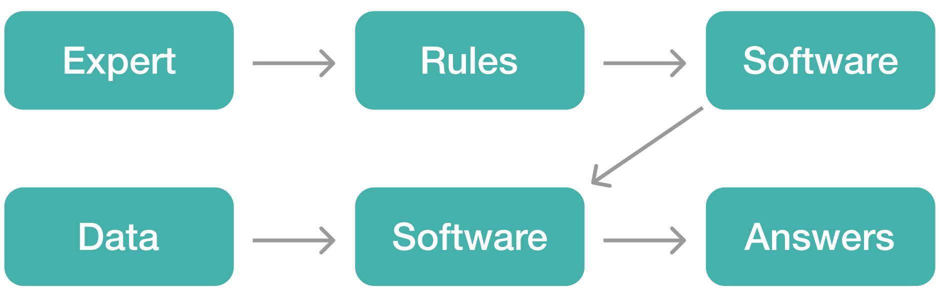

In order to alter the results (as environment is changed and new patterns are discovered in the past data) you would need to re-write the rules. More importantly, in many cases formulating the rules can be prohibitively complex, the most common example is how to write down all rules necessary to tell dog from a cat.

Machine Learning algorithms work in a different way. They can extract the rules without being explicitly programmed by human.

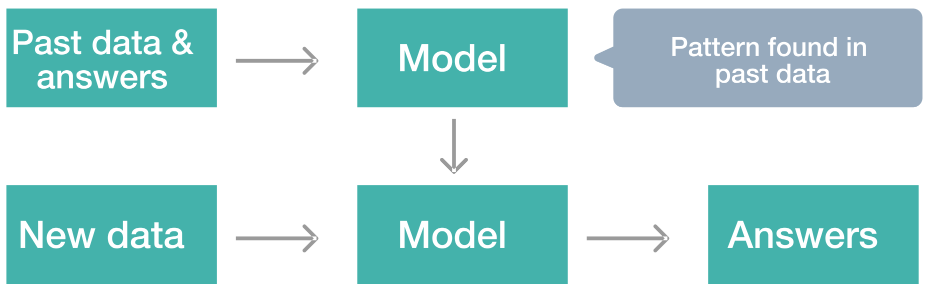

All patterns discovered in the past would draw new conclusions to form new data sets. This makes possible to process a huge amount of information without regular intervention by experts. Instead, experts can only analyze a small number of samples and make ad-hoc adjustments.

Types of Machine Learning

ML algorithms can be commonly divided into several categories.

- *Supervised learning* (learning with a teacher)

Requires past data or labeled dataset to make predictions based on them. A supervised learning algorithm takes a known set of input dataset and its known responses to the data (output) to learn the regression/classification model. A learning algorithm then trains a model to generate a prediction for the response to new data or the test dataset. Supervised learning is used in visual object recognition, email filtering, recommendation systems, etc.

- *Unsupervised learning* (learning with no teacher)

Produces output from unlabeled data (for example, shopping data). A Machine Learning system tries to unearth a certain pattern in the data, the results to be generated are unknown). The most known example is data clustering (used in every case where you need to classify objects based on specific patterns).

- *Semi-supervised learning* (a mix of both)

Used for expensive data (when getting labeled quality data can be costly). For instance, where you need to have an expert.

The principle is that data with known labels is mixed with unlabeled data obtained from a similar source. The result is based on combining clustering results with known labels to predict the labels.

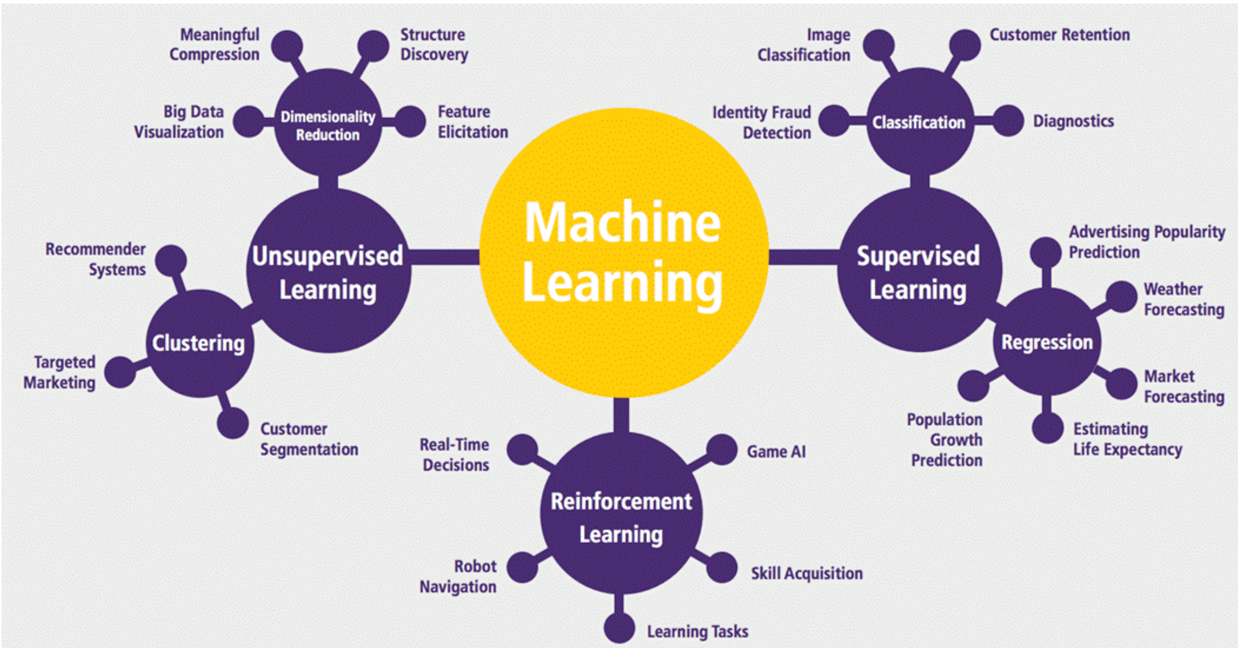

Picture credit: [guru99]

- *Manifold learning* (dimensionality reduction)

In case of the high-dimensional data, there are many input parameters (features), and we need to visualize them in some way. As human brain can hardly perceive more than three dimensions, there are some specific ways that allow reducing the data to two or three dimensions keeping the same structure in order to visualize the data. This provides for making better decisions and more accurate predictions.

- *Reinforcement learning* (learning ‘on-the-fly’)

A Machine Learning model learns to respond to changes – obstacles, changing traffic conditions and similar – and learns to take decisions “on the fly”. The most known examples of reinforcement learning are self-driving cars and high-quality bots for computer games.

Deep Learning

What is Deep Learning?

Deep learning is a branch of machine learning based on a set of algorithms that attempt to model high level abstractions in data.

Deep Learning (DL) and Neural Network (NN) are driving some of the most significant inventions in last decade. Their outstanding ability to learn complex dependencies from data makes them the enabling technology for many machine learning tasks, such as self-driving cars, image recognition software, automatic text recognition and translation, recommender systems and so forth. Apparently, being a powerful algorithm, it is highly adaptive to various data types as well.

Neural Network, also referred to as Artificial Neural Network, is named after its artificial representation (or, rather, mimicking) of functioning of a human's *nervous system*, especially the visual cortex. Nervous System comprises of millions of nerve cells or neurons, each having the following structure:

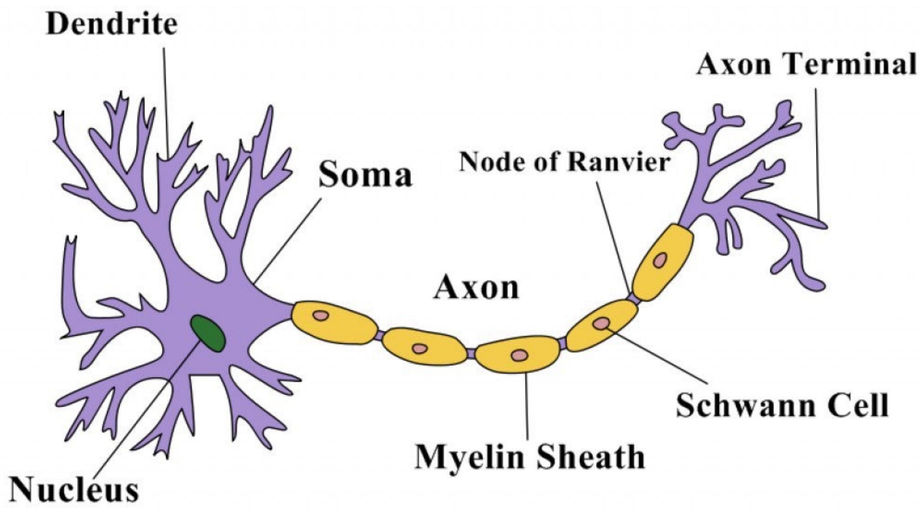

Source: Quasar Jarosz CC BY SA 3.0, via Wikimedia Commons

In simple words, each neuron receives input information in form of electrical impulse from other neurons, does some processing and then transmits it to other cells. Though a single neuron performs only very basic operation, collective actions of all neurons in the nervous systems results in highly complex behavior and cognitive abilities of animals, including us humans.

Artificial Neural Network works in a very similar manner and can be visually represented as a set of layers of neurons. At each layer, the individual neuron is taking in some input information, weighting it accordingly and then passing on a modified version of this input to the next layer. Thus, the output of each layer forms the input of next one.

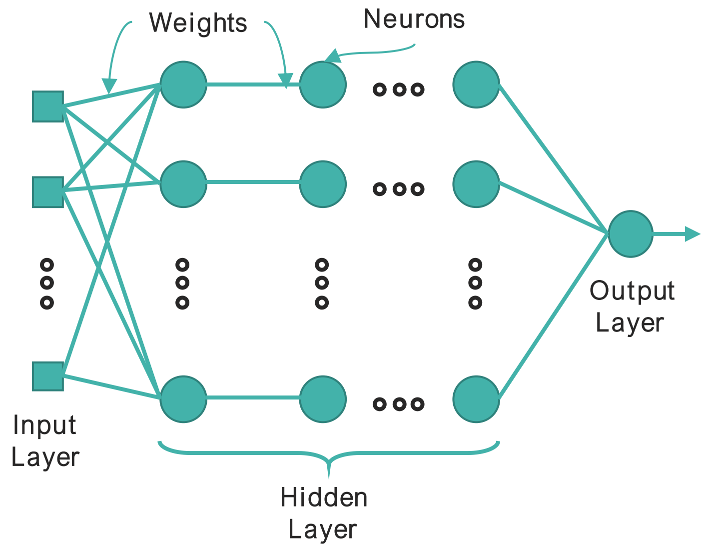

Usually, more complex tasks requiring a bigger number of elements in a neural network. In a **deep** neural network, there are many layers between the input and the output, which allows the algorithm to use multiple processing steps each implying numerous linear and non-linear transformations.

Deep Learning Example

To explain how Deep Learning works let us take an example of CNN – the Convolutional Neural Network – which is currently the standard approach in Deep Learning for analyzing image data.

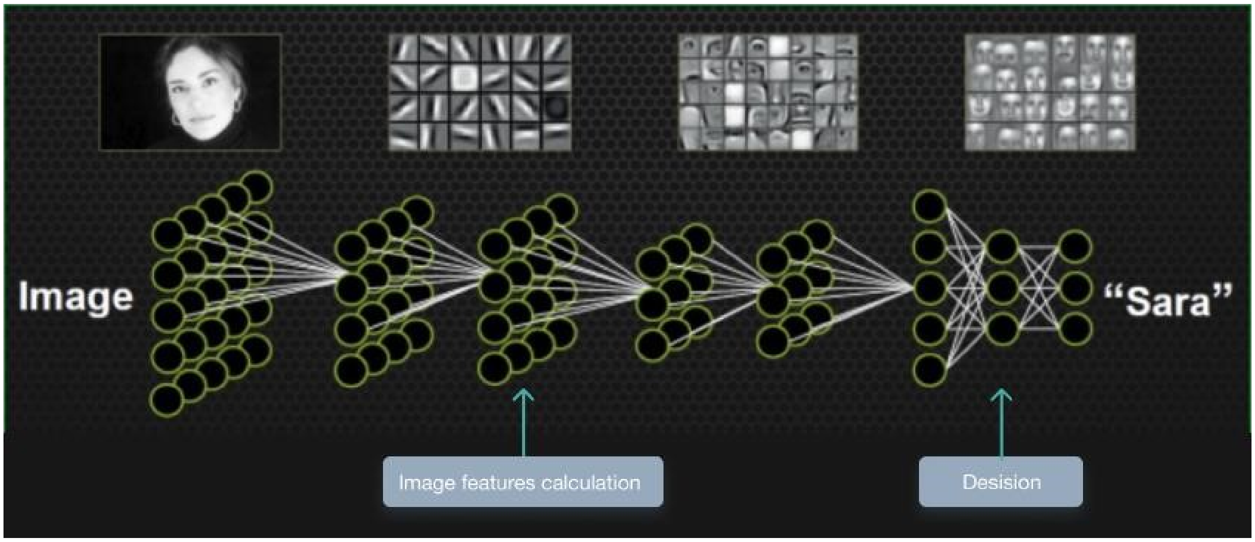

Source: *Jason's Machine Learning 101* presentation, Dec 2017

If we submit an image of a human face as an input data, the first layers will detect some high-level features like outlines and edges. The following levels will then face-shaping object like eyes, nose, and lips be recognized. In the end, the neural network will "assemble" the attributes of an individual which enables the system to figure out who is presented on the picture.

Risks and Challenges

Machine Learning is a very powerful tool that doubtlessly holds an enormous potential for organizations, but it also poses a substantial risk, especially if not treated properly.

Types of Risks

Deloitte[2] splits machine learning algorithmic risks into 3 main categories, which are the input data, algorithm design, and output decisions.

The first group of risks relates to 2 main problems: insufficient, low quality, outdated or irrelevant data sets and biases, that can cause a model to be widely inaccurate.

The first risk area is associated with a volume and quality of a data set. Do you know the origin of the data? Wasn't it faked? Does it provide enough information and variability to be used in the model? Answering these questions is necessary in the preparatory step.

---

[2] https://www2.deloitte.com/us/en/pages/risk/articles/algorithmic-machine-learning-risk-management.html

There might also be different compliance issues depending on your jurisdiction and where does your data come from.

On the other hand, all people have biases, and since machine learning models are made by people (and often using the data produced by humans), we cannot avoid them being biased. Many business cases will require data to be adjusted for special circumstances where new parameters or conditions come into play.

Algorithm design is vulnerable to risks as well. Let alone coding errors, the logic and assumptions we pledge while training a model are often false and training results are disconnected from the reality. That leads to a model being useless in the real world.

Even when the data is flawless and algorithms were chosen and tweaked perfectly, there is a risk the **results of the machine learning process** will be misinterpreted or will disregard its underlying assumptions. Such situations are very common in business. As machine learning models are built to help in decision-making process, it is essential that an understanding of its outcome is fully aligned across stakeholders.

Mitigating Machine Learning Risks

Anticipating your question: given that there are so many risks how can we make sure the machine learning process will go smoothly?

The first thing that all organizations leveraging this technology need to work out is a clear formulation of **goals** they want to achieve through machine learning projects. Having internal **strategy** and management scheme is important, as well as to have clearly outlined principles and policies and to assign accountable employees.

Which is **indispensable**, is to employ a team of experienced data scientists and machine learning engineers (or contracting a trusted service provider) alongside industry professionals. As we previously mentioned, a teamwork of data engineering and business area experts is required to mitigate multiple data and algorithm risks and to help understand how results could differ from expectations.

For some projects (especially on the enterprise level) a group of independent data experts will be needed to provide an unbiased evaluation of the model on its fairness, compliance with existing privacy rules, and explainability. Making models explainable and interpretable are actually one of the biggest challenges that machine learning engineers face nowadays (especially for most Deep Learning that is necessary for most demanding tasks). Let's have a closer look at these and other natural limitations of machine learning.

Limitations of Machine Learning

Machine learning has an enormous potential for solving business tasks, improving processes and products. But one should not ignore its limitations when deciding on running a project or implementing the technology.

1) Lack of explainability

But just like personal computers were not built to interpret their decisions to general users, Artificial Intelligence was not designed to explain its predictions. However, as Machine Learning is becoming more pervasive in all spheres of live, human need a clear understanding of its decision-making process. Car manufacturers need to have a full control over AI that works is in the core of their autopilot systems. Investment brokers that utilize machine learning must be able to explain the grounds for their decisions to a client.

Yet in many cases data scientists still don't have an explanation to how Machine Learning models work, even sophisticated non-linear models tend to become interpretable thanks to several recently developed techniques[3], and this field is expected to progress quickly.

2) It requires big data sets

As we already mentioned, Machine Learning models require vast amounts of data to get trained – moreover, these data must be clean, unbiased, and variable enough and to provide quality results. It may seem that obtaining data is not a problem in the world where 2.5 quintillion bytes of information is created each day. In fact, collecting enough data and perfectly labeling it requires time and money resources. Another challenge that organizations face at this stage is privacy regulations that seriously limit collection and utilization of personal data in certain countries or regions.

3) Talent scarcity

When an extensive dataset is in place, you must inspect and clean the data, select features, then choose or build a custom model,

[3] https://medium.com/@Zelros/a-brief-history-of-machine-learning-models-explainability-f1c3301be9dc

apply a set of machine learning algorithms, pick up and tune parameters, constantly monitor performance and make necessary adjustments. However, not many companies possess such resources or can easily hire them. Tencent estimated[4] in 2017 that there were just 300,000 AI engineers worldwide. That is not to mention high-level and experienced specialists who can manage complex projects and conduct serious artificial intelligence research. There are fewer than 10,000 of those in the world, and the number is not expected to grow exponentially. All it makes machine learning adoption increasingly difficult for companies, let alone startups that are more actively turning their eyes towards external provides and consultants.

Although an effective model for small and middle-sized companies, finding a reliable partner with deep relevant expertise and an appropriate level of compliance emerges as another challenge.

4) It takes time to get desirable results

One of the reasons why so many organizations fail in machine learning projects is because they expect immediate results. A lot of

---

[4] https://www.tisi.org/Public/Uploads/file/20171201/20171201151555_24517.pdf

business owners and decision makers tend to approach Machine Learning projects the same way they do classic software development, which is wrong.

In software development, goals are well-defined, algorithms are transparent, and outcomes are binary, while in machine learning algorithms are designed to improve accuracy (or success rate) over time and outcomes are often subjective. Time and patience are indispensable prerequisites for maximizing an impact of machine learning within an organization. As more and more companies utilize neural networks and ML algorithms to work on running data that they users generate every day, the process may well become infinite requiring engineers to continuously tweak and adapt parameters to new sets of data.

5) Machine learning is not good at understanding context

Machine Learning algorithms are almost brilliant in recognizing objects and classifying images, but it has some troubles with understanding context.

You can say without hesitation whether the image is rotated 180 degrees or not, but it isn’t that obvious for AI. If the model was trained to locate people on images it would do just that. It can learn to recognize when a picture is upside-down given it has seen a large dataset of both normal and inverted images.

Until recently, that was a “mission impossible” for computers but latest achievements give hopes that the solution will be found soon. Several modern approaches, like vectorization of words in natural language processing, are yielding promising results in solving the ‘context understanding’ problem. A breakthrough in this field can make a true revolution in certain areas like legal practice and knowledge management yet it remains a limitation of machine learning.

6) ML algorithms can be fooled

Even a well-functioning mathematical model — one that relies on good data — can still be tricked, if one knows how it works. Like on the given example, overlaying a minor noise on the image can lead to its miscategorization.

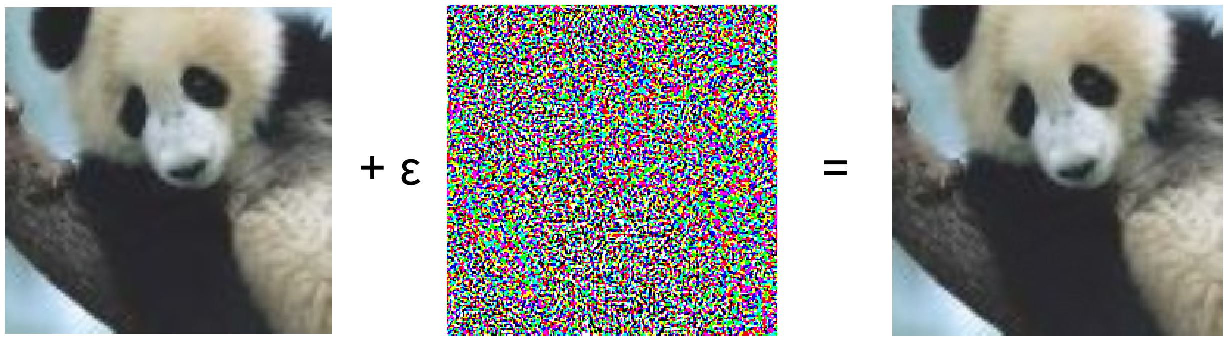

"panda" 57.7% confidence

Random noise

"gibbon" 99.3% confidence

Source: Explaining and Harnessing Adversarial Examples

This also keeps many companies, organizations, and governments away from a large-scale adoption of machine learning, because as long as it can be tricked so easily, the risks would always prevail over benefits.

Conclusion

We are still far away from having an impeccable Machine Learning which is transparent and interpretable yet robust and secure. However, by having business owners and decision makers at every level educated on its underlying risks and limitations, we are making the task substantially easier and bringing us closer to solving these challenges.

Key steps to start with ML

After reading the news on AI progress, one usually expresses a natural concern that using ML implies specialized software and expensive hardware. While it is true to conduct groundbreaking research, practical application of ML techniques can be normally started virtually for free!

First, recent progress in ML techniques have been heavily influenced by the open source movement. Many important software packages and libraries are made free to both private and commercial use.

Quite frequently, starting a ML project may encounter the chicken-or-egg problem. To show the utility of ML on a company level, it is required to have certain amounts of data available but, as was previously mentioned, collection (and annotation) of data is often expensive and time-consuming process. Sometimes, to convince stakeholders you need to demonstrate initial result with very limited or even no data. Fortunately, the are multiple open data sets that can be used for model in tasks like image classification or advanced text processing, which makes it possible to build a useful prototype in a very short term with limited resources available.

Finally, there have always been questions: it necessary to have an advanced powerful computing cluster to start a ML project? Typically, the answer is **not**! In many cases you can develop a proof-of-concept by using typical office (or gaming) PC, with an option to use cheap (or even free) cloud computing solutions.

That said, the most important resource in starting a machine learning project people, either an internal team of AI and ML experts or a trustful solution partner.

Thank you for making it till the end!

We've thrown a lot on you but hope it was time worth investing. Now it's your turn to put these ideas into practice. It must be way easier now to see which parts of your product offering or business processes can be improved with Machine Learning and consciously evaluate all benefits and risks connected to implementing the technology. Our team of ML experts is ready to provide you a comprehensive consultation and assist in planning and realization of your projects. Send us a message to hello@silkdata.ai with the promo-code stated below, and we'll grant you 3 hours of expert level Machine Learning consultancy for free.

Your personal code to claim 3 hours of free Machine Learning consultancy

SILKMLGUIDE3

Silk Data is a European Machine Learning and IT consulting company founded by PhD titled experts with an average of 25 years of academic and business experience. Our mission is to help companies solve their problems by using full potential of Machine Learning, Artificial Intelligence, and advanced digital technology. For more information, please visit www.silkdata.ai.