Re-imagining Customer Engagement with AI:

An Introduction to AI

Written By:

Allan Frank

Capgemini

CONTENT

01 ARTIFICIAL INTELLIGENCE: THE NEXT 'BIG THING' IN DATA-DRIVEN CUSTOMER ENGAGEMENT

It used to be the stuff of science fiction films – an aspirational otherworldly concept that captured our imaginations. Today it is the foundation of competitive customer engagement.

Artificial intelligence (AI) – the secret sauce of personalized experiences from Siri and Alexa to news feeds and chatbots -- is poised to impact business and our economy as dramatically as the Industrial Revolution, the Information Age, the Internet, and the advent of mobile devices.

Most enterprises are familiar with the concept of AI – usually in the context of machine learning. But universally, organizations don't understand the the full scope of AI and how it can be practically applied in their enterprise.

We hear the same questions repeatedly:

- What is AI?
- Is it the same as cognitive technology?
- Why is AI suddenly relevant now?
- What does AI mean for my industry and line of business?
- What are the relevant use cases?
- How can I incubate AI in my organization?
- How does it fit into my overall digital strategy? What are the new skills, technologies that we need to develop?
- It must be really expensive! Can I afford it?

Every business that has successfully integrated AI into their processes started out with the same questions. We developed this paper to provide some clarity around AI technologies and applications.

02 A NEW IDEA THAT'S AS OLD AS THE BABY BOOMERS

Artificial Intelligence is not a single technology; rather, it's an umbrella term for technologies inspired by human biological systems. AI constitutes a variety of software, algorithm-driven approaches with data, that together simulate any number of human cognitive functions.

To the lay person this is "dare to dream" stuff – but the truth is that AI has been around since the 1950s. In the early years, there was a somewhat naïve view regarding what it would take to build an "intelligent" machine, so progress was slow. Consider the tools computer scientists had to work with at that time: Programming required very detailed instructions and rules to get the computer to do anything useful. Computers filled entire rooms, and the Internet was just a thought in the back of someone's mind.

Fast forward to today: Over 1 billion people and millions of computers and sensors are interconnected across the world. Computers have shrunk – and many are not much larger than a grain of sand. A convergence of forces over the last 5-10 years has enabled a virtual explosion of possibilities for AI that is quickly moving into every facet of how we work, live, and play. The most influential drivers include:

- Availability of data
- Cheaper and faster hardware
- Better algorithms

03 AVAILABILITY OF DATA

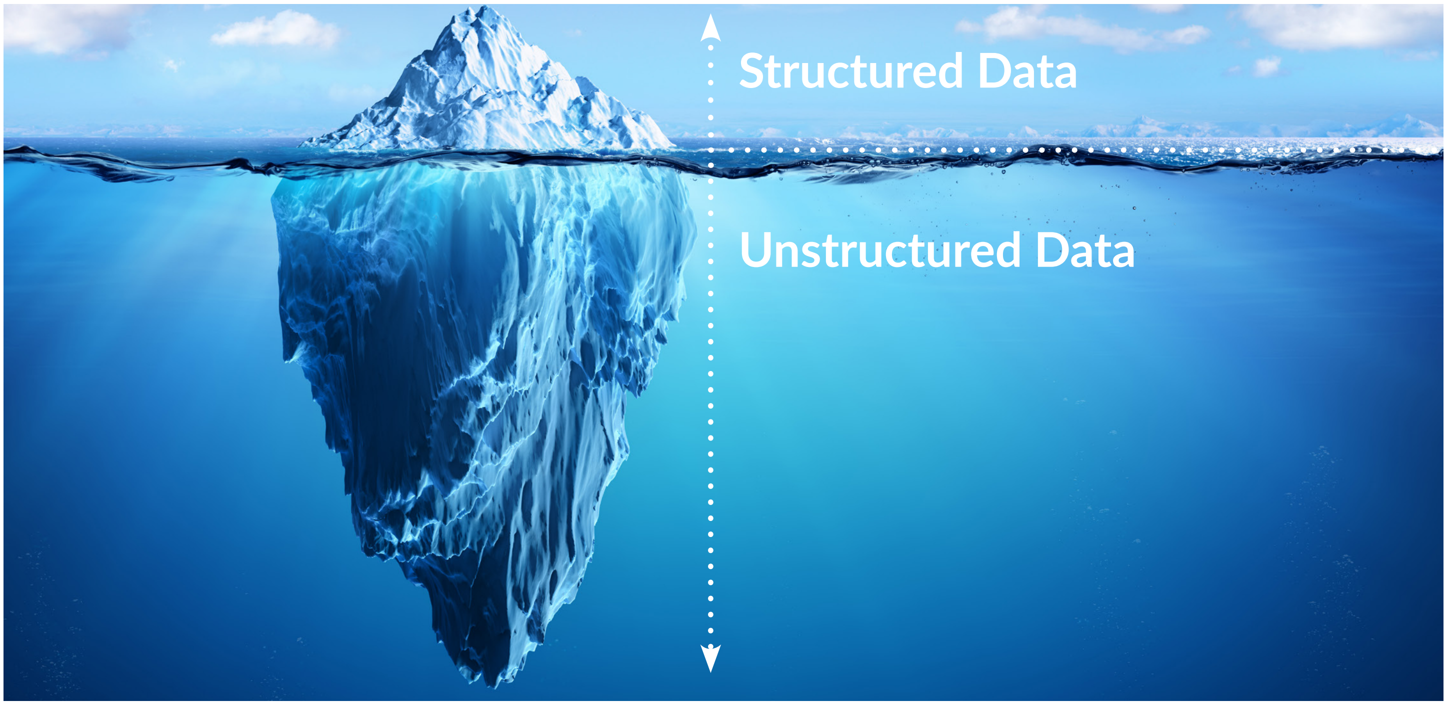

Quality training data (text, image, audio, user activity, or knowledge graph) is the new competitive advantage in AI.

As its name implies, artificial intelligence is about replicating what the human brain knows. The learning is iterative, and requires processing thousands or millions of examples of something (like photos, music, words in a dictionary, etc.) to become competent in a task. The level of data available significantly impacts the speed of learning, and the ultimate competence AI can attain. For example, Google's success rate in delivering more precise results in website and image searches is reliant on the unprecedented volumes of data fed into Google's AI algorithms.

The Big Data generated as a byproduct of our increasingly interconnected existence is a vital resource in AI development. Businesses are not only able to collect, analyze, and filter specific information at an amazing speed – but also recognize patterns that drive decision-making policies based on this analysis. Thus, AI provides the missing link – intelligence - that helps businesses further their goals.

04 BETTER ALGORITHMS

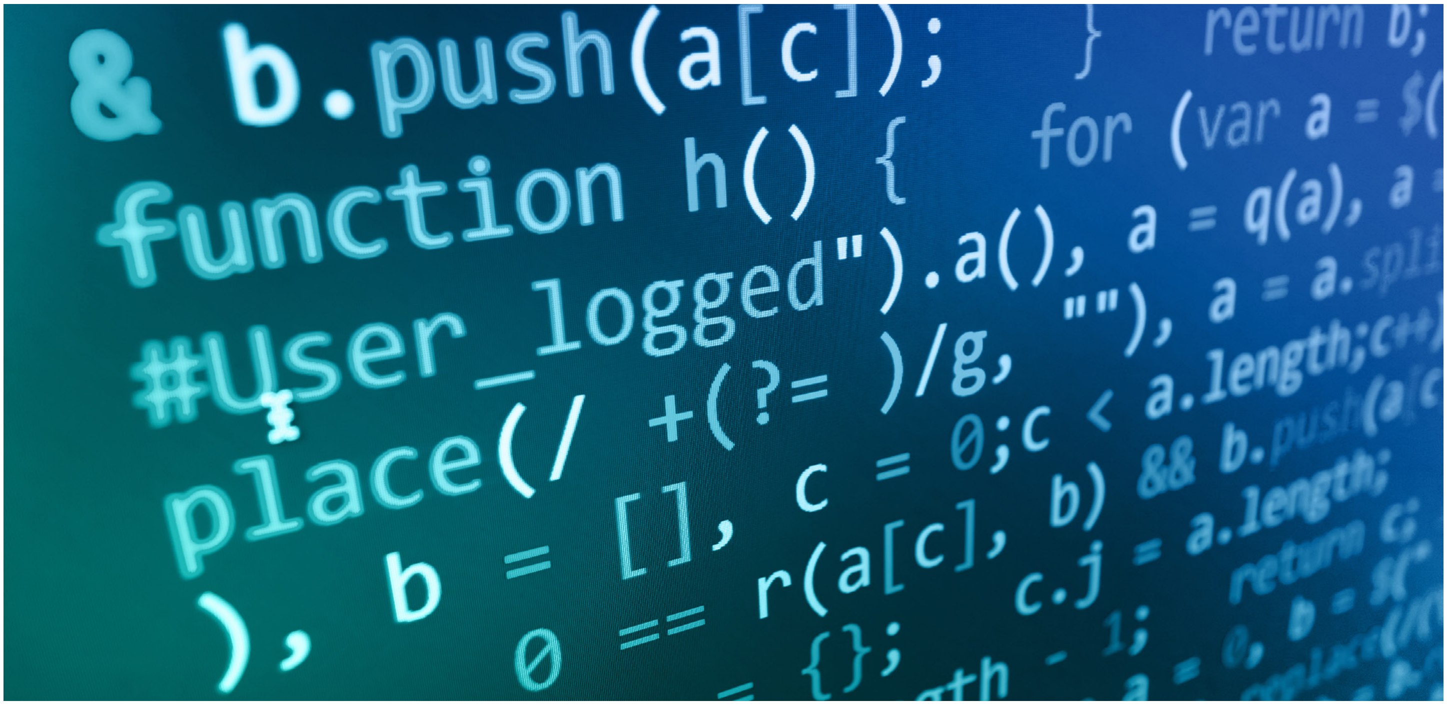

New types of algorithms lie at the heart of the new AI wave and it's safe to say that without these, the data and computing power would amount to nothing. It might come as a surprise that it is all based on 1950s-era technology known as the Artificial Neural Network (ANN) - an attempt to create a computer model of the network of nerve cells in a human brain. Loosely speaking, a neural network is an interconnected web of artificial neurons that either fire or not based on the input to the neuron.

A key part of building these neural networks is to train them to do the correct thing when they see data. The study of these algorithms has now spawned its own sub-field of AI known as deep learning, a reference to the number of neuron layers in the neural networks. We'll talk more about deep learning in the next section.

05 CHEAPER AND FASTER HARDWARE

AI applications require more computing capability than conventional computing architectures are able to provide.

Traditional computer processors can only process information linearly – whereas many aspects of human intelligence are of a more parallel nature. Serial computing - which is what CPUs (Central Processing Units) are built for - cannot tackle these tasks with any efficiency. CPUs are commonly referred to as “the brain” of the computer and are responsible for executing a sequence of stored instructions.

Architecturally, the CPU is composed of just few cores with lots of cache memory that can handle a few software threads at a time. In contrast, a GPU (Graphics Processing Unit) is composed of hundreds of cores that can handle thousands of threads simultaneously. This enables GPUs to process neural networks 25 to 40 times faster than serial CPUs, and accelerate some software by 100x over a CPU alone, with increased power, and cost-efficiencies.

The TPU (Tensor Processing Unit) takes this to the next level. A TPU is a custom-built chip developed specifically for machine learning and tailored for TensorFlow, Google’s open-source machine learning framework. On AI workloads that utilize neural networks, the TPU is 15 to 30 times faster than contemporary GPUs and CPUs – and is much more energy efficient. For example, in Google Photos, an individual TPU can process over 100 million images a day. To put this in perspective, Google claims that this is roughly equivalent to fast-forwarding technology about 7 years into the future (3 generations of Moore’s Law).

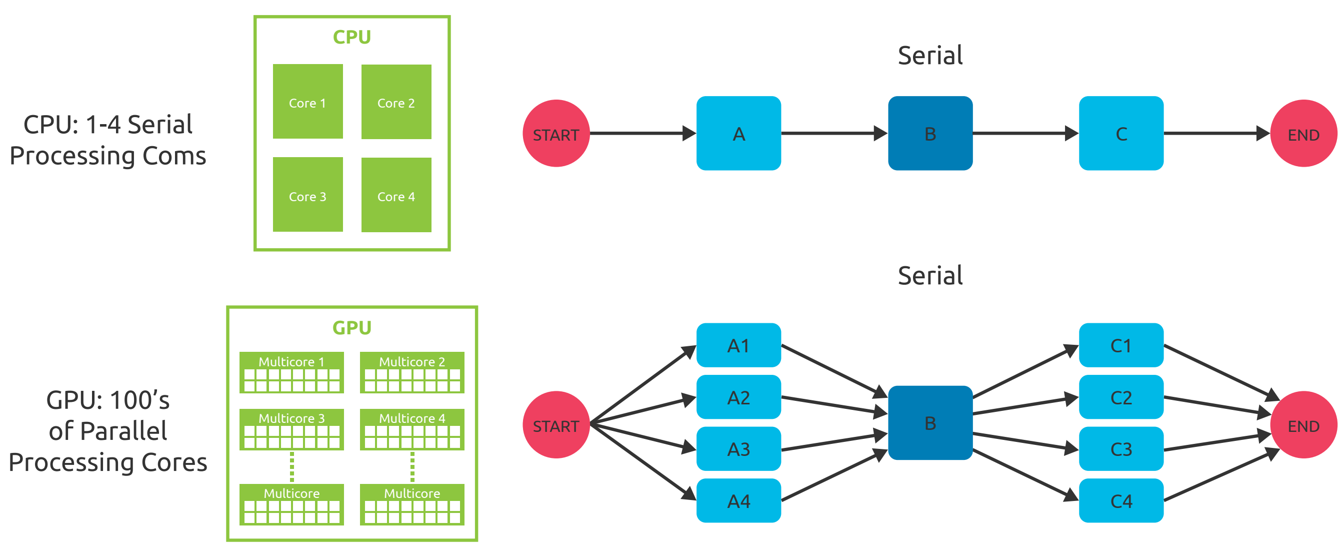

06 AI: ONE CONCEPT ENGENDERED IN MULTIPLE WAYS

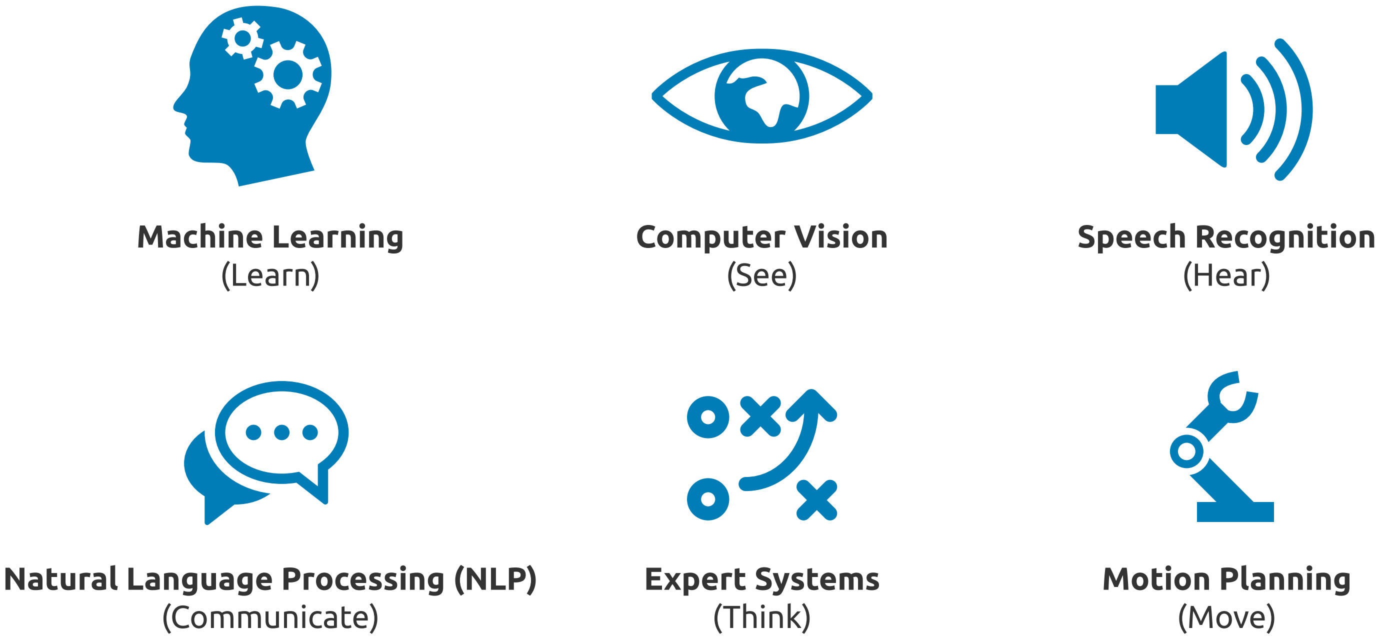

Artificial Intelligence (AI) is the overarching concept which broadly refers to machines exhibiting human intelligence. It is an umbrella term for technologies inspired by human biological systems. Under this umbrella, there are variety of software and algorithm-driven approaches (combined with data) to simulate human cognitive functions as shown in the below figure.

While I am not going to dive into each of these approaches, I will examine Machine Learning, Computer Vision, and NLP as some of the more pervasive approaches of AI being applied in business today.

AI is an information system that is inspired by a biological system designed to give computers the human-like abilities of hearing, seeing, reasoning, and learning. **- Tractica**

07 MACHINE LEARNING

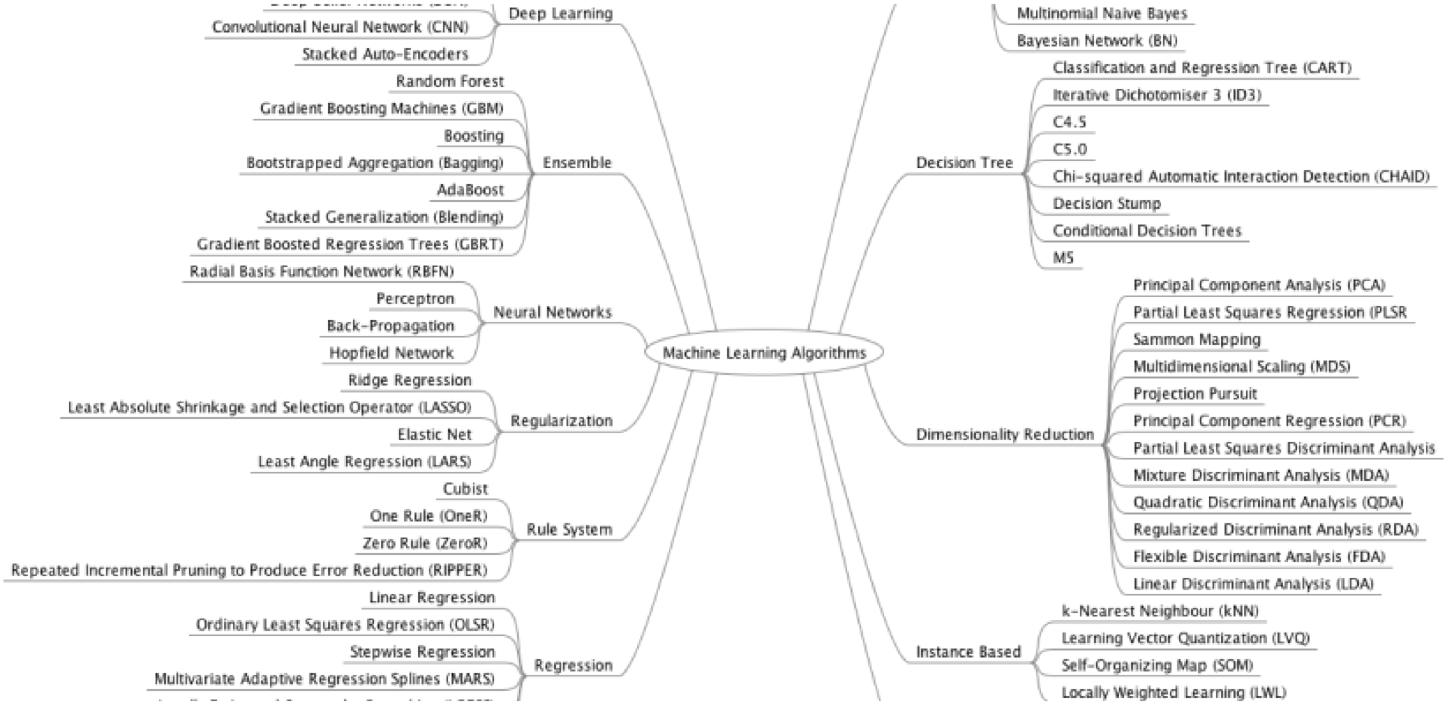

Source: https://jixta.wordpress.com/2015/07/17/machine-learning-algorithms-mindmap/

Machine learning is one of the most widely recognized and successful applications of AI.

Until recently, a computer needed to be fed detailed procedural instructions in order to solve a problem. For this to happen, the programmer needed to fully understand not only the problem, but also the specific path to the solution.

Today, machine learning algorithms make it possible to “train” a computer to think in remarkable ways -- from being able to understand what a cat looks like from a collection of cat pictures, to differentiating between spam and legitimate email. Instead of the human programmer mapping out the thought process, the machine learns from massive volumes of data to make connections and associations.

> “Machine Learning is the science of getting computers to act without being explicitly programmed.” **– Andrew Ng**

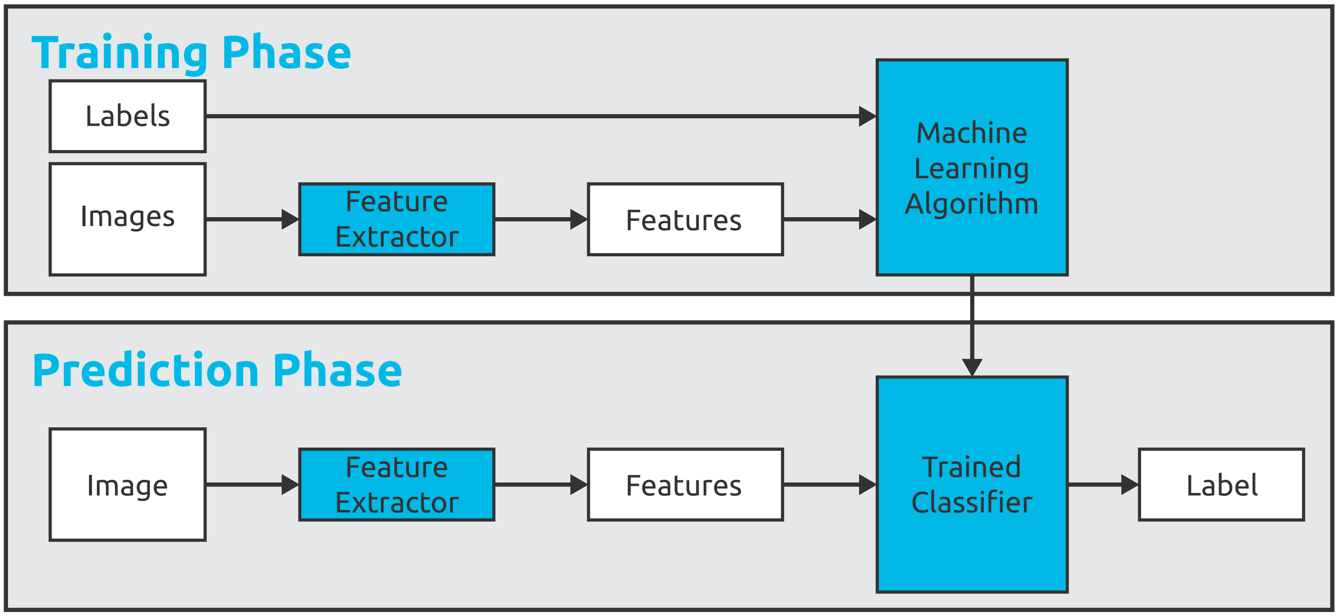

Machine Learning Phases

The significance of machine learning cannot be overstated. For the first time in human history, a machine is being given the capability to learn without the solution being specifically understood and programmed. The implications of this are huge – and illustrates why machine learning is the most powerful part of the AI landscape to emerge so far. Machine learning consists of three steps:

- Collect - lots of data
- Train - a model on that data
- Predict - use the trained model to make predictions on new data

Collect Phase

There is already lots of data within your enterprise: in your systems of record, within your websites and mobile apps, as well as partner data and data you have purchased. Based on the business problem you are trying to solve, determine what data you need, and what's available. Where are the gaps? Start addressing those.

Training Phase

During the training phase, the model continuously iterates and gets feedback on its accuracy. This process requires massive amounts of computing power to get the model just right.

Prediction Phase

The "training" phase is followed by the "prediction" phase where the model is put to actual use.

07 MACHINE LEARNING: TYPES OF MACHINE LEARNING

There are three major categories of machine learning. They represent three distinct learning models: Supervised Learning, Unsupervised Learning and Reinforcement Learning.

In **Supervised Learning**, the learning algorithm is "trained" using example inputs along with ideal or desired outputs. Through this process, the algorithm eventually learns to accurately map random inputs to correct outputs. There are two main types of supervised learning:

- Classification: Classifies output in two or more distinct categories: for example, is this a spam email or not? Customer Segmentation
- Regression: Predicts a numeric valued output: e.g., propensity – say, a number between 1 to 100 – for a customer to buy this product

In the **Unsupervised Learning** model, the algorithm is not provided with desired outputs. It must self-discover a structure or patter to the unlabeled input data.

Reinforcement Learning relies on the algorithms interacting with the environment in trial-and-error mode, receiving positive and negative feedback in the process.

Supervised Learning	Unsupervised Learning	Reinforcement Learning
Training data includes inputs and desired outputs	Training data is input only; does not include desired outputs	Training data has input and critic/reward from the sequence of actions
Separate learning and prediction phase	Separate learning and prediction phase	• Learning and prediction phases simultaneously • Online learning • Learn by trial and error

Common applications of machine learning:

UNSUPERVISED

Using clustering analysis to segment customers by behavoral characteristics. Analyze prospects by social structures and other key demographic properties.

SUPERVISED

Predict credit worthiness of credit card applicants and credit card holders. Shape attributes around deliquent behavior.

REINFORCEMENT

Generate offers and options for banking users as theit behaviors change over time. Shape clustering offers as user accept or reject offers.

Deep learning is a type of advanced machine learning algorithm at the forefront of new breakthroughs in AI. It is largely responsible for the sudden resurrection of the entire field.

Deep learning models the human brain's neural networks of connections to make predictions about new data sets.

Deep neural networks learn by adjusting the strengths of their connections that are propagated through the layers. A neuron weighted more heavily than another will exert more of an effect on the next layer of neurons. The final layer puts together these weighted inputs to come up with an answer.

As an example, let's say we want a to recognize pictures that contain a dog. To achieve this, we start the training process by compiling a positive set of images – thousands of examples of dog photos, which are labelled (by humans) as "dog", and negative set of images of other animals, objects and things that are not dogs, which are labelled as "not dog." When these images are fed into the neural network, the image is converted into data which moves through the network, and various neurons assign weights to different elements: features that identify a dog. At the end, the final output layer puts together all the pieces of information – pointed ears, tail, black nose – and spits out an answer: dog.

Learning From Experience

Deep neural networks learn by adjusting the strengths of their connections to better convey input signals through multiple layers to neurons associated with the right general concepts.

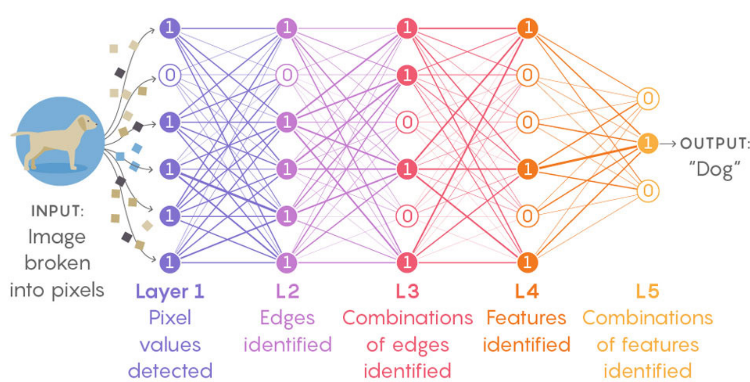

When data is fed into a network, each artificial neuron that fires (labeled "1") transmits signals to certain neurons in the next layer, which are likely to fire if multiple signals are received. The process filters out noise and retains only the most relevant features.

Source: https://www.quantamagazine.org/new-theory-cracks-open-the-black-box-of-deep-learning-20170921/

07 MACHINE LEARNING: DEEP LEARNING

Traditional Machine Learning Flow

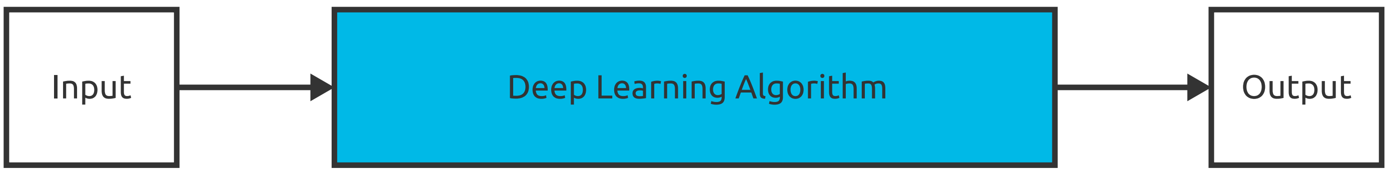

Deep Learning Flow

Deep learning, as a subset of machine learning overall, is also a combination of two steps – training and inference. Training involves teaching a neural network to recognize objects, voices, etc., just like the neurons in a child's brain are taught to do so as they learn about life.

But, how does machine learning differ from deep learning? Here are some key distinctions between the two.

	Machine Learning	Deep Learning
Training data volumes	Thousands of data points	Requires Big data – millions of data points
Output	Numerical value, like a classification or a score	Anything from numerical value to free-form elements like text and sound
How it works	Uses various types of algorithms that learn and then predict future actions	Uses neural networks that pass data through many processing layers to interpret data features and relationships
How it's managed	Algorithms are directed by data analysts to examine specific variables (features) in data sets	Algorithms are largely self-directed on data analysis

08 COMPUTER VISION

We humans use our eyes and our brains to see and visually sense the world around us. Computer vision is the science that aims to cultivate a similar capability in a machine or computer.

As is the case with deep learning, computer vision uses algorithms to achieve automatic visual understanding. It identifies images of visible objects (i.e., people) as well as patterns that cannot be seen (i.e., time-lapsed images, infrared images, etc.).

Some well-known examples of vision application include self-driving cars, object detection, facial recognition, Mobile Augmented Reality, optical character recognition, and games such as Xbox Kinect.

Computer vision is concerned with the automatic extraction, analysis and understanding of useful information from a single image or a sequence of images.

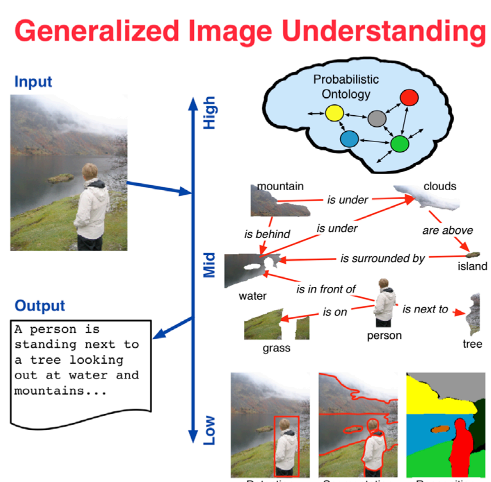

Source: http://web.eecs.umich.edu/~jjcorso/r/career/index.html

09 NATURAL LANGUAGE PROCESSING (NLP)

NLP enables computers to understand written and spoken human language – including complex utterances such as accents and slang terms. NLP applies intelligence to different aspects of language: phonology, similar sounding words, syntax, semantics, intent and context. Traditionally, NLP has taken a rules-based approach -- such as keyword matching (“equal to” or “contains”). More recently, NLP uses machine learning/deep learning to understand the context from the text.

Some common examples of NLP include:

- Information Extraction
  - E.g.: An application that automatically detects events in email (from the text) and prompts adding them to your calendar
- Language Translations
- Summarizations – automatically summarizes a long article
- Sentiment analysis – identifies the general emotion expressed in a piece of text – e.g. Is the product review comment positive or negative? Is the comment a compliment or a complaint? Is the article bullish or bearish towards a particular stock?
- Text Classification
  - E.g. Routing of inbound customer email to the appropriate department based on email content
  - E.g. Email classification for spam detection
- Conversational Agents (chatbots)
- Information Discovery
  - E.g. Parsing news articles for specific information such event, topic, person, place, etc.

NLP engines are particularly useful for building conversation applications that rely on speech-to-text and text-to-speech.

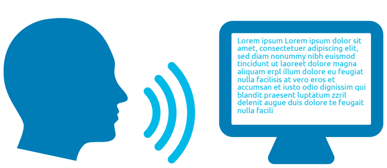

Speech-to-Text

Hardware devices such as Amazon Echo capture human speech (as analog) and convert it to text (digital) in a particular language (like English).

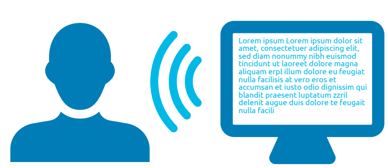

Text-to-Speech

Text found in documents, emails and other interactions are read and translated using NLP algorithms, then converted to speech. E.g. Amazon Echo’s backend cloud-based system generates a response based on available content/data/API that is then converted to speech.

10 APPLYING ARTIFICIAL INTELLIGENCE IN THE REAL WORLD

Given the quickly evolving power of AI, its transformational impact will extend into all facets of how we work, live, and play. Machine learning, computer vision and natural language processing are distinct capabilities that help to explain the what of artificial intelligence. As for how AI can be used to actually solve real world problems, consider three functional areas:

- **Human & Machine Interaction** – Enhancing the communication and interaction between human beings and computers including: conversational interfaces; chatbots; virtual assistants; improving the customer journey; website personalization; recommendations, etc.
- **Process & Machine Interaction** – Leveraging AI to enhance the execution of operational processes by automating complex knowledge worker tasks (sometimes referred to as robotic or smart process automation). E.g.: Automating customer onboarding process or order processing
- **Advanced Analytics** – Using advanced machine learning algorithms to discover, classify and identify patterns in data (commonly referred to as deep learning). E.g.: Analyze payment transactions to identify fraudulent ones.

11 CONCLUSION

AI such as Machine Learning and Deep Learning, NLP, and Computer Vision – all point to a future in which our platforms and systems are smart enough to learn from human interactions and data. These tools will intuitively assist with our requests, anticipate our needs, take care of mundane and forgotten tasks and remind us of important ones. AI can connect the various nodes of our lives (home, work, travel) into one experience that moves seamlessly with us from house to car to office and back again. Today our smartphones channel the magic; tomorrow a new generation of devices will be the conduits. And by then, AI will be the fuel that powers everything digital.

People Matter, Results Count.