

WHAT IS MACHINE LEARNING?


MACHINE LEARNING

One of the most promising and active fields in artificial intelligence (AI) is machine learning. Rather than describing to the machine how to perform a task, the idea is to give the machine many thousands of examples and let it learn by itself. For example, by showing it the many ways of writing the number 4, the machine will eventually become very good at recognizing it (and many post offices in the world now analyze hand-written addresses on mail using AI). There are many tasks at which machine learning is far better than humans – from analyzing x-rays for potential signs of cancer to catching fraudulent credit card transactions as they happen – and the list continues to grow.

4 → 4 2 → 2 3 → 3
 4 → 4 9 → 9 0 → 0
 5 → 5 7 → 7 1 → 1
 9 → 9 0 → 0 3 → 3
 6 → 6 7 → 7 4 → 4

DEEP LEARNING AND NEURAL NETWORKS

Most of the ideas behind AI are not new. One type of machine learning is based on the idea of neurons in the brain that get “fired” by certain triggers and send signals to other neurons. These “neural networks” with millions of nodes can tackle complex problems by splitting them up in many smaller steps. For



instance, a neural net that recognizes cats might have a few nodes that determine whether the object is furry. The next layer may look for eyes or legs or color patterns. These layers are developed automatically by training the system with thousands of examples. Neural nets with many layers are sometimes referred to as *deep learning*. While a decision tree model (“Does it have eyes?” Is it furry? Does it move?”) is understandable to humans, in the case of neural nets that have automatically created tens of layers on their own, the meaning of each layer is incomprehensible to us. This can lead to ethical questions, as we can no longer explain a decision made by a neural net in terms of a step-wise deductive model.

instance, a neural net that recognizes cats might have a few nodes that determine whether the object is furry. The next layer may look for eyes or legs or color patterns. These layers are developed automatically by training the system with thousands of examples. Neural nets with many layers are sometimes referred to as *deep learning*. While a decision tree model (“Does it have eyes?” Is it furry? Does it move?”) is understandable to humans, in the case of neural nets that have automatically created tens of layers on their own, the meaning of each layer is incomprehensible to us. This can lead to ethical questions, as we can no longer explain a decision made by a neural net in terms of a step-wise deductive model.

Neural networks are good in environments that are:

1 Highly complex

2 Constantly changing

WHY IS MACHINE LEARNING A FUNDAMENTAL CHANGE IN COMPUTING?

We have long passed the stage where the amount of data collected from computer systems, from monitoring of human activity to internet-connected devices (Internet of things or IoT) can no longer be processed manually and must be done by computer. As data continues to increase exponentially it is becoming unfeasible for humans to write the automated routines that dictate how to draw conclusions from this data. Machine learning automates this step, which makes it possible to obtain benefits from the avalanche of data available.

Another reason neural networks are a game-changer is that they are scalable. Traditional software programs get problematic when they



become larger – they can no longer be easily overseen and are prone to bugs (errors) or security vulnerabilities. Neural nets can add more layers of nodes without increasing complexity. It is possible that in the future a large portion of all software will be replaced by neural nets.

APPLICATIONS FOR THE UN SECRETARIAT

There are many potential applications for the UN; opportunities arise when ample data (textual, numerical, images or any type of signals) is available. Examples are classification (tagging documents or images with particular topics); image recognition (e.g interpreting satellite data to assess disaster-hit areas or understand the effects of climate change); prediction of events (political, economic or security related); and tools to aid translation or automated conversations (“bots”) that can provide assistance to staff or to the public.

Interestingly, machine learning, particularly neural nets, are good at two types of problems: (a) systems that are so *complex* that we cannot describe them in a model (therefore writing out an algorithm for the task is not an option) and (b) systems that constantly *change*, making it hard to describe them to a computer. The UN’s challenge is to understand the dynamics of human society, which is a challenge that contains both characteristics. Machine learning should be an excellent tool to help us in that respect - one step at a time.