



**DINO ESPOSITO
SIMONE MASSARO**

ARTIFICIAL INTELLIGENCE
for executives and policy makers

*Vision and tools to predict
the unpredictable*

**Dino Esposito
Simone Massaro**

ARTIFICIAL INTELLIGENCE
for executives and policy makers

*Vision and tools to predict
the unpredictable*

Dino Esposito, Simone Massaro

ARTIFICIAL INTELLIGENCE
for executives and policy makers
Vision and tools to predict the unpredictable

© Dino Esposito, 2019
© Simone Massaro, 2019

No part of this publication may be reproduced, transmitted, transcribed, stored in a retrieval system, or translated into any language or computer language, in any form or by any means, electronic, mechanical, magnetic, optical, chemical, manual or otherwise, without the prior written permission of one of the authors.

A monumental *Thank You* goes to the engineering team at BaxEnergy for their general support and the effort put in reviewing the text.


An even bigger *Thank You* is reserved to a couple of special people for their relentless help in making sense of the trickiest aspects of machine learning algorithms: **Francesco Esposito** (Youbiquitous) and **Tiago Santos** (Smartwatt).

If you will find this reading easy and pleasant, well, that was because of them! If you still find it hard, well, that's only our fault.

For a deeper technical overview of machine learning algorithms and programming techniques, including a basic explanation of the mathematics behind and some programming examples in Python and .NET, here's a further reference.

Introducing Machine Learning

Dino Esposito, Francesco Esposito
Microsoft Press, 2019



ISBN-10: 0135565669

ISBN-13: 978-0135565667



Artificial intelligence would be the ultimate version of Google.

The ultimate search engine that would understand everything on the web. It would understand exactly what you wanted, and it would give you the right thing.

We're nowhere near doing that now.

However, we can get incrementally closer to that, and that is basically what we work on.

LARRY PAGE

GOOGLE FOUNDER

INTRODUCTION

Let's say it upfront: Artificial Intelligence, or just AI, is today a wildly abused term that many uses without the faintest clue of its true essence. 9

So, C-level executives use to sprinkle AI on top of their presentations, like icing on the cake, to showcase the company and services.

So, media representatives and know-it-all pundits use AI as the thread to weave their web of words and look smart and up-to-date.

So, gurus, futurists and science popularizers draw from the well of AI to build fascinating doomsday's stories, half fiction and half fantasy.

In the middle, slowly parachuting to ground from the heights of mathematics and statistics, are true scientists and software engineers.

Why is AI being talked a lot?

AI triggers in humans the same reaction as Sci-Fi stories: curiosity, wonder and scare. It raises a morbid desire to burn in the hell of future and an irrational disappointment for not being part of it. In a way, AI represents the closest we can go today to the realiza-

tion of our most futuristic dreams – those same dreams captured in Sci-Fi books and movies.

Dreams that now really seem real.

Artificial Intelligence is only an umbrella term for a new branch of software systems but nothing you can get without dedicated software investments. You can imagine AI as a city map with lots of attractions all around and a big red dot saying “You are here”. You can even pay a cheap ticket to visit attractions, but the real cost is getting there and it is up to you!

Machine Learning is the solid part of what nearly everyone calls AI and it is just software, problem solving and consulting. Any executive and every engineer should look into it and much deeper than into, say, microservices. Why is that? Because machine learning is a breakthrough and allows to do things that aren’t likely possible otherwise.

10 The challenge is figuring out just those things that are most relevant for your business.

This book sets the (rather ambitious) goal of explaining what is artificial intelligence and how companies can realistically take advantage of it, especially in the context of renewable energy.

“

**THE PROBLEMS OF THE WORLD
CANNOT POSSIBLY BE SOLVED
BY SKEPTICS OR CYNICS WHOSE
HORIZONS ARE LIMITED BY THE
OBVIOUS REALITIES.
WE NEED MEN WHO CAN DREAM
OF THINGS THAT NEVER WERE,
AND ASK WHY NOT.**

John F. Kennedy

Speech to the Irish Parliament,
Dublin, June 1963

Dartmouth.
August 31, 1955.

A PROPOSAL FOR THE DARTMOUTH SUMMER RESEARCH PROJECT ON ARTIFICIAL INTELLIGENCE

¹² We propose that a 2-month, 10-man study of artificial intelligence be carried out during the summer of 1956 at Dartmouth College in Hanover, New Hampshire. The study is to proceed on the basis of the conjecture that every aspect of learning, or any other feature of intelligence, can in principle be so precisely described that a machine can be made to simulate it. An attempt will be made to find how to make machines use language, form abstractions and concepts, solve kinds of problems now reserved for humans, and improve themselves. We think that a significant advance can be made in one or more of these problems if a carefully selected group of scientists work on it together for a summer.

The following are some aspects of the artificial intelligence problem:

- 1) Automatic Computers
- 2) How Can a Computer be Programmed to Use a Language?
- 3) Neuron Nets

- 4) Theory of the Size of a Calculation
- 5) Self-Improvement
- 6) Abstractions
- 7) Randomness and Creativity

In addition to the above collectively formulated problems for study, we have asked the individuals taking part to describe what they will work on. Statements by the four originators of the project are attached.

The originators of this proposal are:

1. C. E. Shannon, Mathematician, Bell Telephone Laboratories
2. M. L. Minsky, Harvard Junior Fellow in Mathematics and Neurology
3. N. Rochester, Manager of Information Research, IBM Corporation, New York
4. J. McCarthy, Assistant Professor of Mathematics, Dartmouth College.

PART 1 – The origins of Artificial Intelligence

“

ARTIFICIAL INTELLIGENCE...
I'VE BEEN FOLLOWING THAT SINCE
I WAS IN HIGH SCHOOL

Paul Allen

American business magnate,
co-founder of Microsoft

1 HUMANS AND MACHINES

When we think about Artificial Intelligence (AI) the mind jumps beyond any observable reality and leads to imagine a world in which robots and cyborgs control every aspect of life.

Artificial intelligence scares and wonders us.

On one hand, we can't wait for it to come in our everyday life but, at the same time, we're terrified by the dramatic changes it may bring. Anyway, the wish of a machine that irons out troubles and issues is somehow innate in the human nature.

As humans, we have always dreamt of artificially built machines capable of reasoning – and possibly thinking – like us and even deeper than us. In the plot of works by ancient Greek authors Aeschylus and Euripides, a device was sometimes used to simulate the intervention of the god. Latins later called it *deus-ex-machina* to indicate an entity able to resolve human conflicts that humans were unable to settle down alone.

Interestingly, in the course of history, polymaths and philosophers, many times, attempted to provide a theoretical formulation for the mechanics of the human thought.

From Leibniz to Gödel

The first polymath who imagined something close to what we today call artificial intelligence was Raymond Lull – a Spanish theologian who lived in the 13th century in the island of Majorca.

The legend says that Lull had a vision in 1272 during a solitary spiritual retreat on the top of a mountain. He envisioned a series of basic principles that, properly combined together, could lead to the core principles of every science and every form of knowledge. In his works, Lull used symbols to represent each basic principle so that each basic truth could be expressed in a formal way. Way ahead of his time, Lull didn't go further than using the system as an aid for exposition and memory.

18

In the middle of the 17th century, Gottfried Leibniz built on the foundation of the Lull's work and conjectured that the human thought could be systematized in an array of algebraic rules so that any argumentation could be reduced to some basic mechanical calculation.

The Leibniz's work inspired further development of mathematical logic that George Boole and, in the early 20th century, David Hilbert and Bertrand Russell carried out.

In particular, Hilbert devoted his entire course of studies to reconsider the relationships between the five axioms of Euclid. More precisely, Hilbert speculated about the relationship between the first four axioms and the fifth. Hilbert challenged all mathematicians of his time to prove the statement "*Can all mathematical statements be expressed through a set of well-defined rules?*" His ultimate goal was finding a way to formalize all known mathematical reasoning in much the same way Euclid well did for his time.

Put another way, Hilbert questioned whether a single set of ax-

Raymond Lull
Palma
1232 – 1316



19

Gottfried Leibniz
Lipsia
1 July 1646
—
Hannover
14 November 1716



ioms could exist that allows to derive all mathematical statements from them.

No relevant answer came for a couple of decades and to many it just looked like an abstract, meta-mathematics problem. Then, all of a sudden, from a side field of research a definitive answer popped up.

It was 1931.

The Theorems of Incompleteness

Born in today's Czech Republic, Kurt Gödel is unanimously considered one of most impactful logicians in the history of mankind. In spite of being nowhere near as famous as, say, Aristotle, the impact of his work on the progress of mankind is, at the very minimum, immense and as John Von Neumann – the father of modern computers – said “*a landmark which will remain visible far in space and time.*”

Winner of the US President's National Medal of Science in 1974 for mathematical, statistical and computer sciences, and professor at Princeton in the same years as Albert Einstein, Gödel made it to history when he was 25 and published the paper “*On Formally Undecidable Propositions of Principia Mathematica and Related Systems*”.

In the paper, Gödel demonstrates a couple of propositions of mathematical logic collectively immortalized as the Theorems of Incompleteness.

In any formal system, expressive enough to model the arithmetic of natural numbers, Gödel proved, there is at least one statement that can't be proven true or false following the axioms of the system. Furthermore, Gödel proved that even though the statement is assumed true (or false), then any further reasoning from the statement will ultimately lead to another undecidable statement.

The whole scientific community interpreted the incompleteness theorems as a negative answer to the Hilbert's original question: “*Can all mathematical statements be expressed using a set of well-defined rules?*”

David Hilbert
Königsberg
23 January 1862
—
Göttingen
14 February 1943



Kurt Gödel
Brno
28 April 1906
—
Princeton
14 January 1978



It's fairly obvious in the end: if in an axiomatic system an undecidable statement always exists, then the dream of a unified set of rules to explain the entire mathematics just vanishes.

Formalization of Thought

Great, but where does the greatness of Gödel ultimately come from? In which way, are the two incompleteness theorems such a landmark in logic?

Well, on one hand Gödel proved that there are things that mathematical logic can't just prove. On the other hand, though, Gödel proved that, within the limits of a consistent formal system, any reasoning can always be expressed as a set of formal rules and then, in some way, mechanized.

22 This aspect is monumentally relevant as it sets the theoretical foundation for mechanical reasoning – the ancestor of modern computer-based reasoning.

All this was happening in the early 1930s and the horror of the second world war was still to materialize. In those years, there were no things like computers. Moreover, any such thing like a modern computer was completely beyond imagination. Yet, Gödel set the ground for mechanical computing which later developed into electronic calculators and ultimately in forms of artificial intelligence.

The results achieved by Gödel gave the spark to three parallel and independent research paths that conveyed to the same result in a few years, around the mid-1930s.

The same Gödel in 1933 formulated the concept of *general recursive functions* – a computable logical function that takes a limited array of natural numbers and computes a natural number. In 1936, Alonzo Church, a logician from the Princeton university, defined *lambda calculus*, a formalism able to express any computation based on natural numbers. At nearly the same time, in a fully independent way, British mathematician Alan Turing built



Alonzo Church
Washington
14 June 1903
—
Hudson
11 August 1995

23



Alan Turing
London
23 June 1912
—
Manchester
7 June 1954

the theoretical model of a computing machine, the popular Turing machine, to perform symbolic calculations on an infinite tape.

The three classes of computable functions were later unified in the Church-Turing thesis. As a result, a function is computable in the lambda calculus if and only if it is computable in the Turing machine if and only if it can be defined as a general recursive function.

What's the point, anyway?

The Church-Turing thesis makes conceivable building a mechanical device able to reproduce the process of mathematical deduction through the manipulation of symbols. Seen with the eyes of today it seems foregone, but it was the late 1930s and these people – Gödel, Church and Turing – were just three men who dreamt of things that never were. And, better yet, they not just dreamt of those things, but also gave a substantial contribution to make them happen!

Well, sort of!

While Gödel and Church remained in the realm of logic and abstraction, Alan Turing physically built a few electromechanical and programmable prototypes during his war years at Bletchley Park.


Imagine for a moment to be in the shoes of Alan Turing. The war is over and with it also ended the dramatic urgency of breaking Nazi ciphered codes to save lives and change the outcome of the war. However, you now know that you can build a machine that can calculate numbers from numbers. But you also have strong theoretical evidence that you can also build a machine that can compute anything you can express through a consistent grammar of symbols.

This is, well, unbelievable!

In the shoes of Alan Turing, you'd probably feel like god and you too would probably foresee somewhere ahead of you, but clearly identifiable, a machine that can behave in much the same

way as humans do. And at that point, you too, probably, would ask the same crucial question.

Can machines really think?



BLETCHLEY PARK

A mansion in Buckinghamshire, 50 miles north-west of London, Bletchley Park was established in 1938 as the central site for British and allied code breakers. Best mathematicians and cryptanalysts of the country were conveyed to try to decrypt the various German codes. In particular, the teams at work focused on the Enigma and Lorenz ciphers used by U-boats and high command respectively.

It has been estimated that the effort conducted at Bletchley Park shortened the second world war by at least two years and, more importantly, set the result of the war. For example, it is known as a fact that decrypted communication played a key role in the D-Day. Alan Turing was a prominent member of the Bletchley Park personnel.

Birth of Artificial Intelligence

The term Artificial Intelligence was officially coined in 1955 and used for the first time in the paper quoted at the beginning of the book.

In the summer of 1956, John McCarthy – professor at Dartmouth College – organized a six weeks summer research workshop in New Hampshire. He invited a dozen of colleagues with very different backgrounds, but a common interest for formal logic, computation, and thinking models. Invited members came from research fields like mathematics, engineering, psychology, neurology.

The workshop had the purpose of a brainstorming session around the idea of thinking machines and each participant was supposed to present a work. In the summary, McCarthy deliberately used a new, kind of neutral term just to stay away, and ideally unify, the two (partly conflicting) souls of ongoing academic research.

At the time, the abstract theme of intelligent machines was researched in two distinct contexts. One was the automata theory, directly based on the work done by Church and Turing. The other context was based on cybernetics, directly descending from the theoretical work done by Charles Babbage in the 19th century that the team led by John von Neumann was turning into concrete hardware.


The ultimate purpose of the Dartmouth workshop was rather ambitious: devising the theory of an artificial brain. Looking back, it is interesting to notice that while computers – as we know them today – were not available yet, the idea of a special application, called artificial intelligence, was already around.

Today, Artificial Intelligence is (again?) only an umbrella term under which we can spot two macro areas: expert systems and autonomous systems. An expert system is a system that knows how to effectively react to its input suggesting or making a thoughtful de-

John von Neumann

28 December 1903

—
8 February 1957



FIRST-EVER NEURAL MODEL

In 1943, Warren McCulloch and Walter Pitts devised a computational model inspired by the known structure of the brain neuron. Central in their vision was the artificial neuron—a function that receives one or more weighed inputs and sums them to produce an output. The model consists in a concatenation of multiple neurons. Still today, their seminal piece of work represents the foundation of the simplest (but working) neural networks. Two facts are remarkable. One is that at the time McCulloch and Pitts proposed their model there was no other concrete computational model and no physical computers. The other is that further development from the notion of artificial neuron took several years to materialize and for some (mostly business) reasons the whole neural idea was put aside until it resumed in the 1980s around the theory of neural networks.

cision. In a way, an expert system is the software counterpart of a human expert and a sort of decision aid system. An autonomous system is a system designed to autonomously recognize a number of scenarios and successively work as an expert system on data they have never seen before.

Autonomous systems are the new frontier of Artificial Intelligence. We globally refer to them as Machine Learning.

QUICK FACTS

- Gödel incompleteness theorems gave the spark to a research that culminated in the nearly simultaneous formulation of recursive functions, Turing machine and lambda calculus.
- Unification of the three computing models set the ground for programming languages and machines. The strict necessities of war fastened development and turned it into real progress after the end of the war.
- Artificial Intelligence is an umbrella term, originally coined to unify research fields and today called to unify machine learning and intelligent software applications.

NEXT UP

We all abuse of the words artificial intelligence and label with it nearly any solution we build or suggest in the industry. What companies really need, instead, is intelligent software solutions.



**A GREAT DEAL OF INTELLIGENCE
CAN BE INVESTED IN IGNORANCE
WHEN THE NEED FOR ILLUSION
IS DEEP**

Saul Bellow

Canadian-American writer, winner of the Nobel
Prize for literature in 1976

2 INTELLIGENT APPLICATIONS

C-level and decision makers, in any industry, are today under the attack of a (viral) form of artificial intelligence hype. Artificial intelligence is sensed as a broadly available kind of service powerful enough to push companies in exciting new territories, full of business opportunities.

Doubtless, there's a lot of expectation these days around artificial intelligence and sometimes it even seems unstoppable and pathological. The substance that lies under the surface of such a great level of expectation, however, is nearly the same it was back in the 1980s.

What is artificial intelligence (AI), exactly?

- AI is not a service you can buy somewhere at some IT store.
- AI is not a plug-and-play business solution.
- AI is just software you plan, design and build.

AI, however, is different from any other type of software you may have bought or built in the past. The funny thing about AI is that it really looks brand-new, it effectively forces implementors to learn

and practice new things but, at the very end of the day, it's nothing completely new that didn't exist before.

For the part that today mostly affects the industry, AI consists in the combined application of optimization and data mining techniques, just in a broader scope than only a few years ago. If it looks like magic, it's essentially because of the technical superficiality of most of the circulating stories, artificially built for marketing purposes.

AI is not magic; but it isn't unreal either!

The Timeless Dream of Emulating the Human Brain

Largely responsible for the hype around AI is the term “intelligence”. The term amplified beyond imagination the perception of the power hidden in AI and created unrealistic expectations by the sole word of mouth.

AI doesn't deliver anything per se; AI just allows companies to build end-to-end solutions with more inherent software intelligence inside than today. Yet, the word *intelligence* brings our minds straight to figuring out the electronic counterpart of the human brain and to seeing as real, or just upcoming, the highly automated world that Sci-Fi fiction imagined already a century ago.

**WITHOUT BEATING AROUND THE BUSH, IN 2019
NO ONE KNOWS HOW TO BUILD THE SOFTWARE
COUNTERPART OF THE HUMAN BRAIN.**

One reason is the effective lack of knowledge about the internal functioning of the brain. Another reason is the objective difficulty due to the different architecture of the biological brain compared to the architecture of today's computers.

As long as neuroscientists learn deeper about structure and be-

havior of the brain, mathematicians have a better chance to devise the abstract model that could reproduce it. To be implemented, any abstract model devised so far requires a computing architecture radically different from that of modern computers. Remarkably, the computational model in use today is the same model that John Von Neumann figured out in the 1950s.

The Von Neumann's model was designed to be a concrete implementation of the abstract Turing Machine – a machine made of one processor only capable of processing symbols sequentially.

The little-known part of the story is that John Von Neumann outlined a profoundly different model that premature death left half-done (1957, aged 53). According to his unfinished studies, such a model envisioned a universal constructor machine able to process copies of itself rather than symbols. This vision, also known as *self-replicating automata*, was inspired by the shared knowledge of the functioning of the human brain of the time. Other researchers later resumed his work, but ended up producing, even in recent years, only lab prototypes.

So, modern computers found on the principle that processing and memory units are separated on the chip making transfer of data unavoidable. This is quite different from how the brain is known to work.


The brain is made by distributed neurons each of which performs local computing, therefore having memory and processing units integrated in the same body. In addition, the brain can perform primitive operations that are fairly advanced for computers – for example, pattern recognition.

In addition, it is known that neurons and synapses (the physical links that connect two neurons) observe a stochastic behavior, meaning that their action is not characterized by a fixed, reliable and deterministic set of principles.

Individually, each neuron may be sluggish and imprecise, but the dense network each neuron is part of compensates the poor indi-

vidual performance and overall allows for fairly effective operations, also in terms of energy demand and consumption.


In humans, a neuron is a cell designed to receive, process and transmit information to other nerve cells as well as muscles and glands. A neuron consists of a body that ends in an array of thin filaments called *dendrites* at one end, and a nerve fiber called the *axon* at the other end. The figure shows it.



34

Dendrites exit at one end of the neuron body and the axon exits at the other end

Dendrites are like input devices and serve as the receptors of any transmitted information. The axon is responsible for the transmission of information towards receptors of other connected neurons. The transmission occurs through electrical impulses. The diagram in the figure below illustrates a small segment of the dense network of neurons and their synaptic connections.



35

Dendrites from multiple neurons connect to the same axon of another neuron, thus forming a synapse.

It is reckoned that the link between two (or more) neurons results in storage of information and then memory.

By the way, note that the figure only captures and illustrates the most common type of synapse in the mammalian nervous system – the *axo-dendritic* synapses. In nature, other types of synapses exist such as the *axon-to-axon* synapse.

In a nutshell, the brain features an inherently distributed architecture made of low-speed and low-precision building blocks. If we compare the single neuron to a computer, we would then have a very slow and unreliable computer. At the same time, though, we experience every day that the human brain is a marvelous example of efficiency and computing power (i.e., consciousness, thinking, but also image recognition and real-time reaction).



THE BRAIN CLOCK

The human brain is quick at tasks like recognizing images that would be quite intensive even for the fastest today's computers. How fast is the brain?

The brain has no centralized system clock and neural activity occurs over multiple frequencies, the fastest of which is below 100 Hz. A neuron is therefore hardly able to light up for transmitting an impulse more than 100 times a second.

This is nothing compared to the 3 GHz of today's personal computers and smartphones.

36

What's the bottom line?

A logical gap exists between two distinct entities, brain and computer, we insist to relate together. Brain and computer are equally powerful but in different ways because they are based on different architectures.

What about computer intelligence?

The effective intelligence we can build today in computer applications depends on how good we emulate with software the complex neural learning model of the brain.

Primordial Software Intelligence

Which forms of intelligence are realistically and concretely actionable in modern software?

Intelligent software is software capable of sensing the surrounding environment (including end users) and reacting to observed facts and changes.

This definition incorporates many aspects. For example, as humans we can easily recognize a phone number regardless of how it's written. Most web sites, instead, just fail at that. Many times, the software is too rigid when it processes input data or too lazy. Why, for example, should not "01" be understood as "1" if it is known that the context requires a number? Going further, why not the computer should try to guess the age of the user and adjust font size accordingly?

This is software intelligence, but no techniques commonly referred to as AI are involved in it.

As minimalistic as it may sound, a conditional statement, an *IF-THEN-ELSE* rule, is already a (rather primordial) form of software intelligence. And any coded flow diagrams are other primordial forms of software intelligence too. The more software architects understand the business domain and the natural processes under the surface, the more intelligent they can make the final product. Today's most sophisticated forms of software intelligence revolve around the idea of giving users a better experience. The real challenge, though, is bringing intelligence in the processes.

37

THIS IS THE REASON WHY ANY CEO NEEDS TO GET INVOLVED WITH ARTIFICIAL INTELLIGENCE AND NEEDS TO FOCUS MORE ON THE ESSENCE OF IT THAN THE AVAILABLE AI TOOLS AND SERVICES OFFERED BY RAMPING VENDORS.

Most software applications today limit to mirror a superficial, and partly blurred, vision of reality. Most applications are devised through the lens of trendy architectural principles and buzzwords (e.g., microservices, cloud, Blockchain and, of course, AI and deep learning) without focus on processes. It's the way the company works and does business that can be improved by artificial intelligence.

SOFTWARE IS ONLY THE PHYSICAL MEAN THROUGH WHICH AI IS ULTIMATELY DEPLOYED.

Companies don't need artificial intelligence per se; companies need more intelligent software. And artificial intelligence is the fastest way to reach higher levels of intelligence.

38

For many years, research focused on creating the counterpart of a human expert – a software program able to extricate itself in a fixed fold of rules and scenarios. For years, *expert systems* have been the cutting edge of intelligence applied to a number of industry fields. Still largely used today, they're destined to become legacy systems soon.

Expert Systems

An expert system is a software system applied to a given field. If competently programmed, an expert system can give nearly the same answer a human expert would give in front of similar input data. They are smart and fast if the provided input lets the engine find the path to an answer. Otherwise, they fail or, like a human, just surrender unable to provide an answer. Furthermore, unlike humans, expert systems *are not able to learn*.

The first system labeled as an expert system was built at Stanford in the 1960s and was designed to analyze the spectrographic data of chemical substances to conjecture their underlying molec-

ular structure. The overall performance of the system was reckoned analogous to that of human experts. This finding gave the spark to many other similar projects in both the industry and universities.

The key aspect of an expert system, however, is that its processing power is limited to the fixed set of decisional pathways and input parameters that have been hard-coded in the supporting software. Architecturally speaking, an expert system results from the composition of two sub-systems: a repository of known facts (knowledge base) and an inference engine. The system works processing input against the known list of facts and applying inference rules repeatedly. A good example of a common-use expert system is the auto-pilot mode of aircrafts.

Expert systems can be generally employed with proven success in all fields where a human expert would be too expensive or simply not available 24x7. Expert systems have a successful record in aerospace, healthcare, criminology, networking, geology and financial portfolio management.

39

Unfortunately, an expert system only works well within the boundaries of a given cognitive process and may not respond reliably in situations outside the codified range of expertise. Maintenance is then a relevant shortcoming of expert systems. Once released in production the set of answers an expert system can give is fixed and can be extended or modified only through a new release.

Even though expert systems are still largely and successfully used today, they're considered an old-fashioned approach for greenfield developments. Recently, neural networks have been rediscovered and data mining algorithms have been examined from a different angle. The result is that more autonomous systems, capable of being trained to learn, have become possible and came to the rescue.

This is where *machine learning* comes into play.

When today someone mentions solutions based on artificial intelligence, all that she means is solutions based on some machine learning pattern.

Before having a crash course on machine learning, we need to focus on the preliminaries of any machine learning projects such as availability of high-quality data and, even before that, the business perspective of machine learning solutions.

QUICK FACTS

- The behavior of a computer is sequential whereas it is distributed and parallel in the brain. This is a key impedance mismatch.
- The need for artificial intelligence is ultimately only the need for more intelligent software.
- A distinction between *regular* software and *artificial intelligence* is not the right way to make software smarter and better.
- Although still functional, expert systems are set to be overtaken by end-to-end solutions extensively based on machine learning.

NEXT UP

Humans always dreamt of creating the electronic counterpart of the human brain and expert systems is the closest we could go to it. Now, the technology in some way is offering new opportunities.