Artificial intelligence: what are the advances?

• 03/15/2016

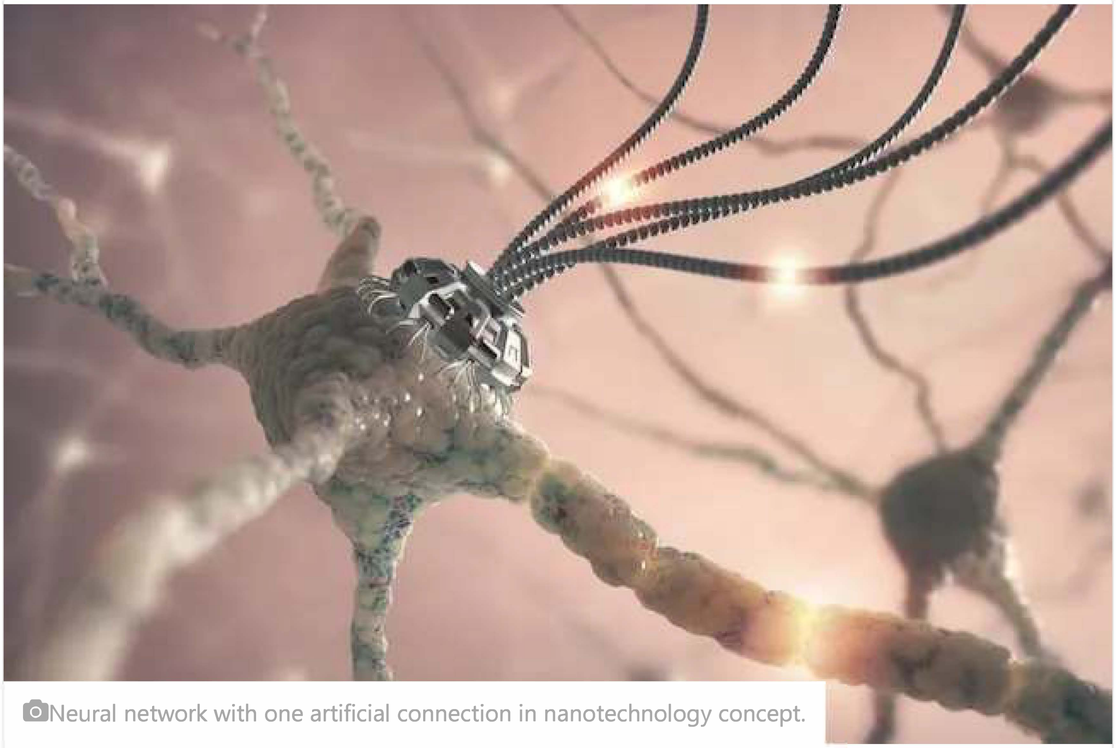

Neural network with one artificial connection in nanotechnology concept.

After a little introduction to what artificial intelligence is or is not, let's go to a more down to earth side by making a small approximate inventory of AI techniques. It is always about popularizing and restoring my process of

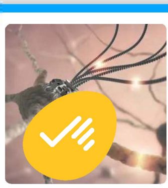

Artificial intelligence: what are the advances?

Read also : Advances in artificial intelligence

Winters in the spring of AI

The modern history of artificial intelligence started as we saw in the first part in 1957. This was followed by an important period of fundamental research, notably at MIT AI Lab, at the origin in particular of language **LISP** (1958) which has been used for two to three decades to develop AI software solutions. This research was funded by ARPA, the Pentagon research agency then became DARPA, the equivalent of French DGA, but obviously much better funded. Research on AI was mainly funded by public funds, notably in the USA and the United Kingdom.

The AI experienced its first " winter " with a reduction of a good part of these budgets from 1973, both in the United Kingdom and in the USA. This was the consequence of the publication of the Lighthill Report intended for the British public body **Science Research Council** – equivalent of the French National Research Agency – which called into question the merits of the research of the time in robotics and language processing. This winter lasted until 1980.

In question, somewhat overly optimistic promises from industry experts. As often, forecasts can be correct on all or part of the bottom but next to the to prove a new and important mathematical theorem. 30 years of error for the first forecast and as much for the second knowing that it is still largely in the making to be generic ! This pitfall is the same in current forecasts around singularity and transhumanism (the computer smarter than humans in 2030, immortality for our children, etc.).

IBM researcher Herbert Geernter had succeeded in 1958 in using geometry theorem demonstration software operating in rear chaining (from the solution to the problem) on an IBM 704 and from a base of 1000 rules. It was a rather simple combination. The same goes for the incompleteness theorem of **Godel** which says that “*in any recursively axiomatable, coherent theory capable of « formalizing arithmetic, one can construct an arithmetic statement which can neither be proven nor refuted in this theory* ” or the last theorem of **Fermat** (x ^ n + y ^ n = z ^ n impossible for an integer n > 2) which have never been demonstrated via AI.

Fermat's theorem was demonstrated in the mid-1990s and after years of effort by several mathematicians including Andrew Wiles. His demonstration published in the mathematics annals is 109 pages and uses many concepts ! A challenge was launched in 2005 by a certain Jan Bergstra to demonstrate Fermat's theorem with a computer and it is still to be noted. It's up to you to play if you feel like it !

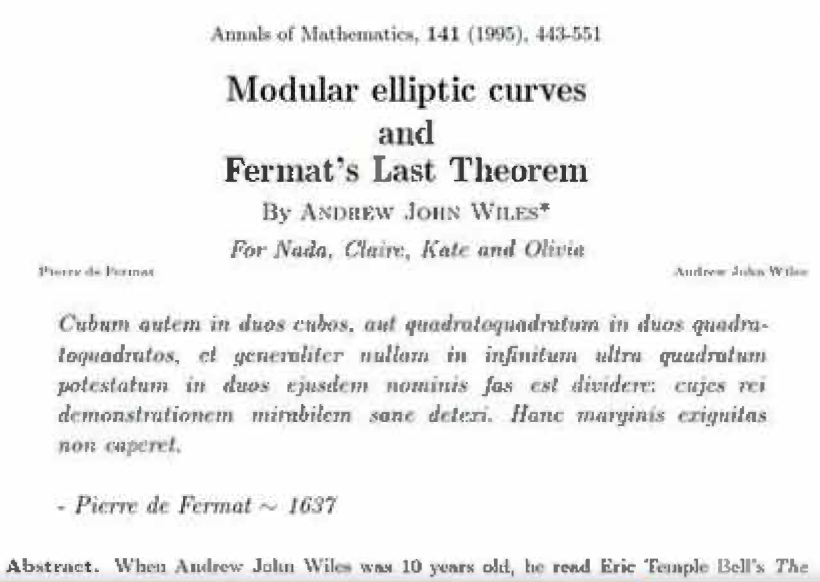

Annals of Mathematics, 141 (1995), 443-551

Modular elliptic curves
and
Fermat's Last Theorem

By Andrew John Wiles*

For Nada, Claire, Kate and Olivia

Pierre de Fermat

Andrew John Wiles

Cubum autem in duos cubos, aut quadratoquadratum in duos quadratoquadratos, et generaliter nullam in infinitum ultra quadratum potestatum in duos ejusdem nominis fas est dividere: cujes rei demonstrationem mirabilem sane detexi. Hanc marginis exiguitas non caperet.

- Pierre de Fermat ~ 1637

Abstract. When Andrew John Wiles was 10 years old, he read Eric Temple Bell's *The*

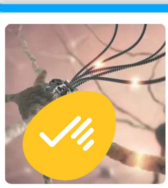

Artificial intelligence: what are the advances?

Herbert Simon also planned – still in 1958 – that in 1978, machines would be able to carry out all human intellectual activities. And Moore's law did not yet exist since it was stated afterwards, in 1965 ! In 1967 Marvin Minsky believed that in a generation, all AI-related problems would be resolved. Two generations later, we are still discussing it. He also predicted that in the mid-1970s, computers would have the intelligence of an average man. It remains to be seen what an average man is. The delays were also evident in machine translation and speech recognition. Note that Herbert Simon was awarded in 1978 with the Nobel Prize in economics for his work on rationalities of decision-making, after winning the famous Turing Medal in 1975.There is no Nobel Prize in Forecasting ! In general, they should be attributed to people who have already died !

This was followed by a period of enthusiasm in the early 1980s, fueled in particular by the wave of expert systems. Language **Prolog** French Alain Colmerauer contributed to this wave. A new wave of disillusionment followed around the 1990s. One reason was that the material could not keep up with the needs of AI, in particular to deal with two key needs: speech and image recognition, very greedy in computing power (see the source of the diagram below).

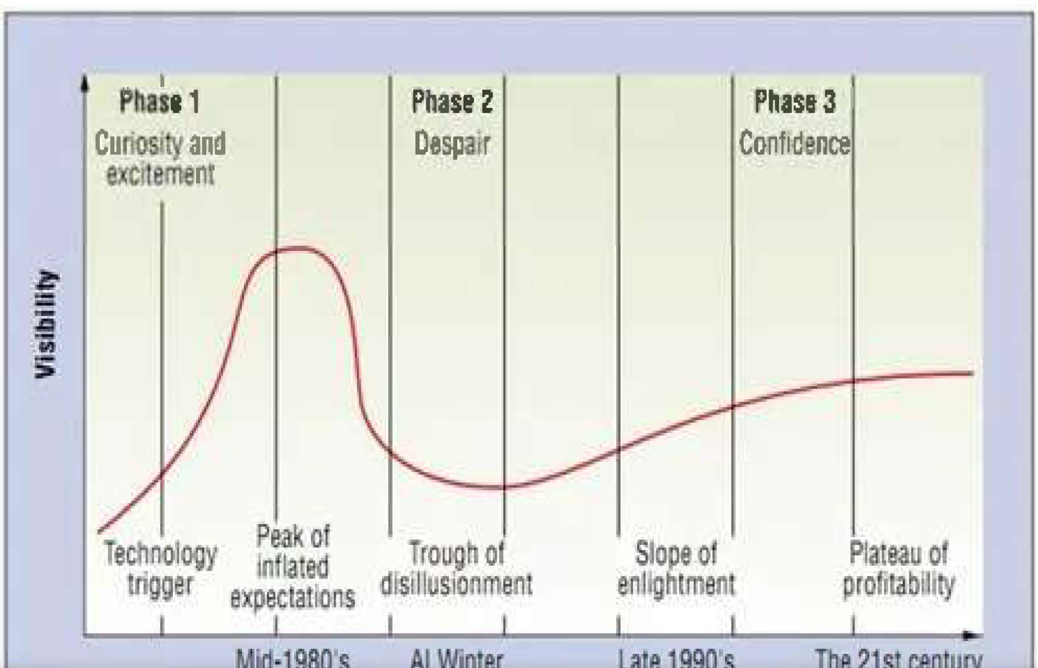

In the 1980s, various fifth generation " computer gosplans dedicated to IA applications were launched. It started with that of **Japanese MITI** launched in 1981 with spending of a billion dollars, then with the English project **Alvey** launched at £ 350 million and finally the Strategic Computing Initiative from DARPA. All of these projects went downhill and were discreetly completed. MITI's project aimed both to advance the state of the art on the hardware and software sides. They sought to process natural language, demonstrate theorems and even win the Go game. The project probably suffered from an overly traditional, linear and centralized organization. The late 1980s also saw the collapse of the computers dedicated to LISP language.

During the 1990s and 2000s, many projects emerged **HPC** (high-performance computing) very far from IA and focused on gross power and finite element calculations. They were and are still used for simulation, including nuclear weapons, air flows on aircraft wings or for weather forecasting. HPC from **Cray Computers** had been created for this ! This company exists always. She is one of the few survivors of the 1970s !

Since the early 2000s, the AI has been revived thanks to various developments :

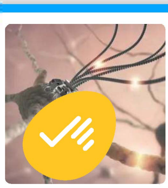

- The increase in **power of equipment** which has made it possible to diversify the implementation of many previously inaccessible methods. And in particular, the use of statistical methods that can exploit the power of machines as much on the calculation side as storage and then, more recently, neural networks.
- The attack of various **symbolic steps** as the victory of Deep Blue against Kasparov in 1997 then of IBM Watson in Jeopardy in 2011. Finally, a few days ago, DeepMind's victory at Go's game.
- L '**Internet** which created new needs such as search engines and enabled the establishment of massively distributed architectures.
- The availability of very **large volumes of data**, via the uses of the Internet and mobiles, connected objects or genomics, which makes it possible to combine brute force methods and neural networks and other learning machines or statistical methods.
- The **needs** in robotics, in the space conquest (Curiosity, Philae ...), in assisted or autonomous driving vehicles, in computer security, the fight against fraud and scams.
- The **many business applications** AI crossing the learning machine, connected objects, mobility and big data.
- The adoption of **methods** scientific and pragmatic – based on – and transdisciplinary experimentation, by researchers and industrialists.

Like any complex scientific field, AI has never been a field of unanimity and this may continue. Various schools of thought argue over the approaches to be adopted. We have seen the opposition of the supporters of connectionism – using the principle of neural networks and self-learning – in the face of those of computationism who prefer to use concepts of higher level without trying to solve them via biomimicry processes. We find this dichotomy in the battle between « neats » and « scuffies», the first, notably John McCarthy (Stanford), considering that the solutions to the problems should be elegant and square, and the second, notably Marvin Minsky (MIT) that intelligence works more

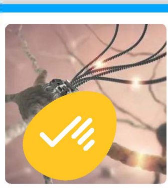

These debates have their equivalent in cognitive sciences, in the identification of innate and the acquired for language learning. Burrhus Frederic Skinner is behind the linguistic behaviorism that describes the conditioning operating in language learning. Noam Chomsky had questioned this approach by highlighting the innate, a kind of preconditioning of children's brains before birth that allows them to easily learn languages. Basically, the functioning of human intelligence is always the subject of scientific disagreements ! We continue, as we will see in the last article in this series, to discover about neurobiology and the functioning of the brain.

Other debates have taken place between declarative programming languages and inference engines using rules. Then came the statistical methods based in particular on Bayesian networks and optimization techniques. To date, the most commonly used methods are more mathematical and procedural, but methods based on neural networks and self-learning are gaining ground. The integrative artificial intelligence that is developing aims to jointly exploit all approaches.

Today, the debates are about the dangers of AI. AI is the subject of a social, philosophical, economic (debate on the future of employment) and therefore political. Debates tend to go too far from the scientific and technical sphere, to the point that sometimes we no longer know what we are talking about ! AI is a big thing or everything is put in the same bag. The AI is excessively anthropomorphized by imagining that it imitates, replaces and exceeds humans. This is one of the reasons for these papers to try to set a few clocks on time !

With that, I will now start from the lowest abstraction layers (expert systems, neural networks, machine learning, statistical methods, ...) and then go up to the higher layers which generally use the lower layers, as in the recognition of speech or images. For each of these techniques, I will discuss if necessary their seniority, the most recent advances, the flagship applications as well as some players in the corresponding markets.

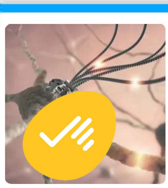

Brute force

Brute force is a common way to simulate or overcome human intelligence. For a game like chess, it aims to test all the possibilities and identify the most optimum paths among zillions of combinations. This can work if it is within the range of the computing power of the machines. These mechanisms can be optimized with pruning algorithms that evacuate " dead branches " from the combinatorics that cannot lead to any solution. It's easier to achieve in chess than in Go's game ! Brute force was used to win first with computer **Deeper Blue** IBM in 1997, calculating 200 million positions per second. Neural networks have been exploited to win Go recently with the solution created by **DeepMind**, the subsidiary in IA of Google.

Brute force is used in many areas such as search engines or password discovery. Many parts of AI can be considered to use it, even when they rely on modern neural network or machine learning techniques that we will deal with later. It only works if the combinatorics stay in the computer power envelope. If it is too high, methods of simplifying problems and reducing the combination are necessary.

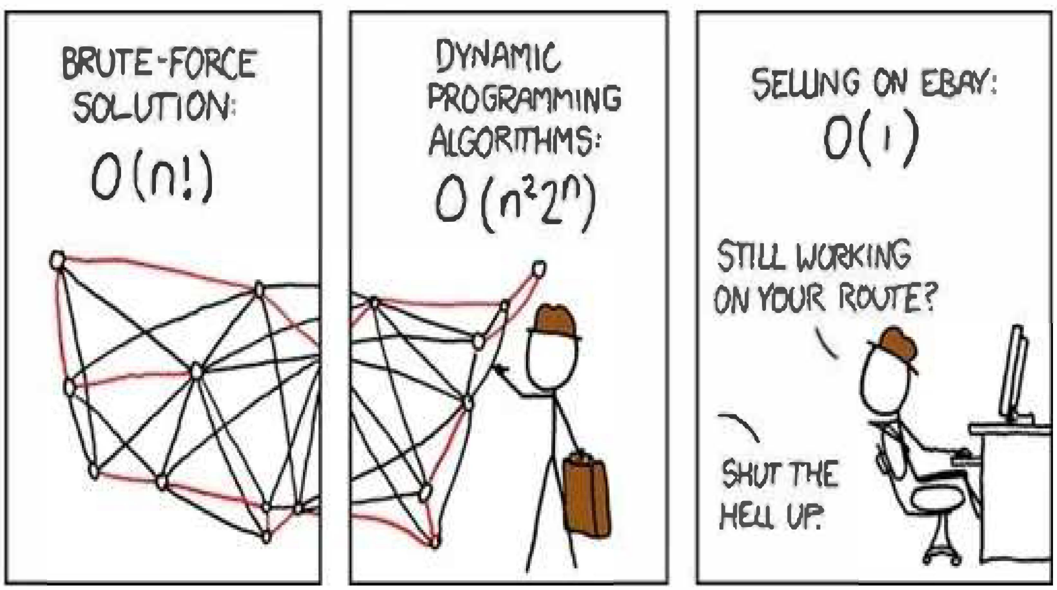

(image source)

The brute force has also become generalized because the power of computers from smartphone photo / video devices to sensors of various connected objects.

Rules engines and expert systems

The beginnings of the rule motors date back to 1957 when Alan Newell and Herbert Simon developed the General Problem Solver (GPS), problem solving software using rules modeling the possible inferences of a domain and solving a problem starting from the expected solution and going back to the hypotheses.

The drivers of rules are based on the concept of reasoning constrained by rules. The engine is provided with a set of rules which may, for example, represent the knowledge of experts in a given field. With rules close to logical programming of the genus " if X and Y are true, then Z is true " or " X entrain Y ". We can then question the system by asking it questions like " is W true " and it will manage to exploit the rules recorded to answer the question. Rule motors use graph theory and stress management.

This branch of AI was introduced by John McCarthy in 1958. In the 1970s, it culminated in the work of Robert Kowalski of the University of Edinburgh, Alain Colmerauer and Philippe Roussel who are at the origin of the programming language **Prolog** who experienced his hours of glory in the 1980s. The **LISP** has also been used in this area. A small industry has even developed with specialized computers from **Lisp Machines** and **Symbolics** (1979-2005), and software d '**Intellicorp** (created in 1980 and now specialized in application management software for SAP).

The drivers of rules are used in expert systems, a field and a market that has developed since the 1980s. Expert systems have been notably theorized within the framework of Stanford Heuristic Programming Project in 1980. They answer questions in specific areas of which knowledge has been codified. This allowed AI to make itself useful in specific areas, such as health. However, the approach encountered the difficulty of capturing the knowledge of experts. This explains

IBM Watson uses brute force to operate large volumes of unstructured knowledge databases.

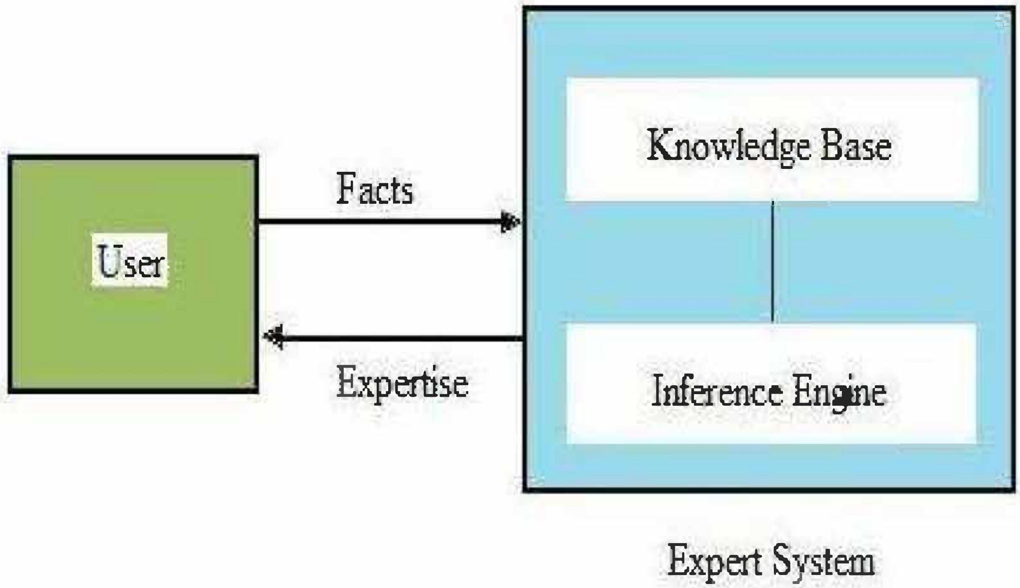

An expert system is based on two key components: a knowledge base, often generated manually or possibly by operating existing knowledge bases, then an inference engine, more or less generic, who will use the knowledge base to answer specific questions. Expert systems can explain the rational response. Traceability is possible up to the knowledge codified in the knowledge base.

There are still tools and languages in this area and in particular the French offer **ILOG**, acquired in 2009 by IBM and whose R&D laboratories are still in Gentilly near Paris. The ILOG Jrules inference engine has become IBM Operational Decision Manager. For its part, ILOG Solver is a C + + constraint programming library, which has become IBM ILOG CPLEX CP Optimizer. A less effective branding strategy than that of IBM Watson, as we will see in the next article in this series.

Statistical methods

Statistical methods, and in particular Bayesian, make it possible to predict the

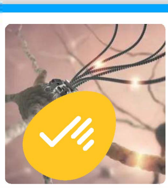

There are many applications such as detecting fraud potential in bank card transactions or analyzing the risk of incidents for policyholders. They are also widely used in search engines to the detriment of more formal methods, as Brian Bannon pointed out in 2009 in Unreasonable Effectiveness of Data.

Most scientific studies in the field of biology and health generate corpus in the form of statistical results such as Gaussian efficacy of new drugs. Exploiting the mass of these results is also a Bayesian approach. The brain also implements a Bayesian logic for its own decision-making, in particular motor, the associated centers being also located in the cerebellum while in the cerebral cortex manages the memory and the explicit actions (source : Stanislas Dehaene).

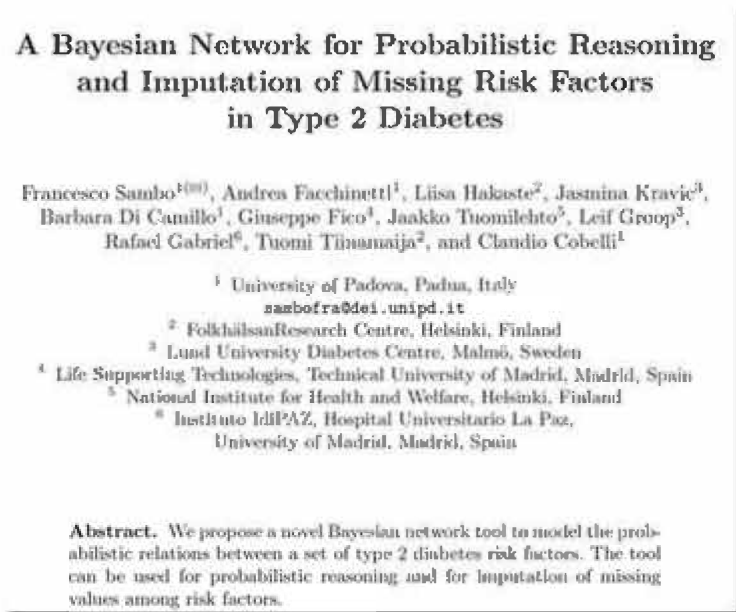
A Bayesian Network for Probabilistic Reasoning and Imputation of Missing Risk Factors in Type 2 Diabetes

Francesco Sambo[1(✉)], Andrea Facchinetti[1], Liisa Hakaste[2], Jasmina Kravic[3], Barbara Di Camillo[1], Giuseppe Fico[4], Jaakko Tuomilehto[5], Leif Groop[3], Rafael Gabriel[6], Tuomi Tiinamaija[2], and Claudio Cobelli[1]

[1] University of Padova, Padua, Italy
sambofra@dei.unipd.it
[2] FolkhälsanResearch Centre, Helsinki, Finland
[3] Lund University Diabetes Centre, Malmö, Sweden
[4] Life Supporting Technologies, Technical University of Madrid, Madrid, Spain
[5] National Institute for Health and Welfare, Helsinki, Finland
[6] Instituto IdiPAZ, Hospital Universitario La Paz, University of Madrid, Madrid, Spain

Abstract. We propose a novel Bayesian network tool to model the probabilistic relations between a set of type 2 diabetes risk factors. The tool can be used for probabilistic reasoning and for imputation of missing values among risk factors.

Fuzzy logic

Fuzzy logic is a concept of logic invented by the American Lofti Zadeh (" Fuzzy Logic ") in 1965. I had the opportunity to hear her present at a conference at the Central School in 1984, when I was on computer option in the third year. It doesn't rejuvenate us !

The fuzzy logic allows you to manipulate vague information that is neither true nor false, in addition to Boolean logic, but to be able to do operations on it

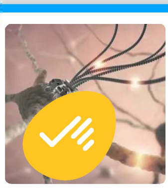

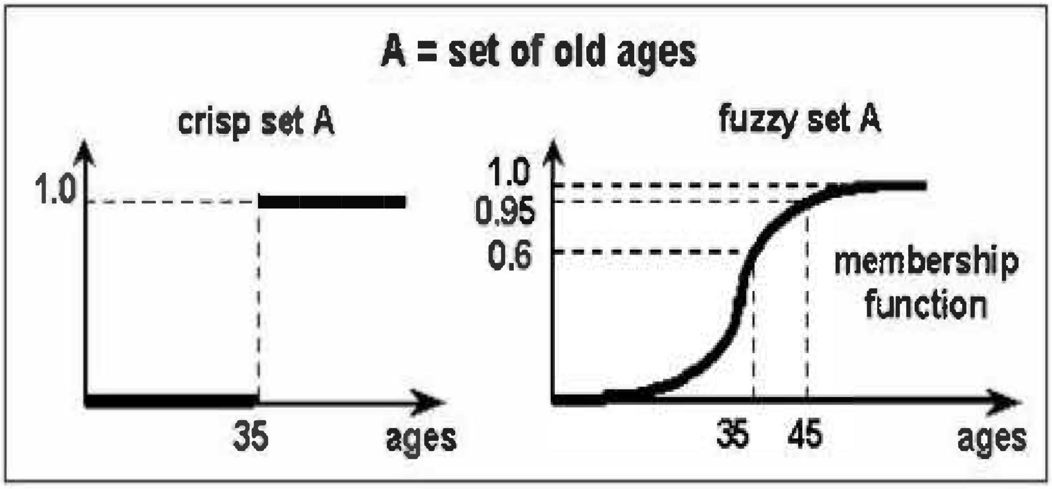

What about applications ? They are relatively rare. They are found in industrial control, in gearboxes at Volkswagen (to take into account the driver's intention " blur "), to manage traffic lights and maximize the flow, in speech and image recognition, most often, in addition to the Bayesian. Of the tens of thousands of patents would have been deposited to protect technical procedures using the theory of fuzzy logic.

Neural networks

The neural networks aim to reproduce approximately by bio mimicry the functioning of living neurons with hardware and software subsets capable of making calculations from some input data and generate an output result. Combined in large numbers, artificial neurons make it possible to create systems capable, for example, of recognizing shapes. The most interesting neural networks are those that can do self-learning.

The basic concept was born in 1943 from the works of Warren McCullochs and Walter Pitts. Donald Hebb added the principle of modulation the connection between neurons in 1949, allowing neurons to memorize experience. Knowledge is acquired through the interconnections between neurons and through a learning process. It is materialized in the form of weights of synaptic connections between neurons which vary according to the experience acquired, for example in the recognition of images.

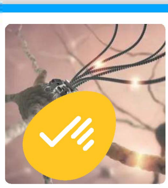

of **perceptron** in 1957 which was a fairly simple neuron in principle. The first perceptron was a single layer artificial neural network running in software form in a IBM 704, the manufacturer's first computer with magnetic torso memories. It was a linear classification tool using a single characteristic extractor.

In 1969 Marvin Minsky published with Seymour Papert the book " Perceptrons " which severely criticized the work of Frank Rosenblatt. Besides, on a very specific point relating to the logical doors XOR of the perceptrons. This book put an end to these developments, much like the Lightfill report a few years later. Always, in the dynamics of the rivalry of *neats* vs *scuffies*. This halt caused the whole of AI research to waste considerable time, especially since neural networks have since become, a fundamental part of the progress on all floors of AI. Marvin Minsky admitted his mistake in the 1980s, after the death of Frank Rosenblatt.

For the past twenty years, neural networks have been put to all sauces, the last being the victory of **DeepMind** against a Go champion in mid-March 2016. Neural networks have progressed step by step, with the creation of countless conceptual variants to improve their learning and memorization skills. The AI is also progressing regularly and in a rather decentralized manner, with dozens of researchers helping to advance the state of the art. The last years have however seen research efforts move from work in the basic logic to its applications.

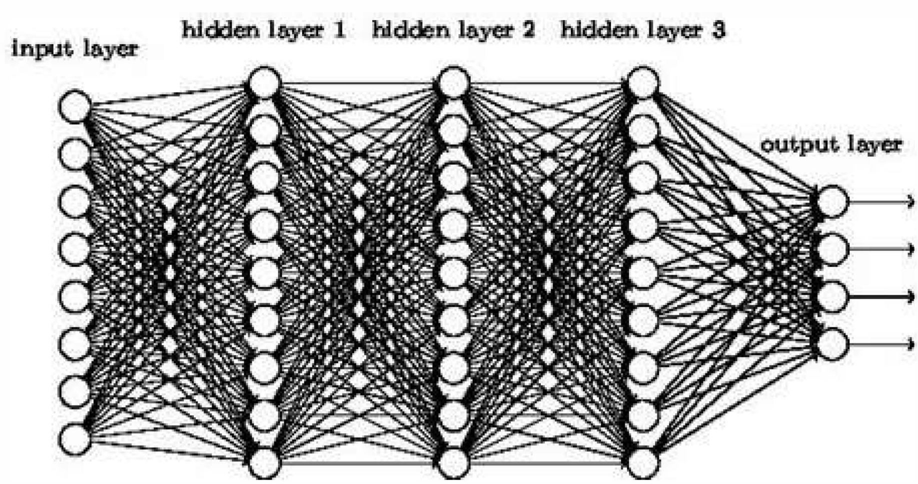

it possible to modify the weight of synaptic links between neurons according to the errors observed in the previous evaluations, for example in the recognition of images. How does this learning loop work ? It is an apprenticeship either assisted or automatic by comparing the results with the correct answer, already known. This is one of the key debates today: are we really able to create networks endowed with self-learning faculties ? It seems that we are still far from it.

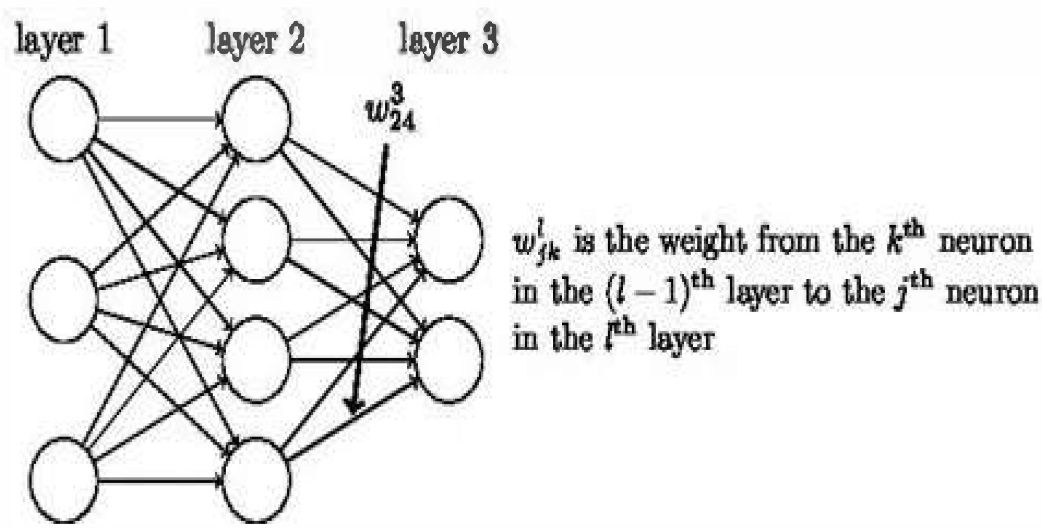

20 years after the revival of neural networks, in 2006, the Japanese Osamu Hasegawa created the autocontally organized neural networks (“ Self-Organizing Incremental Neural Network ” or SOINN), usable in self-replicable neural networks capable of self-learning. In 2011, his team developed a robot using these SOINNs capable of self-learning (video), masterfully illustrating the applications of neural networks. We are 10 years later, and we see that autonomous robots are still far from the mark, even if companies such as Boston Dynamics, a subsidiary of Google, scald us with robots that are very flexible in their approach and resistant to adversity.

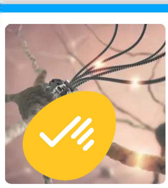

Neural networks have also progressed thanks to their implementation in specialized material architectures allowing them to properly parallelize their treatments as the brain does. The ideal electronic component for creating a neural network is capable of integrating a very large number of processing pico-units with inputs, outputs, calculation logic if possible programmable and non-volatile memory. In addition, the connections between neurons (synapses) must be as numerous as possible. In practice, connections are made with adjacent neurons in the circuits.

The **memristors** appeared in 2008 at HP after being conceptualized in 1971 by the Sino-American Leon Ong Chua. These are electronic components capable of memorizing a state by varying their electrical resistance by applying a voltage. Kind of like the bistable liquid crystals that serve in (fire) electronic scissors. The modifiable value of the resistance makes it possible to store information. Memristors can also be integrated alongside conventional active components in processing units. It's very well explained in Memristor: From Basics to Deployment by Saraju Mohanty, published in 2013, from which the two diagrams below are extracted. The second presents the different types of memristors currently explored. These components are integrated into silicon chips using more or less traditional manufacturing processes (nanoimprint

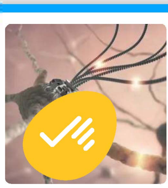

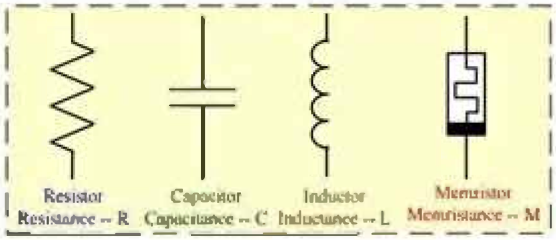

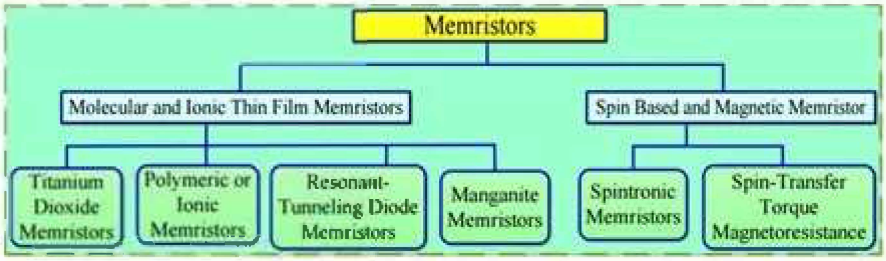

Memristors have been developed as part of program research projects **SyNAPSE** from DARPA. **HP** was the first to prototyp in 2008, with titanium oxide. There are several types, which can generally be produced in traditional CMOS chipset production lines, but with specific vacuum deposition processes for thin layers of semiconductor materials. HP even launched a partnership with the memory manufacturer **Hynix**, but the project was put on hold in 2012. The rebus rate would be too high during manufacturing. This is a key parameter for being able to manufacture components in industrial quantity and at an affordable selling price. In addition, the number of writing cycles seemed limited for chemical reasons, in the oxygen release / capture cycle for titanium oxide memristors.

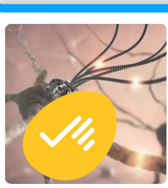

Artificial intelligence: what are the advances?

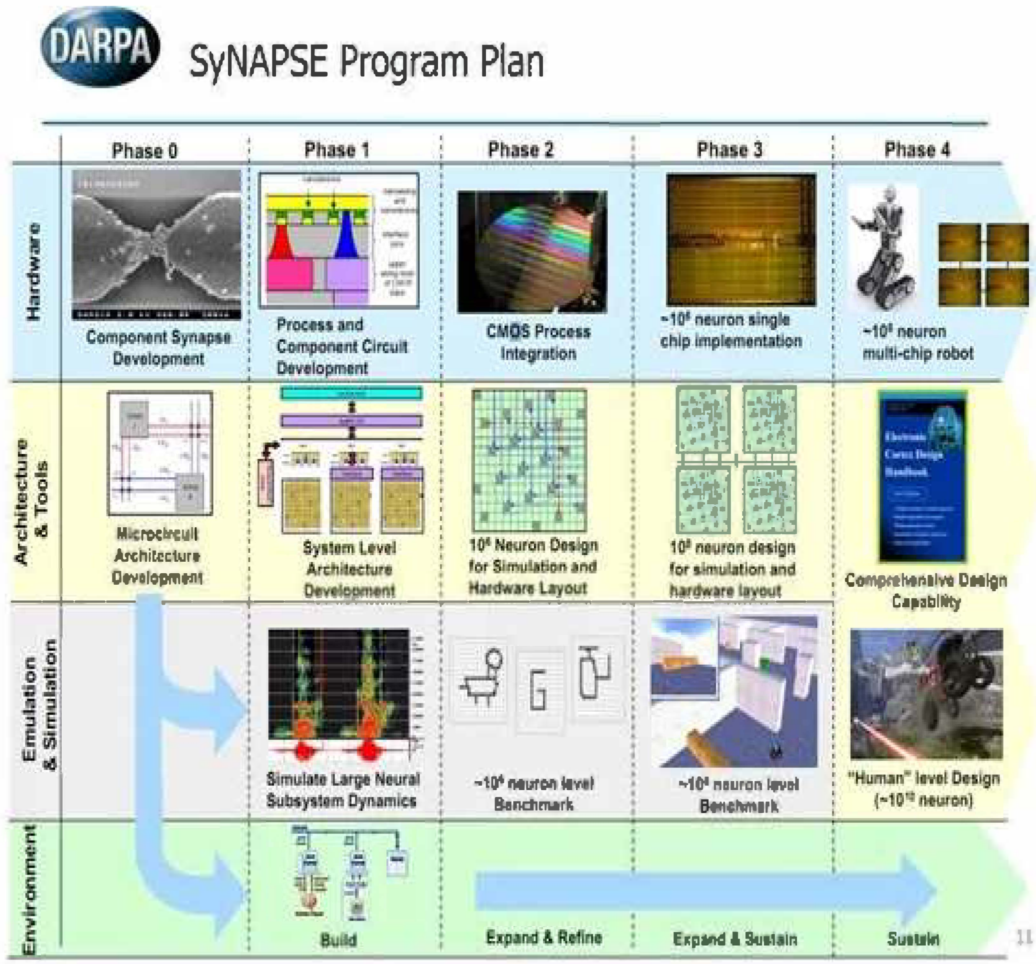

In October 2015, HP and **Sandisk** have, however, announced a partnership to make volatile, non-volatile memristor-based memories, which are said to be 1,000 times faster and more enduring than traditional flash memories.

Other research and industrial laboratories also work on memristores and networks of material neurons :

- **IBM** board with l '**ETH** de Zurich (the Swiss CNRS) on computers based on memristors. This same ETH develops a memristor capable of storing three states based on perovskite (calcium titanate) 5 nm thick. This could be used to manage fuzzy logic.
- Researchers at the Michigan Technological University have announced in early 2016 have created memristors based on molybdenum bisulfite which have a more linear behavior.
- Of the MIT researchers announced in early 2016 their work on the Eyeriss chipset using specialized neurons distributed in 168 hearts with their own memories. But obviously without memristors. The intended application is image recognition. The project is funded by DARPA.
- The project **Nanolitz** also funded by DARPA as part of Atoms to Product (A2P) projects and relies on microscopic wires to connect hearts and neurons more effectively in specialized circuits.
- ANR funded the collaborative project **MHANN** associating INRIA, IMS of Bordeaux and Thales to create ferric memristors. The project was to be completed in 2013 and had benefited from an envelope of 740 K €. Difficult to know what happened online.
- Finally, the California start-up **Knowm** launched the first commercial component based on memristors, manufactured in partnership with Boise State University, based on silver or copper and at the price of $ 220. It is primarily intended for research laboratories in neural networks.

The DARPA SyNAPSE program in any case led in 2014 to the creation by IBM of its neural processors **TrueNorth** capable of simulating a million artificial neurons, 256 million synapses connecting these neurons and performing 46 billion synaptic operations per second and per Watt consumed. All with 4096 hearts. The chipset was manufactured by Samsung in CMOS 28 nm technology and with a SOI insulation layer (from French SOITEC!) to reduce power consumption and speed up treatments. The chipsets includes 5.4 billion transistors in total and is more than 4 cm2 of surface. Most importantly, it consumes only 70 mW, which allows you to consider stacking these layered processors, something impossible with the usual CMOS processors that

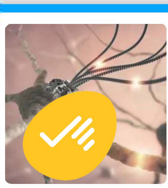

IBM's goal is to build a computer with 10 billion neurons and 100 trillion synapses, consuming 1 KW and holding in a volume of two liters. By way of comparison, a human brain contains around 85 billion neurons and consumes only 20 Watts ! The organic remains at this stage a very energy efficient machine !

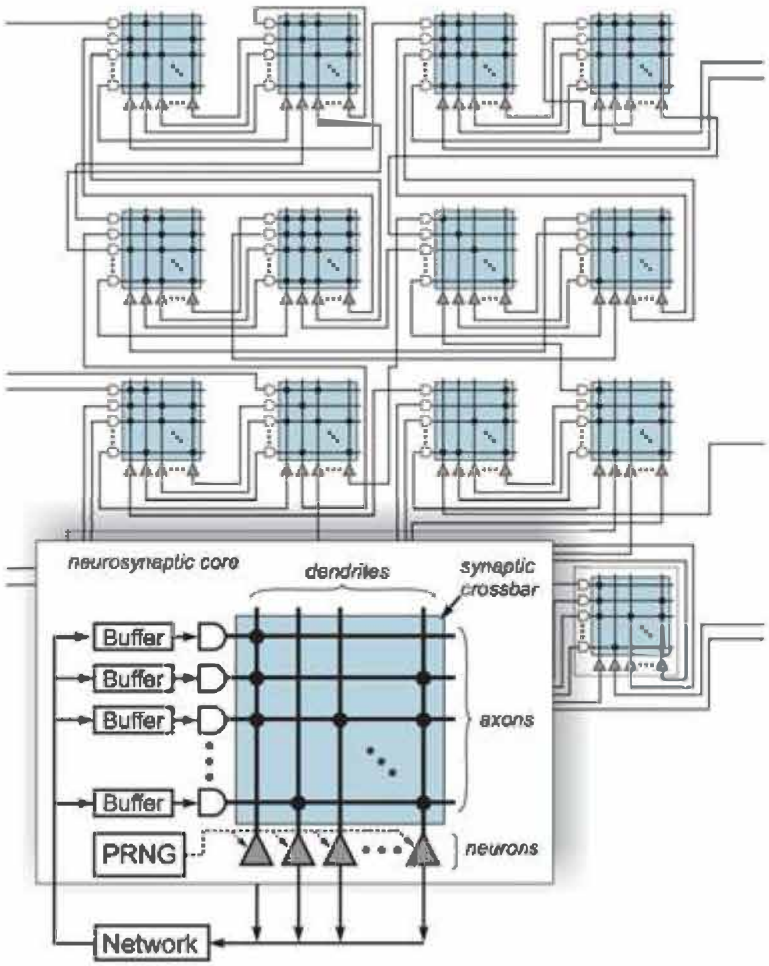

There are other synaptic computer projects based on neural networks. We can notably cite the project of Jeff Hawkins, the founder of Palm, that of Stanford, who works on the chipset **Neurocore** currently integrating 65,536 neurons and operating at very low consumption.

There is also the project SpiNNaker by Steve Furber (University of Manchester, UK), which aims to create a chipset of a billion neurons. However, it is based on a classic material architecture, with 18 hearts 32 bits ARM per chip. We are

Finally, in the commercial field, the **CogniMem** CM1K is an ASIC chipset incorporating a network of 1024 neurons which is used for pattern recognition applications. Only costing 94 dollars, it is used in particular in the **Braincard**, from a French start-up.

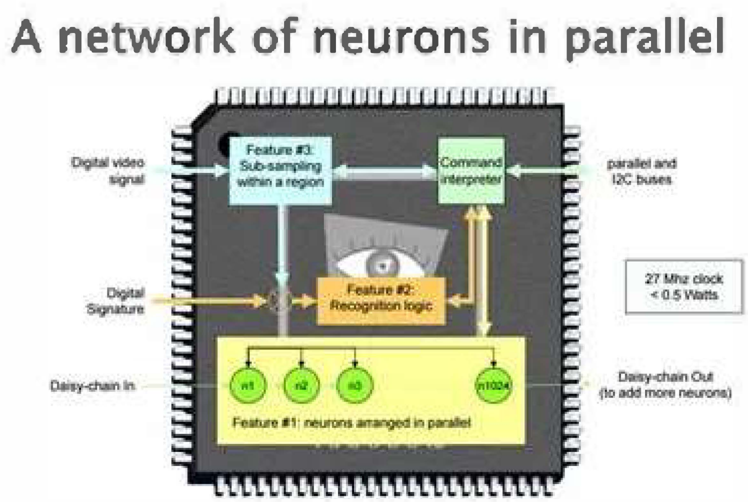

More recently, **Nvidia** presented to CES 2016 its PX2 card for the automobile, which includes two X1 processors comprising 256 GPU. GPU Nvidia is used to simulate neural networks. It's good but probably not as optimal as real networks of neurons and artificial synapses like IBM TrueNorth. What's more, the PX2 card must be refrigerated by water because it consumes more than 200 W. As explained Tim Dettmers, a GPU can only be used for neural networks if memory is easily shared between the hearts of GPU. This is what Nvidia offers with its architecture GPUDirect RDMA.

We can therefore see that all this is bubbling, rather at the level of research laboratories at this stage, and that industrialization will take a little longer, but that the material neural networks probably have a bright future ahead of them.

Learning and deep learning machine

The vast field of the learning machine, or automatic learning, aims to make predictions from existing data. This is an area that is intimately linked to that of neural networks, which serve as a substrate for treatments. Indeed, machine

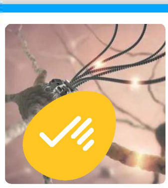

The learning machine is mainly used today for the recognition of shapes in images and that of speech, therefore in artificial directions. It can also be used to exploit unstructured data and manage knowledge bases. IBM lists some of these applications in its marketing. There are case studies in education to create self-adaptive MOOCs, in detail with a purchasing assistant, in health with the personalization of treatments against certain cancers or in the analysis of various data in smart city.

Neural networks experienced a revival in 2006 with the work of Canadians Geoffrey Hinton and Simon Osindero and Singaporean Yee-Whye Teh published in A Fast Learning Algorithm For Deep Belief Nets which optimize the functioning of multilayer neural networks. The concept of the learning machine was then formalized by Georges Hinton in 2007 in Learning multiple layers of representation. He relied himself on the work of French Yann Le Cun (in 1989) which now directs the AI research laboratory of Facebook and German Jürgen Schmidhuber (1992) of which two of the former students created the start-up **DeepMind** now a subsidiary of Google. Small world ! Geoffrey Hinton works for **Google** since 2013, not far from the legendary **Jeff dean**, arrived in 1999 and which now also works on deep learning.

To understand how deep learning works, you have to have a lot of time and a good mathematical and logical background ! We can start by browsing Deep Learning in Neural Networks of this Jürgen Schmidhuber, published in 2014 which makes 88 pages including 53 of bibliography or Neural Networks and Deep Learning, a free online book that outlines the principles of deep learning. He explains in particular why self-learning is difficult. Well, that's still more than 200 pages in body 11 and we're dropped to the fifth page, even with a good developer background ! There is also Deep Learning Methods and Applications published by Microsoft Research (197 pages) which starts by popularizing the subject fairly well. And then the excellent Artificial Intelligence A Modern Approach, a reference sum on the AI which makes the trifle of 1152 pages. I finally found this rather synthetic presentation A very brief overview of deep learning by Maarten Grachten in 22 slides ! Phew ! In short, you have to hit the equivalent of several Las Vegas CES Reports !

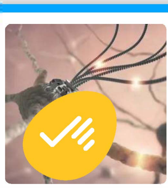

**Olivier Ezratty** is a consultant in new technologies and author of 'Free Opinions, a blog on digital media (digital TV, digital cinema, digital photo), and on entrepreneurship (innovation, marketing, public policies …). Olivier is an expert for FrenchWeb.

Photo credit:* *Fotolia, image bank, vectors and free rights videos

Experts

Experts are independent contributors to FrenchWeb.fr.

To communicate on FrenchWeb or the HR Journal, become a partner, click here.

#Olivier Ezratty **#free opinions**

WHAT ARE THE SERVICES TO STRENGTHEN