THERE IS NO SUCH THING AS ARTIFICIAL INTELLIGENCE

Luc JULIA
(in collaboration with Ondine Khayat)

THERE IS NO SUCH THING AS ARTIFICIAL INTELLIGENCE

Foreword by Jean-Louis Gassée

FIRST ÉDITIONS

Table of Contents

Part Two
The Misnomer

The Dartmouth Conference

It all started with a misnomer. The genesis of artificial intelligence can be traced to 1956 during the famous Dartmouth conference. Scientists who were studying the automata theory introduced by Alan Turing some 20 years earlier were beginning to think that it would be possible to re-create the human brain inside these machines and mechanisms. The objectives of the conference were extremely ambitious and were presented as follows:

> "We propose that a 2-month 10-man study on artificial intelligence be conducted in 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 the 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."

When the term "artificial intelligence" was first used, it was credited to John McCarthy, one of the organizers of the conference and signatory of the program, and it was accepted by everyone. The results turned out to be quite far from the working group's ambitions. They made progress in what would become the field of expert systems but failed to define the steps needed to simulate intelligence. The use of the word "intelligence" for this discipline is in fact a sham because it is based on what one wishes computers could do, rather than what they can actually do. Attempts to copy brain functions were in fact just beginning. In 1957, Frank Rosenblatt invented a learning algorithm called the perceptron which was purported to simulate closely neuronal function. This invention created a lot of excitement around the idea of neural networks, which are still the basis of Machine Learning (ML) today. The initial purpose of the perceptron was to classify images, but it can be generalized to all types of senses (sight, touch, smell, etc.). One can give weights to the various values supplied by the sensors, then define a function linking all the parameters, and finally use the function to classify the results. This was a method for supervised learning. One would give the system several known examples and the associated correct classification conclusion from which a function would be created. Once the system has been trained, an input signal is processed by the function which creates an output that is the probability that the input belongs to one of the defined classes. As a simple example, we can define a parameter such as "number of paws", assuming that we have a sensor that counts paws, and another such as "has scales", assuming that we have a sensor to detect scales. During the learning phase, we give the machine dogs and snakes as input, we define the classes that we are interested in for the output, and the function will converge toward what we call a model. When we subsequently give a particular dog as input, the system will probably correctly classify it as a dog, and if we give it a type of snake, it will probably identify it as a snake. But if we give it a crocodile, the probability that it classifies it as a snake (or a dog) is close to 50%. This is a straightforward case that can be handled with a rule-based system, also called an expert system, because the decision tree is simple. But neural networks include statistics so that ambiguities can be handled better, which means that the machine is working using a system that is not based on just true/false or 1/0 as machines were previously confined to do. On the other hand, this doesn't mean that machines are intelligent, because all they are doing is parroting what they have been taught in a particular domain. The perceptron involved tens of neurons, while the human brain has about one hundred billion neurons. So from the beginning, the results were off. The perceptron and neural networks aren't at all like our brains. They are nice abstractions but we don't even know today how the brain works exactly. This simplification has led many people to believe that there was some sort of equivalence between this technology and intelligence. People got very excited about this breakthrough technology. But this excitement was based on wishful thinking and it wasn't long before the bubble burst.

From time to time, people working on artificial intelligence, especially those who were waiting to use the technology, would get discouraged with the slow progress. It was during that period that a mysterious "hibernation" episode occurred. In its short lifetime, AI had already lived through two winters, a phenomenon known as the "AI Winter" (a time of reduced interest and funding in AI). The first occurred in the early 1970s, after more than a decade of AI failure to solve problems in the areas of language, automatic translation and representation of complex systems using simplistic neural networks. The second episode took place in the late 1980s. AI had made a comeback thanks to the personal computer boom, and the hope was that with increased computing power, the creation of an intelligent system using complex neural networks would be possible. The systems could have more parameters, more layers in the network so that the results would be more refined. In addition, with the arrival of ubiquitous computing, a new Japanese initiative called the "5th generation of computer systems" was tasked with supplying a new AI platform. But once again, the results were slow and a new AI winter began. The key element that was missing emerged in the middle-1990s: Internet and its massive amount of data.

Computing capacity on the Internet is a very important element for AI, but data is essential to all learning algorithms. Internet quickly became the largest database in the world as well as a platform for exchanging information. This enabled the emergence of "big data" and along with it, a renewed interest in AI. For example, to get a machine to recognize a cat with 95% accuracy, we need something like 100,000 cat images. This is a lot, in fact a lot more than what a human being needs to recognize a cat, which proves that our brains use methods that are very different than the ones we program into our machines. Machines employ methods similar to brute force techniques, while human beings use their innate intelligence. Psychologists tell us that children only need two images of cats to be able to recognize cats for the rest of their lives, under any circumstances and with nearly perfect accuracy. Machines are incapable of understanding context. If the machine has not been shown the image of a cat at night during the learning phase, there is almost no chance that given a picture of a cat at night, it will identify it as a cat. We can of course increase the number of parameters and give the algorithm more data during the learning phase, but aside from the fact that we just can't model every possible circumstance, we would eventually run up against limits in computer memory and calculating power.

The democratization of computers and the fundamental role of the Internet in allowing access to more sources and a greater volume of data enabled the resurgence of neural networks. Another very important element linked to computing power appeared at about the same time. Until 2001, CPUs (central processing units) were mainly used to run algorithms on computers. CPUs had up to twelve computing cores to carry out rapid calculations. In the 1970s, co-processors began to be used for tasks such as sending information to the screen – simple, but a huge consumer of computing resources since all the information must be displayed simultaneously. By doing this computation in parallel on another processing unit, a considerable load was taken off the CPU. GPUs (graphics processing units) were developed to manage screens with more and more pixels, more and more colors, and more features (e.g. 3D). They became what is known as massively parallel (many computations being performed at the same time), with thousands of highly specialized core computing units. In 2001, scientists began to use GPUs for matrix calculations (unrelated to displaying information on a screen), but there was no easy way to access the GPU resources. It wasn't until 2006 that Nvidia, the leading GPU manufacturer, made a library called CUDA for programmers to access these functions. The neural network community, a heavy user of highly parallel matrix computations, soon realized that using this architecture would lead to enormous performance gains for computations on their increasingly complex networks. This opened the door to a new type of learning, "Deep Learning" (DL). Adapting hardware to suit ML and DL algorithms had just begun. In 2016, Google came out with a new type of processor called a TPU (Tensor Processing Unit) which was specialized in computations for their automatic learning tool TensorFlow. Thus, since 2007, many new calculation methods have emerged, calculation capacity has grown considerably, and this has renewed hope for AI and its unrealistic promises. But computing is still nothing more than computing. Alarmists who are predicting gloom and doom and saying that robots will take over the world and dominate us, or people who say that artificial intelligence will solve all our problems, are talking nonsense. Furthermore, they are damaging the field and may well contribute to the start of the 3rd AI winter that threatens us today. Because of them, there is a real risk of arresting research and progress in ML and DL precisely when the period is actually productive – though we are still in the early stages of this work.

And all of that because of a misnomer, a term used to name a discipline that has nothing to do with intelligence. I argue that there is no such thing as artificial intelligence. The AI acronym should stand for Augmented Intelligence. This is the term

I will use for the rest of the book, and I will explain why. But let's take a moment to review a little history…

Reality: a brief historical review

The myth of creating intelligence is part of a historical context fueled by popular fantasies about creatures turning against their creators, and alarmist predictions that robots will take our jobs. It is also a time when we have seen IBM, Facebook and others employ their own AI programs. Since ancient times, stories and rumors have been created about artificial beings endowed with intelligence or consciousness. It all began with the desire to imitate or become God.

- The Golem: myths typically speak about the depth of reality through tales symbolizing energies and aspects of the human condition. The myth of the Golem is the story of an artificial humanoid creature who is unable to speak and has no free will. It was made by a human being who created it in his own image to defend and assist him but Golem becomes human and can no longer be controlled; its creator must fight it to regain power. It is the myth of the stone monster that comes to life and becomes intelligent. It is an anxiety-inducing representation, and like Pygmalion and Frankenstein, one of the most important myths associated with the world of artificial creatures. The myth of the Golem continues to feed the fears and fantasies of those who are afraid of devious robots, and to nourish the concerns currently being expressed about this pseudo "super intelligence". The idea that the robots created by humans will take control and lead to our extinction is not based on tangible facts. It is a tactic to induce fear, which can make an entertaining movie but has no scientific foundation.

- Pascal and Babbage: Pascal's calculating machine, the Pascaline, is the foundation of calculation and its automation. Pascal was looking for a way to imitate the way human beings make calculations in order to help his father who was a merchant. In 1642, he invented this machine based on gears and simple arithmetic rules. It could only do addition and subtraction but it did them much faster than people did and with no errors. It was in effect the very first arithmetic machine, the very first computer. By mechanizing the process of carrying out a calculation, Pascal laid the foundations for computers. It wasn't until the 1830s, two centuries later, when the British mathematician and inventor Charles Babbage, famous for developing plans for two different computers (inspired by Pascal's machine), made great progress towards modern computing. His first invention, the difference machine, was only partially completed in the early 1830s. His second invention, the more complex analysis machine, was never built at all. The two machines, often considered to be the first two computers in history, had great potential for the time. The difference machine could do simple calculations such as additions and multiplication, but its most important feature was its ability to solve polynomial functions up to degree seven. This machine was nearly finished in 1832 (with the exception of the printing mechanism), but funding ran out and it was abandoned. So it was never proven to work as was the Pascaline machine. In 1837, Babbage conceived the analysis machine. More than a calculator, it was the first machine that could be programmed. Unfortunately, once again, because of the political and economic environment, but also because of technical difficulties, it was never implemented and the project was abandoned. Babbage's ideas were well ahead of his time, and until the next century, there was no further attempt to create a computer. Babbage is often called the "uncle" of computers because of his vision, but for me, Pascal is the "grandfather of computers" because of his famous Pascaline.

- Automata: these are mechanisms that reproduce human movements. They were created to make the physical work of humans easier by imitating and improving their movements. One of the famous automata is The Turk which was supposedly able to play chess. In the 1770s, Baron Wolfgang von Kempelen invented a human-like automaton that looked like a turk with a turban and kaftan. It sat on a chair attached to a cabinet with doors. There was a chess game in front of the automaton and it was able to move the pieces around. When the doors of the cabinet were opened, one could see a complex system of mechanical parts and gears that sprung to life when the automaton made a move. But what couldn't be seen was a second hidden compartment in the cabinet, where a person was making everything move correctly like a puppeteer. The Turk automaton played hundreds of games against people, and usually won, including against Napoléon Bonaparte and Benjamin Franklin. This deceptive device remained in action for more than fifty years and led people to believe that intelligent machines were possible.

- Automation in factories: the 1960s were marked by the automation of industrial production thanks to advances in electronics and information technology. Despite attempts described in the 1st century by the Greek mathematician Heron of Alexandria in his work Pneumatica, the first industrial robots date back to the 18th century. It all started in the Lyon region in France, renowned for its centuries-old silk industry. In 1725, Basile Bouchon invented a semi-automatic loom using a perforated paper tape to program the operations and help weavers in their repetitive task. Three years later, his assistant Jean-Baptiste Falcon perfected the invention by replacing the tape with perforated cardboard cards linked together to form an infinite chain. In 1944, IBM borrowed the idea of perforated cards to enter instructions into the Harvard Mark I, considered to be one of the first modern computers. But history books tend to remember only Joseph Marie Jacquard who used the idea to create the famous "Jacquard loom" in 1801. The installation of these machines resulted in the no less famous rebellion of the "canuts" (silk workers) from 1831 to 1848. They saw the arrival of these machines as eventually leading to their unemployment. Industrial robots replaced human beings little by little, even though they were originally merely intended to help them complete difficult tasks. Some jobs were indeed lost but on the positive side, robots have helped reduce strenuous work.

Many questions are being asked about the role of AI in the job market. I am quite surprised that the OECD (Organization for Economic Co-operation and Development) has predicted an 8% job loss. This figure seems absurd to me. In fact, in the

1960s, when robots were introduced into Renault car factories, many of the skilled workers who lost their jobs were quickly moved to other positions. We can improve society with technology but we need to understand it and know how to use it to our advantage. Technology will undoubtedly reduce the amount of time needed for certain jobs (thus causing job loss) but in return, it will allow us to devote more time to jobs that are more interesting and rewarding. The real questions to ask about technology have to do with politics and distribution of wealth. Are the people in Silicon Valley who understand and create technology going to get richer while other people stagnate, or are the riches produced going to benefit everyone? Technology creates jobs which are usually well payed, so the question is how to distribute this wealth. This issue has been around forever but I believe we have come to the point where we can no longer put off finding an adequate answer. We often hear that robots will take our jobs, yet the countries that have the most automation are also those with the least unemployment. In Japan, there are hundreds of thousands of robots. In Europe, Germany has the lowest unemployment rate but the largest number of robots. In a recent study, the consulting firm McKinsey estimated that in 2030, about 15% of the tasks we perform will be automated, with a large discrepancy between countries. They predict 9% in India, 24% in the U.S. and up to 29% in Japan. This automation boom shouldn't be seen as a loss of employment but rather as a transformation of employment, meaning changing jobs or changing field for 375 million workers (14% of the working population). This proportion is expected to be more like 33% for the U.S. and Germany, and 50% for Japan. Taxing the production and use of robots is risky because it is like taxing innovation, and innovation should be supported. In the coming years, drudgery will be reduced, new horizons will open and perhaps we will even have more time to do something other than work.

- Deep Blue vs Kasparov: in the 1990s, IBM developed Deep Blue, a computer that specialized in playing chess. In 1997, it beat the reigning World Chess Champion, Garry Kasparov. Chess is a complicated game, and people who are good at it are known to be high performers, very intelligent and endowed with the capacity to envision the game several moves in advance. According to Game Theory, chess is a strategy game that is combinatorial without cycles and with complete and perfect information, which means that the rules are clear and define the end. Whatever the path taken, the result is always that we have a winner or a tie in the end. The Deep Blue program was just a set of rules that could beat Kasparov, not by "thinking", but by having a phenomenal capacity of memory where thousands of complete chess games were stored, together with the different paths to victory for each configuration. There was no intelligence but a remarkable technology that was capable of finding the right move at the right time in all that data. To give you an idea of the size of the problem, the set of possible legal chess moves is estimated at between 10^{43} and 10^{50}. The machine didn't know all the moves but was capable of anticipating more moves than Kasparov could. Thus the machine defeated Kasparov thanks to "big data" but certainly not because of its reasoning capabilities.

In this case, it wasn't even Machine Learning that was used but just brute force.

- Watson and Jeopardy! DeepQA: created between the time of Deep Blue and DeepMind, DeepQA is a huge database of questions and answers developed by IBM in the late 2000s. Jeopardy! is a televised game show in which contestants are given the answers and must come up with the questions. The IBM engineers thought that participating in this game would be a good test for DeepQA. So they created Watson, a specialized computer capable of answering questions in a natural language (that is, a spoken as opposed to a programming language). Contrary to what one might think, the name Watson had nothing to do with Dr. Watson, Sherlock Holmes' faithful assistant, but was the name of the first IBM CEO, Thomas Watson. In 2011, Watson participated in the game show against legendary Jeopardy! champions and resoundingly defeated them. Once again, it had nothing to do with intelligence but rather a huge memory capacity combined with enormous calculation power. Watson could handle 500 gigabytes of data (the equivalent of a million books) per second! It's very impressive, but it is nothing more than knowledge processing.

- AlphaGo vs Lee Sedol: DeepMind, which Google acquired in 2014, is a British company specializing in artificial intelligence and focusing in particular on developing computer systems that can play video games. It made headlines in 2016 when its program AlphaGo beat the Go world champion, a South Korean named Lee Sedol. Go is much more complex than chess. It is very difficult to estimate the number of moves in chess, and much harder and less accurate in Go, but current estimates range from 10^{172} to 10^{762}. The real number lies somewhere in between, but in any case, it is clear that there is not enough computing power to apply a brute force technique. Here, Machine Learning can do very well compared to a human being. Indeed, by feeding the program 30 million moves extracted from 160,000 played games during the learning phase, and by making two instances of the program play against each other to perfect the model, it was possible to produce a result that far outperformed a human – one that even won 18 world championships. It is important to note that since AlphaGo needed more computing capacity than memory, it had some 1500 CPUs, 200-300 GPUs and a few TPUs. AlphaGo defeated its human opponent by using methods and strategies that were not at all human, but by linking statistical techniques from Machine Learning and Deep Learning. It would be ridiculous to shelve this powerful technology just because people are disenchanted or fearful of these technologies.

- Microsoft's racist chatbot, Tay: in March 2016, Microsoft launched a chatbot that users could talk to freely on Twitter. It was supposed to personify a 19-year-old American girl. During these conversations, the chatbot started to make racist and sexist comments and Microsoft had to shut it down only 16 hours after it was put into service. Microsoft has never commented on the reasons for this disaster, but multiple hypotheses have since been formulated. Chatbots are based on technologies from classic Machine Learning. They are trained on large databases containing conversations chosen

41204340_001-274_enc.indd 133 19/06/2019 11:24

because their topics corresponded to the bot's specialty. In order for the bot to converse normally and appear to be like its interlocutors, it is common to implement the "repeat after me" function in the bot. This function enriches the database of the bot by incorporating current conversations in real time. It is not easy to find databases of transcribed, annotated conversations. But they do exist and often come from transcriptions of conversations recorded with clients during a customer service phone call. One hypothesis is that Microsoft trained its chatbot with conversations held in the South in the 1950s, which were notoriously racist. In that case, it is easy to understand how the system could have been biased in this way from the beginning. Similarly, knowing the propensity of Twitter users to hammer out insulting tweets shamelessly, if the "repeat after me" function was given too much weight in the algorithm, it is obvious that Tay would start to spew hateful messages by imitating other users. Since learning algorithms are generic, the case of the Microsoft chatbot highlights the risk of using biased content and thus the importance of carefully choosing data sources. This example shows once again that there is no intelligence, reflection or critical thinking on the part of the machine and that the responsibility for any failings is that of the human designers. In this particular case, the errors were certainly unintentional, but one can easily imagine someone manipulating data and systems for devious purposes. Why not create a bot with a database that stipulates that white people are nice and black people are mean? The biased bot will then become the source of "fake news" popularized by Donald Trump. The only way to combat this is to continue putting resources into education to make every person capable of critical thinking. If we take images reflecting the frequency of occurrence of objects in photos on Google or Facebook in the American population, we will get about 10% people with rather dark skin and 90% people with lighter skin. If we train an algorithm to perform face recognition without taking this into account, we will end up with one that works far better for people with light skin than for people with dark skin. This has nothing to do with ethnic origin and simply reflects the fact that we have used more data for one category than the other. To correct this, the proportion of images used to train the system must be modified in an intelligent way, taking biases into account. This is a crucial point, and one that has become a subject of research in its own right, which is not entirely related to automatic learning or AI. These are problems that should have been thought about more carefully in the last ten or twenty years. Today the major players in this sector take it very seriously and large companies such as Facebook, Google and Microsoft have created research and development groups to treat the specific problem of bias. This awareness has lead to the creation of several movements, such as the "Partnership on AI", which is an open forum for exchanging ideas about these issues. It is funded by several tech companies and its board of directors is composed of university professors, representatives from non-governmental organizations, and other important players. In terms of regulation, people deliberately using sources that are racist, sexist, etc., should be severely punished by law. Strict guidelines will have to be set up and observed before a bot can be used by the public. On the other hand, the laws must not stipulate which sources to use. Researchers must be left to work on their own, provided they respect the law. If there is too much regulation, we will enter a dictatorial world where we are told which source to use. Regulation will become abusive, as is the case today in certain racist countries. It is also true that following too many rules cripples innovation. In the U.S. in the 1950s, regulation practically didn't exist so that researchers were able to make rapid progress.

The real danger of AI and robots comes from us, humans. Anyone could decide to deliberately create and program killer robots or introduce racist chatbots. The solution is regulation, with the risk of jeopardizing innovation. A subtle balance should therefore be found. We must be able to explain to the public what we are doing in this domain. This question of explainability is another important subject and I strongly disagree with people who say that decisions made by deep neural networks can't be explained. We know perfectly well how our systems function and how to change the variables to modify the resulting decisions. There are about 200 trillion decisions made each day on Facebook. It is such a huge number that we can't supply explanations for every one of them, and in any case, no one would take the time to look at them. Most of the decisions are not important – for example the order in which articles appear in your newsfeed. On the other hand, there is more and more automation in these decisions, and they have an important impact on people's lives. As an example, we can examine the decision to grant a client a loan. There are laws prohibiting the use of ethnic origin as a criterion in such decisions. So, we will build a learning system that doesn't take this factor into account. But it is not as simple as that, because other variables are correlated with ethnicity. In the U.S., for example, ethnic origin is often correlated with home address, which must obviously be provided when making a loan application. How can the system be built so that the decision is not biased? These are very important questions. As far as loans go, it is possible to determine algorithmically what the minimum change in the variables must be for the decision to change. If the decision is to deny the loan, one could, for example, determine that the variable responsible for the decision is linked to revenue or current debt. One can then identify and communicate the operational information that will influence these decisions.

- Facebook chatbot negotiators: in the summer of 2017, Facebook decided to put two chatbots face-to-face and get them to negotiate with each other. In the days and weeks that followed, the scoop-hungry press reported that the experiment had to be stopped almost immediately because the two bots had invented a new language and become uncontrollable. The reality is much less exciting. The two bots had been trained using different objectives and different databases, so they quickly lost interest in each other as neither one could complete the negotiation task they had been assigned. They certainly tried to talk but we are almost sure that they didn't understand each other. A robot that has been given a specific task cannot create another task on its own, and certainly cannot create a new language. It can probably adapt itself to its environment to some extent depending on how it was programmed, and perhaps simplify or enrich its vocabulary based on new data supplied as we saw with Tay. In this particular case, we can speculate that the robots mixed their databases (which were rather small, in the order of 5000 sentences) to create something that was incomprehensible by either party. The fact that they "invented" a language is one of the technology legends that we like to believe in and peddle, but it is really the type of story that does a disservice to the discipline.

- Tesla's Autopilot: there is a lot of talk right now about autonomous cars – vehicles that are capable of driving without intervention from the driver. We are not quite there yet, but there has been considerable progress. In October 2015, Tesla deployed a software update for its Autopilot system (version 7.0) on its Model S cars, which allowed them to drive almost autonomously on highways and to enter and exit a garage on their own. Despite recent high-profile accidents, I am convinced that these cars represent a very important step that will eventually save tens of thousands of human lives. I personally would have a much easier time trusting my children to an autonomous car than allowing them to drive themselves. The first fatal accident in May 2016 happened because of a limitation in the vision system. But it would be interesting to know how many lives have been saved thanks to the driving assistance features of these cars, for example the emergency braking system. There are several levels of assisted driving, from 1 to 5, where 5 is the level at which the car is fully autonomous. Tesla cars are at level 2. The technologies used are essentially a mix of sensors and guidance systems. The sensors are radars, lidars which are similar to radars but use lasers instead of electromagnetic waves, and cameras. The sensors are used to see and understand the immediate surrounding environment and to react to dangers or situations that are not taken into account by the GPS guidance system (which only gives a general idea of where to go). These vision and detection systems are potentially much more precise than human senses, and the reaction times of the control devices are much faster than we are. All this means that systems like Autopilot are much more reliable drivers than we are – or will be! The system will never drive under the influence of alcohol, will never exceed the speed limit and will never send text messages while driving since it is programmed only to execute the task of driving. But that's where the potential problem lies – in the programming. The car does not have the intelligence to react to situations that have not been programmed. It has been programmed with all the rules of the road and has been trained using extensive databases with photos of roads, vehicles and pedestrians. But that may not be enough. When there was an accident with an Uber car during which all the sensor data was collected by the car, analysis showed that a cyclist was hit because the system thought what it was seeing was not a cyclist: the probability of that was too low in that particular spot at that particular time. So the car pursued its course while the driver who was supposed to supervise the car's actions was sending text messages. Despite millions of miles driven by all these vehicles (we still see Waymo[3] cars every day on the streets of Palo Alto and Mountain View collecting as much data as possible), there will always be special cases and unexpected situations, and such incidents will happen. Still, autonomous cars will cause far fewer accidents than people will. One must understand

3. Google self-driving company.

how to use the technology while acknowledging its limits.

The examples I have just cited, as complex as they are, are all based on knowledge and recognition. They execute the tasks for which they have been programmed. They don't invent anything new, they just follow the given rules, examples and codes by using the data we have chosen to feed them. They may be called expert systems or Machine Learning or Deep learning, but they are no more than what we, human programmers, have decided they should be. All these technologies are meant to help us with specific tasks, which are often very repetitive and highly codified. They help us by increasing our own physical or intellectual capabilities but they certainly cannot replace us. There is no artificial intelligence that is beyond our control and that will lead to our extinction; there is just augmented intelligence that needs fair regulation so that it can support and sustain our own intelligence. All the services that are available today can be considered intelligent. AI is there to make decision making easier, not to make decisions for us. I prefer to speak of augmented intelligence, which enables intelligent people to be more capable and to perform better in specific domains. But people are the ones who have control, empathy and common sense, because what we teach machines is just general knowledge, which is an infinitely small part of intelligence.

Augmented intelligence is transforming science in a fundamental way. From physics to cosmology, genomics and chemistry, AI is used today in all these disciplines because of its automatic analysis and classification capabilities; and much more progress is expected. Information technologies have connected the world and made it smaller. Many applications have made our society more efficient, notably by reducing communication costs. Certain economists consider AI to be a GPT (General Purpose Technology) likely to affect all aspects of our economy. Just as the steam engine or electricity have revolutionized the world, it seems increasingly obvious that AI will have the same level of impact on our lives in the long term, and that it is indeed the new GPT. There is a tendency to overestimate the effect of a technology in the short term, and underestimate it in the long term. For this reason, it is important that our politicians introduce laws that favor AI development for the general public, keeping in mind that it is there to amplify and enhance our intellectual capabilities. Thanks to this amplification, we will have better means of transport, improved healthcare, a better environment, a higher level of well-being… a better life! It will undoubtedly bring with it a higher level of satisfaction and allow us to be more in touch with each other and with our environment.