Artificial intelligence: between science and the market

Some socio-historical elements to better understand a strange scientific experiment (1956-1990)

By Jean-Sébastien Vayre
Senior lecturer at the University of Côte d'Azur

Original article in French published in *Gérer & comprendre*, September 2021, n° 145, pp. 55-69.

Artificial intelligence (AI) is all the rage nowadays, with it being enthusiastically promoted by leading political and economic stakeholders involved in the development of digital technology. However, it is also a source of controversy, with some even claiming that it does not exist. It is a truly muddled affair. Is artificial intelligence but a mere illusion? This paper will explain why this is not the case. We will more precisely detail how this confusion surrounding artificial intelligence - very much a reality - has come about. To do this, we will posit that artificial intelligence is a scientific discipline that from its very origins was intertwined with an economic practice, resulting in an imbalance between basic and applied research. We will also build on this by concluding that it is precisely this imbalance that causes a lack of clarity surrounding the scientific discipline and, more generally, the instability of its development.

Foreword

There is much debate today among economic and political stakeholders concerning the notion of artificial intelligence: should we not instead talk about augmented intelligence, biological intelligence or remote intelligence? Some experts go even further by arguing that it would be best to no longer talk about artificial intelligence. From a scientific standpoint, this kind of discussion is interesting because it has been around since the emergence of AI. For example, Herbert A. Simon and Allen Newell were not particularly taken by the expression "artificial intelligence" coined by John McCarthy and preferred "complex information processing system" (Newell & Simon, 1956; Norberg, 2019 [1989]). However, while economic and political stakeholders engage in heated debate over the concept of artificial intelligence, they do not really argue on scientific grounds: the issues at stake are a cause of debate because they have serious consequences from a business standpoint. For example, it is not hard to see why Luc Julia (2019), head of the Samsung research centre, is pushing to replace the term "artificial intelligence" with "smart object": his professional interests are evident as Samsung's innovation strategy completely revolves around the Internet of Things (IoT).[1] As paradoxical as it may seem, claiming that "artificial intelligence doesn't exist" (Julia, 2019) therefore does not reflect, at least not in this case, a simple intention to demystify artificial intelligence: it is above all also a business move. This is why, in order to untangle the web that is AI, we wish to start from the very beginning and ask the following questions. What is AI? Is it a science? Is it a consumer item? Is it a new form of autonomous intelligence that might surpass human intelligence? We believe that trying to solve the problem of defining artificial intelligence is important, in that the virtuous quality of AI development depends on how society understands it, the meaning that we collectively give to it.

To answer these questions, we have conducted a socio-historical investigation using papers, reports, communications and videos produced by two major types of actors: those who have a long-standing interest in AI, and those who directly helped conceive and sustain this scientific discipline. This body of documentation was compiled as part of a doctoral project (Vayre, 2016), and was supplemented with research that we conducted over the last four years on the history of

[1] https://www.strategies.fr/actualites/marques/4027180W/-l-intelligence-artificielle-n-existe-pas-luc-julia.html (in French)

AI. Tables 1 and 2 list the resources used in our thesis which form the basis of the work presented in this paper. As mentioned previously, the work was produced with the aid of additional studies, and is therefore based on several sources not cited in Tables 1 and 2. However, these documents are systematically referenced in the body of text and in the paper's bibliography. In addition, the resources used to conduct our investigation, as a whole, were compiled using a methodology that we could deem abductive (Bruscaglioni, 2016), in that we have searched for and explored the substantive documentation to confirm or reject hypotheses made as our work progressed. In other words, and contrary to, for example, the work of Dominique Cardon and his colleagues (2018) which forms part of the development of what we may call, in reference to the French *Annales* school (Burguière, 1979), a quantitative history of AI,[(2)] our working approach instead follows on from the evidential paradigm proposed by Carlo Ginzburg (1980). The author of this study considers quantitative history, while having the merit of shedding light on the major structures that drive the dynamics of a given phenomenon over the long term, as tending to classify these dynamics under categories of thought that are at times far too general. Ginzburg (1980) therefore posits that the negative impacts of this tendency may be mitigated by adopting this mindset which lies, he believes, at the root of intellectual history, and consists of reconstructing an invisible reality by interpreting traces of the past that are perceptible in the present. In the words of Denis Thouard (2007), this way of "inferring from the facts" is, at least in Ginzburg's view (1980), a paradigm for research and thinking that is particularly useful in humanities and social sciences. We have therefore tried to adopt this model across all stages of searching for, compiling, reading and analysing documents comprising our study material. In short, the investigation findings detailed in this paper are the result of selecting documents and information that reflect fragments of empirical reality that we have gradually reassembled through knowledge and intelligibility effects, characteristic of "sociological reasoning" (Passeron, 1991).

At this point, we would like to specify that although we occasionally refer below to developments in AI over the last 20 years, we are primarily interested in the period from the mid-1950s to the early 1990s, as it was during this time that AI experienced its first waves of success and failure (Cardon, Cointet & Mazières, 2018). We will set out our findings below in two large sections, enabling us to distinguish between the scientific discipline and the economic practice that is AI. However, taking into account the work of Bruno Latour (1987), we are aware that this distinction has an abstract quality: from

(2) While the authors did conduct interviews in view of tracing the history of AI, they mainly relied on the statistical analysis of a corpus of over 27,000 articles compiled in 2018 on Web of Science.

Approach type	References
Symbolic artificial intelligence	(Bickhard & Terveen, 1995), (Bonissone & Johnson, 1984), (Fodor, 1975), (Fodor, 1983), (Fodor & Pylyshyn, 1988), (Forgy, 1981), (Gaschnig, 1980), (Ince, 1992), (Laird, Newell & Rosenbloom, 1987), (Lenat, 1977), (Lenat, 1983), (Lindsay, Buchanan, Feigenbaum & Lederberg, 1993), (McCarty, 1977), (Memmi, 1990), (Minsky & Papert, 1969), (Neumann, 1958), (Newell, 1980), (Newell & Simon, 1972), (Papert, 1988), (Samuel, 1959), (Simon, 1991 [1969]), (Simon, 1992), (Tristan & Abdallah, 2009), (Turing, 1950), (Winston, 1970)
Connectionist artificial intelligence	(Ackley, Hinton & Sejnowski, 1985c), (Bechtel & Abrahamsen, 1993), (Bickhard & Terveen, 1995), (Changeux, 1983), (Fodor & Pylyshyn, 1988), (Hayek, 1952), (Hebb, 1949), (Hopfield, 1982), (Lai, 2015), (LeCun & Bengio, 1995), (LeCun et al. 1989), (McCulloch & Pitts, 1943), (Memmi, 1990), (Minsky & Papert, 1969), (Noduls, 2015); (Numenta, 2011), (Numenta, 2014), (Numenta, 2015), (Papert, 1988), (Rosenblatt, 1958), (Rumelhart, Hinton & Williams, 1986), (Smith, 1999), (Smolensky, 1988), (Widrow & Hoff, 1960)
Hybrid artificial intelligence	(Bonasso, Firby, Gat, Kortenkamp, Miller & Slack, 1997), (Cassimatis, 2005), (DePristo & Zubek, 2001), (Hawes et al. 2007), (Kubera, Mathieu & Picault, 2011), (Langley & Choi, 2006), (Müller & Pischel, 1993), (Reynaud, 2014), (Schmidt, 2005), (Silver et al. 2016), (Smolenky, 1987), (Smolensky, Legendre & Miyata, 1992)

Table 1: Documents produced by researchers or organisations with direct involvement in AI development

Resource type	References
Written documents	(Blanc, Charron & Freyssenet, 1989), (Boise, 2007), (Copeland & Proudfoot, 2015), (Dupuy, 1994), (Hodges, 2014 [1983]), (Pélissier & Tête, 1995), (Varela, 1988)
Video records	(Dammbeck, 2003), (Folgoas, 1976), (Guirardoni, 1981), (Karlin, 1971), (Lallier, 1963), (Moreuil, 1972), (Royer, 1961a), (Royer, 1961b)

Table 2: Documents produced by researchers and journalists who have a long-standing interest in AI

a purely empirical standpoint, the research work that has conceived this discipline has inextricably linked the scientific and economic logics shaping the history of AI. However, we have decided to structure this article around this distinction precisely in order to untangle these two types of logic and thereby make the forms of their entanglement and the related issues easier to comprehend. We posit in the first section of this paper that AI is first and foremost a highly competitive scientific discipline, and broadly speaking seeks to conduct highly experimental research programmes. We shall then argue that AI research was, at least initially, a resounding failure from an applied research standpoint, even if it can be deemed a success from a basic research standpoint. Following this, in the second section, we will show how AI is also an economic practice about which many promises are made. We will then explain how this practice plays a vital social role in understanding and explaining the first waves of success and failure of AI. To conclude, we will outline some considerations to give a better understanding of how AI as a scientific discipline and as an economic practice interlink, stressing that, from a socio-historical perspective, this understanding provides insight into the current success of this discipline.

Artificial intelligence as a scientific discipline

The conference organised by John McCarthy, Marvin L. Minsky, Nathaniel Rochester and Claude E. Shannon at Dartmouth College in 1956 laid the institutional groundwork for artificial intelligence. In their proposal drawn up in preparation for this event, the four authors define this science as follows:

> "[Artificial intelligence] 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" (McCarthy et al., 2019 [1955]).

However, it must be noted that the origins of AI are more generally embedded in the history of computing and cybernetics. For example, the work of Blaise Pascal, Gottfried W. Leibniz, Charles Babbage, Augusta Ada King, George Boole, Friedrich L. G. Frege, Kurt Gödel, and, of course, of Alan M. Turing, John von Neumann, Norbert Wiener, Warren McCulloch and Walter H. Pitts played a vital role in the emergence of this science (Crevier, 1997 [1993]; Pratt, 1995 [1987]; Rose, 1986 [1984]). It also worth noting that, since the dawn of this science, the use of the term "artificial intelligence" has not been embraced by all. McCarthy was particularly taken by this term, who eventually persuaded his colleagues to adopt it. As previously mentioned, Simon and Newell preferred to talk about a "complex information processing system" (Newell & Simon, 1956; Norberg, 2019 [1989]).

Different styles of research

Much like other sciences, AI does not have a perfectly harmonious community: not all stakeholders collectively share the same perceptions of this science. In the words of Pierre Bourdieu (1976), AI is a scientific field the stability of which is dependent on the power struggle driving it, within which various forms of domination emerge and dissipate. This is especially true since AI is highly interdisciplinary in nature: depending on their interests, researchers in this field may stumble into such different areas as biology, psychology, anthropology, logic, philosophy, linguistics, mathematics, electronics and computing. However, the study of AI revolves around one shared goal: each and every researcher in the field has helped to test the hypothesis that a machine can exhibit behaviour that humans would generally deem intelligent. Since the discipline's beginnings, the methods of conducting this experimentation has been the subject of intense debate.

At the Dartmouth Summer Research Project the most prominent researchers in this community were Simon and Newell, and there are many reasons why this was the case. Firstly, in 1956, Simon and Newell were the only ones to have a computer program capable of synthesising one of the aspects of intelligence that academics often consider to be the most respected: solving complex mathematical problems. The Logic Theory (LT) program (Newell & Simon, 1956) is capable of proving half of the Principia Mathematica theorems of Alfred N. Whitehead and Bertrand A. W. Russell. Secondly, the LT program was designed using expertise in the fields of humanities and social sciences because the machine incorporates some of the fundamental concepts of the bounded rationality theory (Simon, 1945).(3) However, most researchers attending the conference believed that there was no point in studying human cognitive processes to design an AI program, such as McCarthy and Marvin L. Minsky, who, in the late 1950s, shared the view that AI must focus on exploring formal logic. This idea however is just as controversial as Simon and Newell's theory. For example, Herbert Gelernter and Nathaniel Rochester (1958), along with Oliver G. Selfridge (1959) understood AI from different perspectives. In their view, AI should not be formed by using human cognition or formal logic as a reference, but rather by using just the information processing capabilities of machines as a basis. This approach enabled them to develop their first AI programs: for Herbert Gelernter and Nathaniel Rochester, this was the Geometry Theorem Prover (GTP, Gelernter & Rochester, 1958), and for Oliver G. Selfridge, the famed pandemonium model (Selfridge, 1959).

Three major tension points

From the outset AI has been characterised by tensions. Over time, these tensions gradually intensified and eventually gave the field of AI a lasting structure. Between the 1960s and 1990s, there were at least three major tension points that played a decisive role in shaping the dynamics of this science.

(3) It was thanks to this theory that Simon was awarded the Nobel Prize for Economics in 1978.

The first point relates to amicable disagreements that quickly arose between McCarthy and Minsky concerning the way the issue of AI can be defined: while for McCarthy the fundamental issue underlying this new discipline was primarily one of logic, Minsky did not agree with this view (Norberg, 2019 [1989]). In 1960, this first point of tension emerged between the two researchers who, from that point onwards, undertook different research trajectories. As a result, in 1962 McCarthy decided to leave the MIT AI Lab to head his own one at Stanford University, the Stanford Artificial Intelligence Lab (SAIL). It was at this point that his work on logic had a significant impact on the AI expert community. For example, thanks to the list processing (LISP) language that McCarthy developed in 1958, Douglas Lenat was able to develop his Automated Mathematician (AM; Lenat, 1977) and EURISKO (Lenat, 1983) programs. Similarly, the "IF, THEN" advice taker program proposed by McCarthy in 1959 played a key role in the development of expert systems, just like his work a few years later on circumscription,(4) streamlining the information processing performed by these systems (Crevier, 1997 [1993]). With McCarthy's departure, Minsky was heading the MIT AI Lab on his own, a rather comfortable arrangement since he was receiving sizeable investments to outdo his new rival: over several years, the Defense Advanced Research Projects Agency (DARPA) gave $3m in funding for the Machine-Aided Cognition and Multiple Access Computer project to MIT (Flamm, 1987; Norberg, 2019 [1989]). With this funding, McCarthy's former colleague had considerable resources to establish an anti-logic approach to AI. Given the economic, technical and human resources available to the MIT AI Lab, this approach quickly became considerably popular among Minsky's peers. Many young and brilliant researchers as a result flocked to work with Minsky, including James R. Slagle, Joel Moses, Patrick Winston and Seymour A. Papert, who respectively developed the symbolic automatic integrator (SAINT; Slagle, 1961), the symbolic integration program (SIN; Moses, 1967), the arch concept learning program (Winston, 1970) and the LOGO programming language (Papert, 1971).

(4) This name refers to a computing process by which the obstacles potentially impeding the logic inference engines can be isolated or minimised, to allow for navigation within a knowledge-based system.

In short, as shown in Figure 1 (see above), from the McCarthy/Minsky split was born two major working approaches to AI. According to Roger C. Schank, up until the early 1990s there were two different research styles in the AI field: the "neat" style that subscribes to the logical approach developed by McCarthy, and the "scruffy" style more associated with the anti-logic approach developed by Minsky:

> "In Schank's view, the neat style is refined, focusing on superficial phenomena like logic and syntax, which can be understood and compartmentalised in pretty little boxes. The scruffy style is haphazard, and revelled in wrestling with tortuous issues such as semantics" (Crevier, 1997 [1993], p. 201).

This first point of tension could be supplemented with another that helps to define the research stream developed by Schank and the stream of his colleagues Simon and Newell. In the 1960s, Simon and Newell were teaching and researching at Carnegie Mellon University, and were highly esteemed among academics and industrialists alike (Norberg, 2019 [1989]). They both therefore quickly gained recognition for their work on problem solving. On the back of the LT program's wow factor, Simon and Newell continued to study and draw on human cognition to develop new computing programs. Their work led them to developing their famed General Problem Solver (GPS; Newell, Shaw & Simon, 1959) which planted the seed for the design of the most well-known expert systems. For example, Edward A. Feigenbaum and Bruce G. Buchanan directly based the development of the DENDRAL (Buchanan & Feigenbaum, 1978) and MYCIN (Buchanan & Shortliffe, 1984) expert systems on the GPS, as did their students Randall Davis and John P. McDermott, who respectively designed the TEIRESIAS (Davis, 1978) and "eXpert

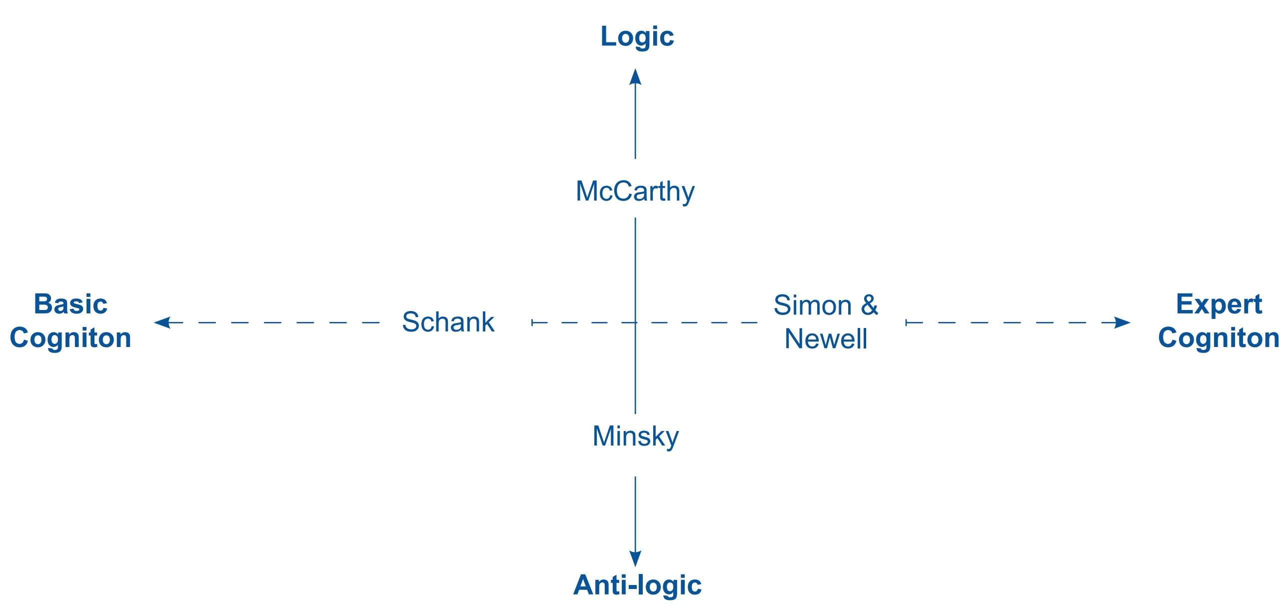

Figure 1. The four main research streams in symbolic AI

CONfigurer" (XCON; McDermott, 1982) programs. However, Schank's (1978) work at Yale University did not relate to solving expert problems. While Schank and Simon and Newell share the view that AI should be based on the study of human cognition, Schank was less interested in so-called high-level intelligence and more so in intelligence applied in daily life. In other words, Schank did not study expert cognition, but was rather focused on basic cognition. His aim was not to develop computer programs to enable machines to assist or supplant highly qualified individuals, but rather to understand how these machines can adapt to humans from day to day to help them live together in a better society.

It should be noted that, as the work of Cardon and his colleagues (2018) demonstrate, there is a third tension point that played a fundamental role in the history of AI. This third point is the opposition between the symbolic and connectionist stances on this discipline. Indeed, the two points of tension discussed just above subscribe to the symbolic stance on AI: the aforementioned researchers all generally hold the view that intelligence is a computational system of symbolic representations of a rather deliberative nature (Fodor & Pylyshyn, 1988). This is exactly why, at this point and following Simon and Newell, many researchers had a particular interest in the expert aspect. However, while this approach dominated the field of AI from the 1960s to the 1990s, a shift occurred starting from the 2000s (Cardon, Cointet & Mazières, 2018; Vayre, 2016). The connectionist stance on AI that Frank Rosenblatt (1958), Wilfrid K. Taylor (1956), Bernard Widrow and Marcian E. Hoff (1960) struggled to defend against the criticism of Minsky and Papert (1969),[5] is now predominant, with the success of deep learning and the neurological theories explaining individual and collective behaviour (Changeux, 1983) being a reflection of this. From this viewpoint, intelligence is formed of a gargantuan network of triggers and inhibitors for basic calculation units, which is rather adaptive in nature. The collective focus of AI researchers is less on expert reasoning and more on what we may call, rather crudely, the "child's learning experience" (Vayre, 2016). This other stance on AI, based on the biological nature of cognition and the mathematical nature of information processing, covers various research styles. So as not to spread the scope of this paper thin, we have decided, at least in this work, to not further touch upon this stance on AI.

Successes and failures: the case of the BAIR Lab

From the outset there have been many successes in AI: the programs developed by researchers in the field are able to perform the tasks they have been assigned to do. For example, the SHRDLU program developed by Terry Winograd (1972), is able to question its motivations and therefore has something which somewhat resembles "self-awareness"; these two properties allow the program to engage in a truly constructive conversation with the user, and this was one of the first instances of this behaviour in the field of AI. For illustrative purposes, here is the sort of discussion that a user can have with SHRDLU:[6]

> "23. Had you touched any pyramid before you put the green one on the little cube? YES, THE GREEN ONE [...]. 24. When did you pick it up? WHILE I WAS STACKING UP THE RED CUBE, A LARGE RED BLOCK AND A LARGE GREEN CUBE [...]. 25. Why? [...] TO GET RID OF IT [...]. 26. Why did you do that? TO CLEAN OFF THE RED CUBE [...]. 29. Why did you do that? BECAUSE YOU ASKED ME TO. (Winograd, 1972, p. 13).

As impressive as that may seem at a basic level, the problem is that this conversation can only truly be considered constructive within the bounds of the micro-world in which SHRDLU develops i.e., a virtual environment formed of several cubes, cones and spheres, a "toy world" you might say (Cardon, Cointet & Mazières, 2018). As explained previously, AI is an experimental discipline. In this respect, any successes, no matter their significance, are quickly overshadowed by failures: progress is systematically a reminder to those who make it of how many of their goals are far from being achieved. To understand this last point, we would now like to focus on the case of the Berkeley Artificial Intelligence Research (BAIR) Lab. The BAIR Lab is an equally insightful and interesting case of understanding the limitations and benefits of AI's experimental nature.

Since the early 1980s, the BAIR Lab had been headed by Robert Wilensky, a former student of Schank. In keeping with Schank's legacy, who was often considered the *enfant terrible* of AI, Wilensky did not care for formalism, whether it be logical or mathematical in nature. Unlike a number of his colleagues who saw in Noam Chomsky's work (1965) the potential to formalise human intelligence, he did not believe language could be reduced down to formal syntax. While he agreed with Chomsky's paradigm – according to which language is at the root of thought – he also believed that language poses a semantic problem and not a syntactical one. In other words, to reuse the Schanksian expression, Wilensky was "scruffy". He had an inclination for tricky problems, and developed a take on AI in his own image, being both original and bold. Indeed, Wilensky was an unusual individual, often considered a non-conformist by a fair number of his colleagues (Rose, 1986 [1984]). He liked originality, and it was probably because of this that he was drawn to the prevailing intellectual climate at the University of Berkeley: while the institution did not have a true computing culture when Wilensky arrived, it fostered an intellectual diversity that he appreciated. At Berkeley you could find anyone: idiosyncratic anthropologists, non-conformist linguists, cognitive psychologists, and, most importantly, Hubert L. Dreyfus and John R. Searle who played an active role in stimulating Wilensky's research. With their unrelenting criticism of AI, the two philosophers in fact fuelled the BAIR Lab in its work, and raised the profile of its director. For Wilensky, who

[5] However, note that in spite of their criticism, the two authors have a certain interest in connectionism. It is worth mentioning that Minsky (1954) wrote his thesis on neural networks.

[6] The passages in upper case are spoken by SHRDLU.

was interested in commonsense reasoning and the issue of basic actions, Berkeley was the perfect testing ground for new ways of understanding AI. In order to fully grasp the experimental nature of the research projects conducted at the BAIR Lab, a presentation of some of the main programs developed by Wilensky and his colleagues is provided below.

Much like the Script Applier Mechanism (SAM) developed by Schank and Abelson (1977), Wilensky's Plan Applier Mechanism (PAM; 1977) had a certain capability of understanding narratives and situations that it was told. In 1980, PAM was able to have the following discussion:

> "[Based on the following description:] John needed money, he got a gun and walked into a liquor store. John told the owner he wanted his money. The owner gave John the money and John left." [... and based on the following question:] "Why did the owner give the money to John?" [... PAM answered with, for example:] "The owner was scared that John would kill him" (Rose, 1986 [1984], p. 71).

Like SAM, PAM was able to exhibit a certain degree of understanding in that it demonstrated knowledge that was not explicitly contained in the statement it was told. However, PAM has a certain edge over SAM: to understand a situation, PAM did not need its creator to provide it with the underlying scenario. Naturally, much like Winograd's SHRDLU program (1972), PAM is only capable of such a feat provided that the statements it was told relate to its micro-world. Furthermore, as demonstrated by Frank Rose (1986 [1984]) with the case of "Plan ANalyzer with Dynamic Organization, Revision and Application" (PANDORA), expanding this micro-world requires a myriad of computing tricks that are just as much ways of questioning the workings of human intelligence. For example, Joe Faletti (1982), a student of Wilensky who developed PANDORA, struggled to make his program understand that the act of going to fetch a newspaper from the letterbox may require different behavioural patterns depending on the weather. Cognitively speaking, such a capacity to understand and adapt relates to significant planning issues, particularly in terms of organising the goals and sub-goals of a particular act and its constituent tasks, but also in terms of memorising relevant information – organising and applying knowledge to perform every task (Faletti, 1982). For instance, in order for PANDORA to put on a coat, it had to know that rain is wet and that being dry is a desirable state, but also that a coat protects from the rain. As odd as it may seem, for Faletti, this sort of problem was equally as important as it was difficult to resolve from a computing standpoint. The threefold benefit of the work of Wilensky and his colleagues is evident in this respect as well. Together, they underscored that:

- as basic as it may seem, an action entails different forms of problem solving which, despite being automatic in nature, are cognitively complex;
- these forms of problem solving are inseparable from the social conventions that existed before the given action;
- the coordination of cognitive and social aspects in completing any human action (even the most trivial ones) entails a form of intelligence that is extremely difficult to identify, describe, understand and formalise.

In opposition to the simplistic discourse on AI that often emerges, the case of the BAIR Lab is proof that AI is not merely a community of researchers wanting to impose their logician and mathematician viewpoints by applying them in the field of humanities and social sciences. For Wilensky and his colleagues, computers are an implement for scientific experimentation, with the heuristic benefit of helping them to question and understand what intelligence is. However, this point of view is not specific to the BAIR Lab: for example, as already noted, Simon, Newell and their Carnegie Mellon students share this viewpoint. The case of the BAIR Lab is also of interest to us for another reason, one that is embedded in this critical and original vision of AI that Wilensky and his colleagues adopted. Following on from the work of Minsky on frames (1974), PANDORA[7] was a method of representing a notion of intelligence with a computer, which used Searle's Background theory (2002). According to this theory, language is a code whose meaning cannot exist without the social conventions that enable its expression. In this respect, history has shown that Wilensky and his team failed in their project to design a computer program capable of simulating basic cognition. Can this lead us to conclude that their research program was a failure? It depends on who we ask. Investors like DARPA or IBM would say yes: it was a computer program that merely worked within the bounds of a micro-world created by a researcher, with no political or economic application. In contrast, a sociologist interested in the history of science and technology would clearly say no. Naturally, with its experiments, the BAIR Lab was unable to confirm the hypothesis that basic cognition can be represented by computers. However, we see this failure to be a huge success, since this unsuccessful venture was a stepping stone for Wilensky to more effectively test out the complexity of the interaction between cognition and culture, the difficulty in representing this complexity with computers, and in particular the concept that the effectuation of this complexity is required to correctly perform, analyse and understand the smallest basic action.

As fragile as it may be, the cognitive value behind this conclusion is particularly high since, following on from the work of Simon and Newell, it raises the question of how we conceive intelligence. The experiments conducted at the BAIR Lab lead to a hypothesis being formed: while Simon and Newell quickly managed to produce satisfactory simulations of expert cognition, this was because, contrary to the belief of potentially most academics, this cognition was probably less intelligent than it seemed. While it may seem outrageous, Wilensky and his colleagues were not so sure that solving half of the Principia Mathematica required more brainpower than going to pick up mail from a letterbox.

(7) As well as the PAMELA program designed by Peter Norvig to supplement it (Faletti, 1982).

AI as an economic practice

We have seen that AI builds on cybernetics and the history of ideas that form it, meaning that AI emerged with the development of the first computers: its origins coincide with those of computing which, we stress, embraces scientific, technological and industrial areas of activity. In this respect, it is important to keep in mind that AI is inextricably linked to the major socio-technological innovations that paved the way for the computerisation of society (Mounier-Kuhn, 2010). In the 1950s, these innovations were prohibitively expensive, which meant pioneers in AI had to quickly partner with political and economic stakeholders to fund their research. This was nothing terribly new: the forefathers of AI had to do the same thing. For example, Turing worked with the British government to crack the enigma code and help the Allied powers defeat the Nazis (Hodges, 2015 [1983]), John von Neumann collaborated with the American government to enhance the explosive power of the atomic bomb and provide the Americans with a tool to intimidate the Soviet Union (Hoddeson et al., 1992) and, more generally, in the aftermath of the two world wars, mathematicians and cyberneticians attending the Macy Conferences wanted to help establish a new world order to guarantee peace among mankind and the "mental well-being" of the people i.e. their autonomy and intellectual freedom (Heims, 1991).

In other words, building on the arguments put forward previously, it must be stressed that to have a career in AI, you cannot just be a renowned researcher who is respected by your peers; you have to also be able to draw in investors to receive funding for as long as possible (Latour ,1987). To do this, AI researchers had to navigate the political and economic spheres, particularly because DARPA was the biggest source of funding for this field for quite some time. It was precisely through this specific form of "economisation" (Akrich, 1989) that AI was able to enter into the public forum. Many controversies have as a result surfaced, leading to spillovers into other areas that have obfuscated the collective understanding of what AI is. We would like to examine the history of AI in relation to the market in order to better understand why this obfuscation was able to take root, and also what its impacts are on the advancement of this science.

Promises to draw in investors

AI is a scientific field torn between the three points of tension detailed above (see section Disappointment and fears). It is important to understand that, in order to compete within such a tumultuous field, AI researchers had to find partners who could provide them with the suitable technological and financial resources to conduct their research programmes. They were therefore collectively compelled to lay down bridges between the scientific, political and economic worlds so that the above-mentioned divergent viewpoints could exist. We have seen that, from the outset, AI breeds a turbulent working environment: while AI researchers know how to play nice, especially when searching for partners who will help them conduct their work more effectively, they are also well aware that in order to achieve their career goals they will have to jostle for position. The leading academic institutions and universities with an AI laboratory foster this competitive environment for at least two reasons. The first one is that, as we have seen with the MIT AI Lab, these laboratories may on occasion pledge several millions of dollars per year to the institutions and universities hosting them. The second is that, given the military-industrial complex's interest in AI, these very institutions and universities strive to draw in the most esteemed researchers. As a result, in AI, science and the market end up sustaining each other to form a particular "opinion economy" (Orléan, 2000) in which the scientific value of the research programme conducted by a given laboratory is not the only factor that matters any more: there is also – and above all – the element of the researchers' ability to flex their muscles before their peers, raise positive public interest, and draw the attention of political and economic stakeholders and build trust with them.

Within this highly competitive environment, arrogance can at times give an edge. This is what at least seems to be the case with Simon and Newell, who, as previously mentioned, always received considerable recognition in the field of AI. The two associates from Carnegie Mellon in fact had a reputation for self-importance. For example, in an interview, Minsky told Daniel Crevier (1997 [1993]) that Simon and Newell came across as aloof during the Dartmouth Summer Research Project. The other attendees believed that the two researchers seemed just as pleased as they were flattered that they were the only ones to present an AI program. Simon himself would go on to confirm this observation (Crevier, 1997 [1993], p. 67). In 1997, 40 years after the famed conference, Simon and Newell's confidence in their work had not at all faltered, and in fact had only bloated. In his 1991 book, Simon said that, with their invention of a computer program capable of processing symbolic data, he and Newell had demonstrated how a system composed of matter can exhibit the attributes of thought. In Simon's view (1991), their work held the key to unlocking the mystery of the dualism of the mind and body. This claim is naturally subject to debate, as demonstrated by the work of Daniel Dennett (1991).[8] Nevertheless, as questionable as it is from a scientific standpoint, the claim is a good reflection of the degree of confidence one has to deal with when working with AI researchers. In the vein of Simon and Newell, AI pioneers are researchers with key expertise in mathematics and computing, but also in humanities and social sciences. To make it in this extremely competitive environment, researchers have to learn to showcase their expertise and unique qualities to political and economic stakeholders. This is why, to impress investors while securing their full trust,[9] AI researchers have to provide a high level of assurance in relation to their work. In this respect, a considerable

(8) It is important to stress that this is not an attempt to disregard the key role played by Simon and Newell in the development of the philosophy of the mind.

(9) Note that the sums of money involved are huge, and it has been known, even in the 1960s, for funding exceeding $1 million per year to be provided.

number of eminent AI researchers since the late 1950s have had to collectively make predictions that often err on the rosy side. Once again, Simon and Newell are among the researchers who did not hesitate to abuse their scientific authority to give credence to overly ambitious promises:

> "1. [That] within ten years a digital computer will be the world's chess champion, unless the rules bar it from competition. 2. [That] within ten years a digital computer will discover and prove an important new mathematical theorem. 3. [That] within ten years a digital computer will write music that will be accepted by critics as possessing considerable aesthetic value. 4. [That] within ten years most theories in psychology will take the form of computer programs, or of qualitative statements about the characteristics of computer programs." (Simon & Newell, 1958, pp. 7-8).

Simon and Newell were of course discerning to some degree, as some of the predictions above were correct. However, in the strictest sense of the word, all the predictions were false and should have been lowered: for example, we would have to wait until 1997 – and not 1968 – for the Deep Blue supercomputer to win against Gary Kasparov in a chess match. For Simon and Newell however, whether their predictions would be proven true or false was not that important. The two colleagues quickly understood that applications of AI could transform into a market brimming with management technology serving all stakeholders in the production and distribution chains of goods and services, including consumers (Cochoy, Smolinski & Vayre, 2016). What mattered to Simon and Newell was that their predictions were equally as reasonable as they were rosy for the military-industrial complex with which they were very familiar. The two researchers knew how business worked, and more specifically how military and industrial business was run: Simon and Newell were also consultants for the RAND Corporation. In other words, even though they knew as researchers that their predictions were not true in the scientific sense, they knew as consultants that the predictions were promises likely to draw in investors.

Disappointment and fears

AI has always been an unsettling field, since it attempts to understand human behaviour from an objective and detached perspective. In this respect, it is important to understand that when Turing argued, in 1950, that a machine has the potential to produce thoughts, his main intention was to shake up the intelligentsia of the time. This forefather of artificial intelligence knew that he was a homosexual at this point, and rebelled against the commonly held beliefs of his time: he was not convinced by the often religious, authoritarian and dubious lines of thinking that, for example, considered women, and to a greater extent animals, incapable of demonstrating intelligence (Turing, 1950). Why was intelligence considered sacrosanct by some? Was there anything that could rule out the theory that a machine can exhibit intelligent behaviour? Simon and Newell quickly realised that the provocative nature of AI research in itself could provide a socioeconomic edge. They knew that the market liked innovation and that it could be an invaluable partner in combating scientific orthodoxy that could hinder the advancement of AI. At least at the beginning, and to establish this discipline as a scientific one, AI researchers could only partially rely on academic institutions: they had to find other means of securing the lasting future of AI. However, Simon and Newell were not alone in realising this. For example, Minsky in his own way helped to publicise AI, painting a more or less realistic picture of what it could produce. He did this most notably through science fiction, advising Stanley Kubrick during the filming of *2001: A Space Odyssey* (1968; Ganascia, 2019 [2016]). There are also other researchers who used other means to give publicity to AI. It is important to bear in mind that the overall goal for these researchers – expressed with varying degrees of clarity – was to promote the market expansion of this science by bringing it onto the communications market.

This basis caused a number of difficulties to emerge from a socioeconomic standpoint. The pioneers of AI most likely had developed communication channels between their discipline and the market too quickly. This hastiness was particularly due to the fact that they felt a sense of urgency since they needed powerful and costly machines to get ahead of their competitors. In view of this, while scientific competition is more strongly influenced by the political and economic dynamics of the capitalist system through its involvement on the market, it is clear that AI was going from strength to strength and was firmly established within academic spheres. In the early 1960s, a flurry of promises were being made, and investors were being hooked in. AI became a media sensation, and it entered into its golden age. However, this era was not set to last. After the highs of great expectations came the equally as notable lows of disappointment. For example, following the rather negative assessment made by the Automated Language Processing Advisory Committee (ALPCA; Pierce et al., 2019 [1966]) concerning the progress made in the field of machine translation, the US government decided in 1966 to halt investments which were initially intended to fund the translation of Soviet Union press releases (Hutchins, 1996). As previously stated, this was only the first in a long line of failures as economic and political stakeholders saw it. The Shakey robot (Nilsson, 2019 [1984]) for example had no military or industrial use, given that the tasks it could perform were slow and essentially a series of jerky movements. What is more, Shakey was very sensitive to changes to its surroundings: With just a slip of the wheels, its perception of its surrounding environment would no longer correspond with the actual situation (Hart & Nilsson, 1972). Speech Understanding Research (SUR) by Donald E. Walker (2019 [1973]) was also another attempt in vain to find an application for AI. This system was not viable since its users had to severely restrict their grammar usage so that SUR could process their request in real time. Ultimately, this technology was more difficult to use than the traditional menu selection systems (Crevier, 1997 [1993]). This was also the case for expert systems which, at least in the 1980s, were a great success however. One such system was XCON (Bachant & McDermott, 1984): after a number of years, updating its knowledge base became a true ordeal. In the words of Cardon and his colleagues (2018), XCON turned into a “cathedral of rules”: this expert system was so complex that the Digital Equipment Corporation (DEC) was required to invest over $2 million per year for its maintenance (Simon, 1987).

Nevertheless, AI has not only just been a disappointment to investors: from the outset, it has also struck fear in consumers. To understand this phenomenon, it must be first noted that the communications market revolving around AI is lucrative, particularly as it draws up fantasies, promises, but also substantiated risks of varying degree. The problem is therefore that, in the eyes of the public, this market generates a mix of information that turns AI into a catch-all and troubling concept. As the following excerpts from interviews show, AI ended up scaring consumers, and this fear can be understood in different ways.

For Michel Melkanoff for example, this fear is irrational, since the risk of AI comes not from the machines themselves but the people designing and using them.

> “There are those who are afraid of machines [...] that [...] will turn into superhuman robots who will take over the world [...]. I have something to say about that. […] [Nobody] can truly have serious concerns [...] over a bunch of wires and metal, it is an irrational fear. Interviewer: “The atomic bomb is a bunch of wires and metal too!” Sure, but it’s not the bomb that people are afraid of, it’s the people dropping them. In this respect, there is perhaps a threat posed by those able to use computers” (Michel Melkanoff, quoted in Moreuil, 1972).

Abraham Moles has a different view on the matter. For the computer and communications science expert, this fear is rather the result of what he calls a “sociological concern”, a fear of varying rationality among people relating to the forms of alienation caused by using AI:

> “The public is afraid of machines […] as they reveal their nature [… and …] pervade our day-to-day life […]. As René de Possel noted, when 47 million French citizens will be classified under 1,000 or 2,000 criteria, each one stored on a punch card, no more police files, no more proceedings, everything stored in a central registry, then you will be able to identify every individual. They will no longer be anonymous. They will be [...] personalised, not able to rely on interstitial freedom or the workings of institutions. They will be prisoners! I believe that this is why [...] people are afraid” (Abraham Moles, quoted in Lallier, 1963).

In reference to the work of Madeleine Akrich, Michel Callon and Bruno Latour (2006), it is therefore through a more or less mastered economic practice that AI researchers transformed their work into promises and applications to interest and bring on board investors. This transformation was then amplified by the communications market revolving around AI. The problem is that, from the outset, this amplification was a distortion obscuring the true nature of AI. The public ended up forgetting for quite some time that AI is first and foremost a science. It is very well possible that this is still the case today: how many people have an understanding of AI research programmes, or the epistemological, social and human issues related to the field? From the 1960s to 1980s, the public saw AI at best as a kind of mechanism or energy, in vague and incomprehensible terms, existing within machines to regulate their operations and developing in a more or less dangerous manner... And at worse, AI was a massive scam.

Criticisms to reassure and shake up the market

For a large section of the population, the concept of AI lost its meaning. It became a source of discomfort, but not really one for the researchers: even those who were not completely happy with this notion got used to it rather painlessly (see the section “Different styles of research”). It was industry stakeholders working in the development of computing who were uncomfortable with the concept. This was the case for IBM for example:

> “The AI projects carried out within the firm [IBM] were eventually a victim of their own success. […] During a shareholders’ meeting, Thomas J. Watson was asked to explain why the company funnelled research investments into such worthless fields. The IBM marketing department had also observed an alarming change in consumer psychology: they considered computers a threat and abandoned them out of fear. [F]or Watson, this was the last straw [...]. The firm’s future marketing campaigns [...] threw away the image adopted from science fiction of a computer acting as a giant brain and replaced it with one that was reassuring, of a machine simply processing figures. Computers, IBM unflaggingly claimed, [...] would only do what they were told. They would never oust an executive, as their sole talent was in quickly processing massive data flows” (Crevier, 1997 [1993], p. 78).

Industry stakeholders’ discomfort with AI worsened with the many disappointments previously mentioned, to the extent that, as previously mentioned, economic and political stakeholders questioned their commitment to developing this field. These stakeholders therefore took a genuine interest in the criticisms of AI. Stuart E. Dreyfus, a consultant from the RAND Corporation, took this opportunity to put his brother Hubert L. Dreyfus into contact with the research organisation. Hubert L. Dreyfus was called upon to assess, from a philosophical standpoint, the viability of the AI project: the RAND Corporation wanted him to predict this field’s ability to confirm the theory that behaviour deemed intelligent by humans can be materially replicated. After his investigations, Dreyfus (1972) gave a resounding no: he believed that intelligence bore no relation to a system that computes symbolic representations and does not entail the performance of logic operations. Dreyfus specifically felt that in contrast to humans, machines crudely perform calculations: they are unable to distinguish between what is relevant and what is not (a problem of restriction). According to the philosopher, even though humans can translate the complexity of the world into simple responses, this same complexity has to be reduced, formalised and made plain for a machine to be able to respond similarly. In Dreyfus’ view, this was an impossible task, at least for symbolic AI which was prevalent at the time. Along these lines, he added that while humans have no trouble adapting to changing environments, this does not extend to machines, which only know how to follow explicit rules (a problem of framework). In summary, Dreyfus believed that AI could not qualify as a science given the irrational nature of its inherent research hypothesis. AI was, in his view, a sort of overambitious alchemy. More generally, Alchemy and Artificial Intelligence (Dreyfus, 2019 [1965]) provides vociferous criticism of AI: it is a provocative and poorly documented report that goads readers into a scrappy debate. This paper rendered such criticism irrelevant and clumsy. Even Joseph Weizenbaum, one of the very few AI researchers to side with Dreyfus, thought that his colleague's report was poor, particularly because it demonstrated a meagre understanding of how computers work (Crevier, 1997 [1993]).

Nevertheless, by seeking the services of Hubert L. Dreyfus, the RAND Corporation legitimised his ideas and contributed to distorting the scientific debate into a socioeconomic controversy which quickly turned into an armchair debate. During such discussions, scientific arguments were thrown out for crude insults. For example, in response to Dreyfus' provocation that a six-year old could beat any computer program at chess – which, at a point in time, was actually the case – Papert publicly challenged the philosopher to beat such a child in a game.[10] Of course, Dreyfus was right in many respects, and his criticism was fundamentally interesting: his poor understanding of AI was offset by his strong knowledge of philosophy. It was this knowledge that allowed him to considerably beat McCarthy to the punch in identifying the two major limitations of AI: the aforementioned problems of restriction and framework. In any case, whether he was right or wrong matters little. The important thing is to understand that starting with Alchemy and Artificial Intelligence (Dreyfus, 2019 [1965]), the controversy surrounding AI spilled over from the field of science to become a socioeconomic issue.

Generally speaking, this problem relates to two major factors. The first concerns a collective form of ramping up commitments (Joule & Beauvois, 2002). This refers to the idea that some researchers, who were probably far too committed to AI development, continued to keep to their promises (so as not to lose face) while also seeking to regain the market's trust (and more specifically investors' trust). The second factor concerns a fictitious reconstruction of reality. It brings together the AI critics who wanted to play the game of industry stakeholders like IBM by denying the existence of AI. The problem was therefore that these critics had forgotten, more or less willingly, something very important: AI is not some form of mechanical autonomous thinking that computer-related technologies could develop. It is a science, the existence of which is hard to deny. The socioeconomic issue of AI is therefore twofold: on one hand (for the escalation of commitments), the issue is linked to the fact that this science's funding was reliant on bluffs that risked perpetuating the cycle of promise/disappointment until this disrupted the positive results of scientific programmes in this field; and, on the other hand (for the fictitious reconstruction of reality), the issue is related to the fact that the strategy consisting of distorting what is AI to have better grounds to deny its existence was a method of masking its development. This concealment gave vendors considerable power: without even truly realising, they became the only stakeholders who could ensure and control not only the dissemination of AI technology applications, but also the means of its funding as a result.

Following the prevalence of the commitment escalation problem for several years (AI winters), the second factor of the socioeconomic problem of AI (the fictitious reconstruction of reality) came to the fore from the 1980s. As shown in Figure 2, during this period, IT companies no longer wanted to talk about expert systems, and even less so AI. They preferred to act as mere IT solution providers in order to seem both far-sighted and diligent in the eyes of their customers. Consequently, AI programs became hidden applications, being discreetly integrated into more traditional computer programs. Patrick H. Winston for example was very familiar with this strategy. In the 1980s, like most of his colleagues, he owned a computer program development business, explaining that the programs were based on what he called a "raisin bread" system:

> "AI is currently integrated into systems like raisins in a loaf of raisin bread: the raisins do not occupy much space, but they often provide the principal source of nutrition. You cannot remove the raisins from the bread; and there are many types of raisins" (Patrick H. Winston, quoted in Crevier, 1997 [1993], p. 252).

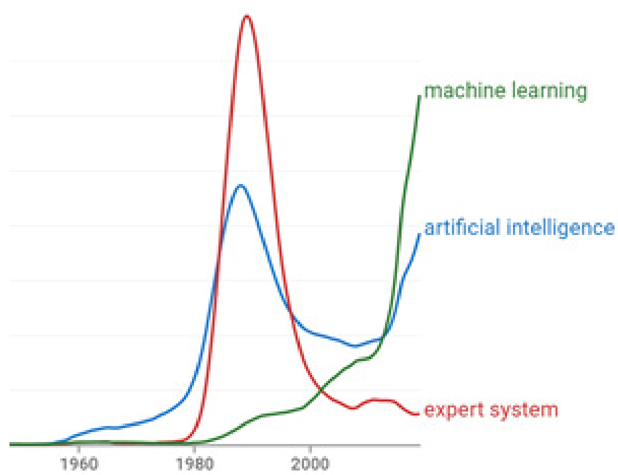

Figure 2. Graph of the number of occurrences of the terms "artificial intelligence", "expert system" and "machine learning" in Google Books Ngram Viewer[11]

A typical example of the "raisin bread" system is the commercial assistance program. Traditionally, this program would just check product availability, record the transaction, draw up the invoice, and notify the shipping service provider. This program was also able to be enhanced with a specially designed expert system

[10] For more details on the concrete forms of this debate, please refer to Papert's report (2019 [1968]), written in response to Alchemy and Artificial Intelligence (Dreyfus, 2019 [1965]).

[11] The lines forming this graph are a sort of sounding board of empirical reality and therefore have a slight lag behind the figures presented in this paper, given that they are generated using statistical analysis of the corpora of texts available on Google Books. For instance, the peak of the occurrences of the term "expert system" is in 1988, while the golden age of expert systems in the US was actually from the late 1970s to early 1980s (Crevier, 1997 [1993]).

that could, for example, make suggestions for substitute products in the event of shortages. The benefit of this new way of conceiving the AI business model was that by concealing the existence of the AI technology, the media frenzy surrounding this science gradually calmed down. The confusion over what is AI dissipated at the same time as any concerns, fears and related risks. There was a drawback however: this discipline and the technological applications it created would continue to exist. As previously mentioned, the concern would still linger because the organisation of the dissemination of AI technology applications shifted between the hands of economic stakeholders whose interest should not be mistaken with those of society. It would take, as we have witnessed in the past decade, the significant and rapid increase in digital data production, storage and processing capacities – which paved the way for a new age of machine learning[12] (see Figure 2) – for society to be aware of this issue and once again question the economic, social and humans stakes of AI development and of the dissemination of related applications (see the big data movement; Cardon, 2015; Vayre, 2016).

Conclusion

AI is a scientific discipline with a research programme that was, at least at the beginning, highly experimental: it tests the hypothesis that the intelligence of humans – and, by extension, of all living beings – can be materially replicated. History has shown us that, from a purely scientific standpoint, AI had the merit of contemplating what intelligence is and therefore brought about major developments in not only the field of cognitive sciences but also humanities and social sciences. We posited that working in AI was not just engaging in scientific activity, but also in an economic practice. To carry out work in this science, there is a need to draw in investors who can fund costly equipment: in AI, science and the market are inextricably linked. This is why, in a socioeconomic context in which the major digital stakeholders (GAFAM – Google, Apple, Facebook, Amazon and Microsoft) tend to draw in leading AI researchers, we believe that it is worth stressing that historically AI was first and foremost an experimental scientific discipline which seeks to better understand what intelligence is and how we can (or cannot) synthesise it. In this respect, our work has the advantage of highlighting that, for pioneers, AI was not a consumer good that leads to the development of automated services like, as is the case currently, the platform economy for example: it was a science that lets researchers ask fundamental questions that could result in successful applications.[13] However, our studies have revealed that, from the outset, AI has also been an economic practice consisting of dressing up scientific ideas in a political and economic fashion so as to link them to socio-technical uses. We have sought to show how this window dressing has historically driven push and pull dynamics with economic stakeholders (i.e. investors and consumers). It is a factor that should be better understood, particularly if we want to have a greater insight and grasp of the actual stakes of AI development.

[12] Major successes in this field include the AlphaGo, Watson and DeepL Translator programs respectively developed by DeepMind, IBM and DeepL.

[13] Even though these applications were actually rather unsuccessful at the beginning.

References

AKRICH M. (1989), « La construction d'un système socio-technique. Esquisse pour une anthropologie des techniques », *Anthropologie et Sociétés*, 13(2), pp. 31-54.

AKRICH M., CALLON M. & LATOUR B. (2006), *Sociologie de la traduction : textes fondateurs*, Presses des Mines, Paris.

BACHANT J. & McDERMOTT J. (1984), "R1 revisited: four years in the Trenches", *AI Magazine*, 5(3), pp. 21-32.

BOURDIEU P. (1976), « Le champ scientifique », *Actes de la recherche en sciences sociales*, 2(2-3), pp. 88-104.

BRUSCAGLIONI L. (2016), "Theorizing in grounded theory and creative abduction", *International Journal of Methodology*, 50(5), pp. 2009–2024.

BUCHANAN B. G. & FEIGENBAUM E. A. (1978), "Dendral and meta-dendral: their applications dimension", *Artificial Intelligence*, 11(1-2), pp. 5-24.

BUCHANAN B. G. & SHORTLIFFE H. E. (1984), *Rule-based expert systems: the mycin experiments of the stanford heuristic programming project*, Addison Wesley, Boston.

BURGUIÈRE A. (1979), « Les Annales, 1929-1979 », *Annales. Économies, sociétés, civilisations*, n°6, pp. 1344-1346.

CARDON D. (2015), À quoi rêvent les algorithmes : nos vies à l'heure des big data, Paris, Seuil.

CARDON D., COINTET J.-P. & MAZIÈRES A. (2018), « La revanche des neurones : l'invention des machines inductives et la controverse de l'intelligence artificielle », *Réseaux*, 5(211), pp. 173-220.

CHANGEUX J.-P. (1983), L'Homme neuronal, Paris, Fayard.

CHOMSKY N. (1965), *Aspects of the theory of syntax*, MIT Press, Cambridge.

COCHOY F., SMOLINSKI J. & VAYRE J.-S. (2016), "From marketing to 'market-things' and 'market-ITing': accounting for technicized and digitalized consumption", *in* CZARNIAWSKA B. (dir.), *A research agenda for management and organization studies*, Edward Elgar, Glasgow, pp. 26-37.

CREVIER D. (1997 [1993]), À la recherche de l'intelligence artificielle, Paris, Flammarion.

Davis R. (1978), "Knowledge acquisition in rule-based systems – knowledge about representations as a basis for system construction and maintenance", *Pattern-Directed Inference Systems*, pp. 99-134.

DENNETT D. (1991), *Consciousness explained*, Little, Brown and Co, Boston.

DREYFUS H. L. (1972), *What computers can't do: a critique of artificial reason*, Harper & Row, New York.

DREYFUS H. L. (2019 [1965]), *Alchemy and artificial intelligence*, available online: https://www.rand.org/content/dam/rand/pubs/papers/2006/P3244.pdf

FALETTI J. (1982), "PANDORA - A program for doing commonsense planning in complex situations", *AAAI-82 Proceedings*, p. 185-188.

FLAMM K. (1987), *Targeting the computer: government support and international competition,* Brookings Institution, Washington.

FODOR J. & PYLYSHYN Z. W. (1988), "Connectionism and cognitive architecture: a critical analysis", *Cognition*, 28(1-2), pp. 3-71.

GANASCIA J.-G. (2019 [2016]), « Marvin Minsky : un des cerveaux de l'intelligence artificielle », available online: https://interstices.info/marvin-minsky-un-des-cerveaux-de-lintelligence-artificielle/

GELERNTER H. L. & ROCHESTER N. (1958), "Intelligent behavior in problem-solving machines", *IBM Journal of Research and Development*, 2(4), pp. 336-345.

GINZBURG C. (1980), « Signes, Traces, Pistes – Racines d'un paradigme de l'indice », *Le Débat*, 6(6), pp. 3-44.

HART P. & NILSSON N. (1972), *Shakey: experiments in robot planning and learning*, available online: https://www.youtube.com/watch?v=GmU7SimFkpU

HEIMS S. J. (1991), *The cybernetics group*, MIT Press, Cambridge.

HODDESON L., HENRIKSEN P. W., MEADE R. A. & WESTFALL C. (1992), *Critical assembly: A technical history of Los Alamos during the Oppenheimer years, 1943–1945*, Cambridge University Press, Cambridge.

HODGES A. (2015 [1983]), *Alan Turing : le génie qui a décrypté les codes secrets nazis et inventé l'ordinateur*, Paris, Payot.

HUTCHINS J. (1996), "ALPAC: the (in)famous report", *MT News International*, n°14, pp. 9-12.

JOULE R.-V. & BEAUVOIS J.-L. (2002), *Petit traité de manipulation à l'usage des honnêtes gens*, Grenoble, Presses Universitaires de Grenoble.

JULIA L. (2019), *L'intelligence artificielle n'existe pas*, Paris, First.

LALLIER J. (1963), *Visa pour l'avenir : des machines à penser ?*, available online: https://www.ina.fr/video/CPF86656441/des-machines-a-penser-video.html

LATOUR B. (1987), *Science in action: how to follow scientists and engineers through society*, Harvard University Press, Cambridge.

LENAT D. B. (1977), "Automated theory formation in mathematics", *Proceedings* of the 5th international joint conference on Artificial intelligence, vol. 2, pp. 833-842.

LENAT D. B. (1983), "Eurisko: a program that learns new heuristics and domain concepts", *Artificial Intelligence*, 21(1-2), pp. 61-98.

MCCARTHY J., MINKSY M. L., ROCHESTER N., SHANNON C. E. (2019 [1955]), *A proposal for the Dartmouth summer research projet on artificial intelligence*, available online: http://jmc.stanford.edu/articles/dartmouth/dartmouth.pdf

McDERMOTT J. (1982), "R1: a rule-based configurer of computer systems", *Artificial Intelligence*, 19(1), pp. 39-88.

MINSKY M. L. (1954), *Theory of neural-analog reinforcement systems and its application to the brain model problem*, Princeton University, Princeton.

MINSKY M. L. (1974), *A framework for representing knowledge*, Massachusetts Institute of Technology, Cambridge.

MINSKY M. & PAPERT S. A. (1969), *Perceptrons*, MIT Press, Cambridge.

MOREUIL F. (1972), *Un certain regard - Les machines et les hommes : l'intelligence artificielle,* available online: https://www.ina.fr/video/CPF86655105/l-intelligence-artificielle-video.html

MOSES J. (1967), *Symbolic integration*, Massachusetts Institute of Technology, Cambridge.

MOUNIER-KUHN P.-E. (2010), *L'informatique en France de la seconde Guerre Mondiale au Plan Calcul. L'émergence d'une science*, Paris, Presses de l'Université Paris-Sorbonne.

NEWELL A. & SIMON H. A. (1956), "The logic theory machine: a complex information processing system", *IRE Transactions on Information Theory*, 2(3), pp. 61-79.

NEWELL A., SHAW J. C.& SIMON H. A. (1959), "Report on a general problem-solving program", *Proceedings* of the International Conference on Information Processing, pp. 256-264.

NILSSON N. J. (2019 [1984]), *Shakey the robot*, available online: http://www.ai.sri.com/pubs/files/629.pdf

NORBERG A. L. (2019 [1989]), *An interview with Marvin L. Minsky*, available online: https://conservancy.umn.edu/bitstream/handle/11299/107503/oh179mlm.pdf?sequence=1&isAllowed=y.

ORLÉAN A. (2000), « L'individu, le marché et l'opinion : réflexions sur le capitalisme financier », *Esprit*, 11(269), pp. 51-75.

PAPERT S. A. (2019 [1968]), *The artificial intelligence of Hubert L. Dreyfus: a budget of fallacies*, available online: https://dspace.mit.edu/bitstream/handle/1721.1/6084/AIM-154.pdf?sequence=2&isAllowed=y

PAPERT S. A. (1971), *A computer laboratory for elementary schools*, Massachusetts Institute of Technology, Cambridge.

PASSERON J.-C. (1991), *Le raisonnement sociologique. L'espace non-poppérien du raisonnement naturel*, Paris, Nathan.

PIERCE J. R., CARROLL J. B., HAMP E. P., HAYS D. G., HOCKETT C. F., OETTINGER A. G. & PERLIS A. (2019 [1966]), *Language and machine: computers on translation and linguistics*, available online: http://www.mt-archive.info/ALPAC-1966.pdf

PRATT V. (1995 [1987]), *Machine à penser : une histoire de l'intelligence artificielle*, Paris, Presses Universitaires de France.

ROSE F. (1986 [1984]), *L'intelligence artificielle : histoire d'une recherche scientifique*, Paris, Payot.

ROSENBLATT F. (1958), "The Perceptron: a probabilistic model for information storage and organization in the brain", *Psychological Review*, 65(6), pp. 386-408.

SCHANK R. C. (1978), "Computer understanding of natural language", *Behavior Research Methods & Instrumentation*, 10(2), pp. 132–138.

SCHANK R. C. & ABELSON R. P. (1977), *Scripts, plans, goals, and understanding*, Lawrence Erlbaum, Hillsdale.

SEARLE J. R. (2002), *Consciousness and language*, Cambridge University Press, Cambridge.

SELFRIDGE O. G. (1959), "Pandemonium: a paradigm for learning", *Proceedings* of the Symposium on Mechanisation of Thought Processes, pp. 511-529.

SIMON H. A. (1945), *Administrative behavior: a study of decision making processes in administrative organization*, The Free Press, New York.

SIMON H. A. (1991), *Models of my life*, Basic Books, New York.

SIMON H. A. & NEWELL A. (1958), "Heuristic problem solving: the next advance in operations research", *Operations Research*, 6(1), pp. 1-10.

SIMON R. (1987), "The morning after", *Forbes*, 140(8), pp. 164-168.

SLAGLE J. R. (1961), *A heuristic program that solves symbolic integration problems in freshman calculus: symbolic automatic integrator (SAINT)*, Massachusetts Institute of Technology, Cambridge.

TAYLOR W. K. (1956), "Electrical simulation of some nervous system functional activities", *Information theory*, n°3, pp. 314-328.

THOUARD D. (2007), *L'interprétation des indices. Enquête sur le paradigme indiciaire avec Carlo Ginzburg*, Villeneuve-d'Ascq, Presses Universitaire du Septentrion.

TURING A. M. (1950), "Computing machinery and intelligence", *Mind*, vol. 59, pp. 433-460.

VAYRE J.-S. (2016), *Des machines à produire des futurs économiques : sociologie des intelligences artificielles marchandes à l'ère du big data*, Toulouse, Université Toulouse Jean Jaurès.

WALKER D. E. (2019 [1973]), *Speech understanding research*, available online: https://apps.dtic.mil/dtic/tr/fulltext/u2/757049.pdf

WIDROW B. & HOFF M. E. (1960), "Adaptive switching circuits", *1960 IRE WESCON Convention Record,* n°4, pp. 96-104.

WILENSKY R. (1977), "PAM - A Program That Infers Intentions", *Proceedings* of the Fifth International Joint Conference on Artificial Intelligence, vol. 1, pp. 1-15.

WINOGRAD T. (1972), "Understanding natural language", *Cognitive Psychology*, 3(1), pp. 1-191.

WINSTON P. H. (1970), *Learning structural descriptions from examples,* Massachusetts Institute of Technology, Cambridge.