

Cross-interactions between AI and epistemology

Jean-Gabriel Ganascia

1. Introduction

The aim of this paper is to show that the mutual epistemological stakes of artificial intelligence (AI) and sciences, both 'hard sciences' and Human and Social Sciences (HSS), are multiple. It specifically addresses two of them. The first is reflexive: it concerns the epistemology of AI itself, which, as a scientific discipline, deserves a philosophical and historical look at its foundations. It is a question of specifying the nature of this discipline, which cannot just be reduced to a technology and which, as a science, is neither a theoretical science, even if it has originally been founded by mathematicians, nor a 'science of nature' strictly speaking, nor really a 'science of culture' that is a discipline of the humanities. To clarify these different issues, we shall first recall the genesis of AI, its history and its definitions, before trying to approach its epistemological status.

The second issue is related to the uses of AI, machine learning and data processing in different scientific disciplines and the major changes that these uses induce in these disciplines by automating tedious tasks. In doing so, our aim here is to show that AI techniques do not only allow to automate certain tasks, but that they also contribute to designing new interpretation operators, new proof procedures and, more generally, new scientific approaches such as *in silico sciences* (cf. Ganascia 2008). In other words, the contribution of AI is not only practical; it introduces into these scientific disciplines what Gaston Bachelard calls an epistemological rupture (cf. Bachelard 1938), that is, a dissociation between the primary evidence of observation and the scientific facts resulting from experimentation.

To address these different epistemological issues, the paper is divided into two main parts framed by this introduction and its conclusion. The first part is dedicated to the genesis, the history and the epistemology of AI, while the sec-

and concerns the impact on the theoretical sciences, on the sciences of nature and on the sciences of culture, i.e., on the humanities.

2. AI groundings

2.1 Prehistory of AI

Attempts to formalize the laws of thought and to automate reasoning are ancient. Born in antiquity, logic aimed to give the laws of right thinking; to do so, it characterized, by means of formal-mechanical rules, valid reasoning as being sequences of inferences — inferences being formal manipulations of symbolic expressions — that correspond to elementary figures listed as being valid themselves.

In classical formal logic, i.e., in Aristotelian or in Stoic logic, the set of elementary inferences was determined from ‘regulatory’ syllogisms, i.e., figures that lead from two propositions to a third. Later, in the second half of the 17th century, Leibniz tried to mathematize logic, i.e., the laws of correct thinking, in order to prove the validity of an argument by a calculation, without having to memorize all the valid elementary syllogisms, as in traditional Aristotelian logic. This was immediately followed by the desire to automate this calculation on a machine. Therefore, we can say that Leibniz is a forerunner of artificial intelligence since he tried — unsuccessfully — to draw the plans of a machine capable of reasoning by itself. This project was taken up in the 19th century by George Boole who created binary algebra to account for the laws of logic and then by William Stanley Jevons, who actually built a machine, the ‘logical piano’, that could mechanically deduce the consequences of logical premises based on the work of George Boole.

Several other attempts to automate reasoning were made in the early 20th century. Let’s think for instance of the mechanical chess player machines of Torres y Quevedos that were built in 1912 and 1920. We must also mention the cybernetic movement with Warren McCulloch and Walter Pitts (cf. McCulloch & Pitts 1943), Claude Shannon and Norbert Wiener, among others, because it was also at the origin of several attempts to reproduce thought on electronic computers. Finally, we must not forget Alan Turing, who wondered, in his famous article *Computing Machinery and Intelligence* published in 1950 (cf. Turing 1950), what it means for a machine to think and how to build such a thinking machine.

2.2 Birth and epistemic assumptions of AI

However, although the idea that it's possible to build a machine that reproduces thought had been around for a long time, the term 'artificial intelligence' did not appear for the first time until 1955 in a summer school proposal submitted by four researchers, John McCarthy, Marvin Minsky, Nathanael Rochester, and Claude Shannon to the Rockefeller Foundation for a grant to organize a summer school at Dartmouth College, New Hampshire, in 1956. For the promoters of this summer school, artificial intelligence was a scientific discipline that aimed to study intelligence with computers. More precisely, to quote them, "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" (McCarthy et al. 1955: 1).

This means that all cognitive faculties, in particular reasoning, calculation, perception, memorization and even scientific discovery or artistic creativity, could be described with such precision that it should be possible to reproduce them using a computer. Let us insist on the epistemological importance of this conjecture: it draws a horizon of tasks to be accomplished, just like Galileo's postulate according to which the Book of Nature is "written in mathematical language". Since then, despite the considerable progress made and the changes in the technologies used, from those based on symbolic logic to numerical and emergent connectionist approaches, and despite the various debates about the parallels between the nature of intelligence itself and the way it is simulated by machines, the study of artificial intelligence has always been based on the same conjecture, which nothing has yet been able to disprove or prove irrefutably. To clarify, it's important to specify that what the philosopher, John Haugeland, has mistakenly called GOFAI ("good old fashioned artificial intelligence", Haugeland 1985) is, as some AI researchers like Drew McDermott have mentioned, a myth, for many reasons, in particular because the seminal text on AI, the Dartmouth College Summer School proposal (cf. McCarthy et al. 1955), explicitly mentioned neural networks as methods that had to be developed by AI. This doesn't mean that the way the machine simulates "every aspect of learning or any other feature of intelligence" is similar to the way they are implemented in nature. For example, just as Frederick Jelinek beautifully put it in his famous "airplanes don't flap their wings", so the learning styles of children and machines are not at all the same at all and the way the computers work has nothing to do with the way our brains work, even if they make use of neural networks.

Nevertheless, the idea that it is possible to reproduce all the cognitive functions of any intelligent being still constitutes the epistemological horizon of AI.

The same scientists, who were trying to reproduce various cognitive abilities such as reasoning, theorem proving, image or speech recognition, knowledge representation in memory, etc. on computers, were in parallel tempted to take a practical advantage of these simulations and to incorporate them into many technological devices. Very soon, for example, Herbert Simon and Alan Newell wrote papers both on the performance of general-purpose problem-solving computer programs (cf. Newell/Simon 1956) and on human problem-solving using AI tools as cognitive models to study human reasoning (cf. Newell/Shaw/Simon 1958).

The incorporation of AI simulations in technologies has been very popular in recent years, giving AI the privilege of being one of the most active fields of applied research in many areas such as medicine, agriculture, geology, etc. To-day, when people talk about AI, they almost always mention the various technological applications of AI. This most often corresponds to the current meaning of the term AI.

Finally, it should be noted that, among the general public, the success of the term 'artificial intelligence' is often due to a damaging misunderstanding according to which AI would produce artificial entities endowed with intelligence and which, as a result, would compete with human beings. This idea, which refers to ancient myths and legends such as that of the Golem, has recently been revived by contemporary personalities such as Stephen Hawking or Elon Musk, by engineers such as Ray Kurzweil, or by the proponents of what is now called 'strong artificial intelligence' or 'general artificial intelligence'. We will not discuss this meaning here, because it only attests to an abundant imagination inspired more by science fiction than by a tangible scientific reality confirmed by experiments and empirical observations.

2.3 Very brief history of AI

Since its birth, even if the seminal definition of the Dartmouth College Summer School has always remained valid, AI has undergone many evolutions that we can summarize in six stages.

2.3.1 The time of the prophets

A few achievements, in particular the Logic Theory Machine (cf. Newell/Simon 1956), which automatically proved logic theorems, the seminal work of Arthur

Samuel on reinforcement learning applied to the game of checkers (cf. Samuel 1959) and the first efficient neural network learning process, the so-called Perceptron (cf. Rosenblatt 1958), aroused enthusiasm. In the euphoria that followed, the researchers let themselves go to some rather unconsidered declarations that they have been much reproached for afterwards. For example, on November 14th, 1957, Herbert Simon delivered a speech at the banquet of the Twelfth National Meeting of the Operations Research Society of America in which he said:

I am willing to make the following predictions, to be realized within the next ten years:

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.

The lecture was then transcribed and the paper was co-signed with Alan Newell and published in the journal *Operation Research* (Simon/Newell 1958: 7).

2.3.2 The dark years

In the mid-1960s progress was not as fast as expected. In particular, a chess-playing machine was defeated by a ten-year-old boy in 1966, which made the first point of Herbert Simon's statement suspect and by contaminating the others, so AI received some bad press, which resulted in some dark years for AI. This corresponds to what is now called the AI winter, a period during which AI research became less popular, although contrary to popular belief, work never stopped altogether. For example, the first chatbot named Eliza was created by Joseph Weizenbaum at MIT between 1964 and 1966 (cf. Weizenbaum 1966), and later, Terry Winograd (cf. Winograd 1971), still at MIT, developed a famous program called SHRDLU for natural language understanding, i.e., for translating simple sentences into logical formulas. Note finally that during that AI winter, Marvin Minsky and Seymour Papert (cf. Minsky/Papert 1969) showed the intrinsic limitations of the Rosenblatt's Perceptron learning algorithm, be-

cause it was restricted to two-layers neural networks while Warren McCulloch and Walter Pitts, in their seminal paper (McCulloch/Pitts 1943) show that only three-layers neural networks were universal, i.e., able to implement any Boolean logic function.

2.3.3 Semantic artificial intelligence

Nevertheless, as previously said, during that AI Winter, work never stopped. Researchers were then focused on new directions and inspired by works in psychology and linguistics, which gave birth to the first cognitive science approaches. Note that interest in human cognition is far older and that cybernetics had already attempted to model social and cognitive processes with information processing mechanisms. However, new interdisciplinary approaches combining artificial intelligence, psychology and linguistics began in the mid-1970s. This corresponds to what has been called the 'semantic turn'. It led to an increased interest in modeling memory, in the mechanisms of comprehension, which was tried to be simulated on a computer as well as in the role knowledge plays in reasoning. This is what gave rise to knowledge representation techniques (cf. Bobrow/Winograd 1976) with semantic networks (cf. Collins/Quillian 1969) and frames (cf. Minsky 1974), to object-oriented programming and to so-called expert systems, because they used the knowledge of human experts to reproduce their reasoning. The latter raised enormous hopes in the early 1980s.

2.3.4 Neo-connectionism and machine learning

In parallel with the rise of artificial intelligence in the early 1980s, the techniques derived from cybernetics and connectionism were perfected, freed from their initial limitations and made the object of multiple mathematical formalizations. More specifically, as mentioned above and as Marvin Minsky and Seymour Papert had shown (cf. Minsky/Papert 1969), Rosenblatt's Perceptron learning algorithm was restricted to elementary logic functions. In the mid-1980s, this algorithm was generalized to multi-layer neural networks (cf. Rumelhart/Hinton/Williams 1986), giving rise to the backpropagation learning algorithm, which wasn't subject to such limitations. This led to distributed parallel processing, which enabled the use of neural networks in many supervised machine learning tasks.

2.3.5 From artificial intelligence to 'animistic informatics'...

Since the late 1990s, artificial intelligence has often been coupled with robotics and human-machine interfaces to produce intelligent agents that suggest the presence of another, whether it be human, or just an abstract entity. This trend of artificial intelligence can be sketchily characterized as a form of computer animism insofar as it seeks to elicit the projection of a breath of life onto the everyday objects of our environment. The current successes of Chatbots and, more recently, of ChatGPT, testify to the vitality, popularity and fashion of this trend.

2.3.6 The renaissance of artificial intelligence

With the massive development of the Web it became necessary to deal with large amounts of data. More specifically, since the rise of Web 2.0 at the turn of the century, the economics of the Web giants were based on targeted advertising, which made profiling critical. It follows that, based on information about individual behavior such as search queries, websites visited, etc., profiling had to scale to the size of the Web, which required dealing with massive amounts of data. This became known as 'Big Data'. The computational power of machines gradually made it possible to use large corpora of data with machine learning techniques, such as SVM, Kernel Machines, or Random Forests, which made AI very popular. Then, since the 2010s, the extension of Neural Network architectures to Convolutional Neural Networks (CNN) corresponding to the techniques currently called Deep Learning (cf. LeCun/Bengio/Hinton 2015), has produced impressive results that have tremendously accelerated the efficiency and the use of AI techniques.

Later, the Generative Adversarial Nets (GAN) enabled significant advances in image generation techniques (cf. Goodfellow et al. 2014), and the notion of transformers (cf. Vaswani et al. 2017) enabled the construction of Large Language Models (LLM) with hundreds of billions of parameters and impressive text generation techniques of which chatGPT is a popular example.

2.4 Epistemology of AI

The different steps of the evolution of AI corresponded to different epistemological views of this discipline that can be characterized as follows.

2.4.1 Logical-mathematical approach

The first works of artificial intelligence in the fifties and sixties were based on mathematical modeling, in particular on statistics and logic. This has been the case of automatic theorem provers (cf. Newell/Simon 1956), problem solving (cf. Newell/Shaw/Simon 1958) and the first attempts at machine learning, in particular reinforcement learning (cf. Samuel 1959). This gave rise to a science of models, to a 'science of the artificial', to use the title of a book by a pioneer of artificial intelligence, Herbert Simon (cf. Simon 1969), which is distinct from both the natural and the cultural sciences.

2.4.2 Semantic approaches

At the same time, there was a scientific current that used behavioral psychology to evaluate the plausibility of cognitive models of thinking or learning. From the end of the 1960's onwards, a new trend was inspired by other approaches from psychology (cf. Collins/Quillian 1969), in particular Charles Bartlett's schema theory and the theory of prototypes, and from linguistics, with the transformational grammars stemming from Chomsky's theories on the one hand and Fillmore's case grammars or Montague's semantic grammar on the other hand, in order to better understand human cognitive abilities before modeling them. Knowledge representation techniques (cf. Bobrow/Winograd 1976), in particular semantic networks (cf. Collins/Quillian 1969), frames (cf. Minsky 1974) and knowledge-based systems or expert systems are directly derived from these works. This led to a tension between two views of artificial intelligence, one focusing more on the logical-mathematical properties required to simulate cognitive processes on machines to be possible, the other on the study of the psychological processes to be modeled (cf. Newell 1982). This tension was resolved in the early 1980s with the logical formalization of knowledge representation techniques, in particular with description logics, that now form the basis of so-called formal ontologies (cf. Brachman/Fikes/Levesque 1983).

2.4.3 Learning theories and deep learning

From the eighties and the implementation of many learning models (Top-Down Induction of Decision Trees, Genetic Algorithms, Reinforcement Learning, Neural Networks, in particular Back-prop algorithms that generalized the perceptron (cf. Rumelhart/Hinton/Williams 1986), Inductive Logic Programming, etc.), there were attempts to theorize machine learning with, in particular, Leslie Valiant's work on the theory of learnability (cf. Valiant

1984) and Vladimir Vapnik's on statistical learning (cf. Vapnik 1999). These approaches were at the origin of new approaches, in particular ensemble methods ('bagging' and 'boosting') and support vector machines (SVM), which appeared to be prominent in AI since the mid-1990s.

2.4.4 Big Data

Since the beginning of this century, the Web Giants have been using aforementioned machine learning techniques such as SVM to process very large masses of data that are counted in gigabytes (10^9 bytes), terabytes (10^{12} bytes), and even petabytes (10^{15} bytes). Some claim that huge amounts of data solve all problems, without the need for theory or knowledge representation (cf. Anderson 2008), although this is highly debatable from an epistemological point of view. However, since 2010, pragmatic approaches using formal neural networks organized in multiple layers, the so-called Deep Learning techniques (cf. LeCun/Bengio/Hinton 2015), have produced statistical results far superior to previous models, without having any mathematical theory to explain them. This seems to be of great interest from an epistemological point of view, which is ours in this paper. However, nothing says that such a theory will not be available in the future.

3. Impacts of AI on sciences

AI does not only aim at better understanding intelligence by breaking it down into cognitive functions, simulating each of them and exploiting these simulations for technological purposes. It also transforms the scientific activity itself. This is the question we will address in the second part of this article.

3.1 Impact on the natural sciences: *In silico* experimentations

Today, almost all facts can be reduced to huge data sets. It follows that it is possible to induce and test theories directly from data using AI and Machine Learning (ML) techniques, without having to conduct experiments in the outside world. These data sets come from collecting information issued from sensors, or from automated analysis such as the sequencing of macromolecules like proteins or DNA. In addition, computer models make it possible to simulate parts of the physical world and conduct experiments on the results of these models. Undoubtedly, this kind of experimentation is changing scientific ac-

tivity, at least in part. This is obviously the case in the natural sciences, since many real-world experiments no longer need to be performed, which seems highly desirable for both economic and ecological reasons...

At the end of the 1980s, biologists who wanted to give a name to this type of experiment performed with computers, or more precisely, with the silicon microchips that make up the core of computers, invented a new Latin idiom: *in silico* (see http://en.wikipedia.org/wiki/In_silico). The term was constructed by analogy with — and in contrast to — *in vivo* experiments, i.e., experiments on living organisms, and *in vitro* experiments, which relate biological mechanisms to chemical processes reproduced in glass test tubes. Of course, this term reflects the growing role of computers in the sciences in general. But a careful study shows that computers are not just new tools here, but represent an epistemological turn in the empirical sciences in general, because they change the status of the experiment.

To be more precise, let us recall that in ancient times, science was first and foremost a question of observation and for Plato the most important sense was that of sight. Later on, in modern times, touch took over from sight: people wishing to understand the natural world spent more and more time provoking the subjects they were studying. Thus, in the 16th century, Andreas Vesalius (1514–1564) renewed human anatomy by dissecting the corpses of people condemned to death. Scientific experimentation in its modern meaning corresponds to this reversal: it is not enough just to observe; a scientist will intervene in the world in order to first understand it and then to transform it. This active intervention in the real world continued relentlessly: soon, autopsies no longer satisfied naturalists, who chose to provoke natural phenomena on the living body in order to understand the life springs. They then went further and started performing what are known as *in vivo* experiments because they are carried out on living beings. And so it went on: investigation was not only a question of touching and provoking nature, but also of reconstructing it. This led to the idea of reproducing *in vitro*, i.e., in glass test-tubes, the chemical reactions that are at the origin of the elementary physiological functions.

Today, this trend continues, not only with glass test tubes, but also with computers: we now think we can imitate all natural mechanisms, especially those of the living, reducing them not uniquely to chemical processes, but also to information processing. This gives rise to *in silico* experiments, which are experiments of a singular form in the sense that they no longer call upon the external senses, whether sight or touch, but only upon the temporal unfolding of logical and/or mathematical operations.

Insofar as the *in silico* experiments take place virtually, without touching their object of study, but by operating only on transformations of its representations, they are similar to 'thought experiments' (cf. Mach 1976; Sorensen 1992), even if they clearly can't be assimilated to them, since they provide objective results. And the detailed examination of *in silico* experiments seems to confirm this intuition. Indeed, their role in contemporary scientific activity is twofold.

The first role is to validate hypotheses on large amounts of pre-recorded data such as those obtained from the sequencing of genomes or proteins or from simulation of physical phenomena. Any experiment is, of course, the confrontation of a hypothesis with reality, but, in the case of *in silico* experiments, the observations are collected before the hypothesis is put forward, whereas in classical experiments, the scientific hypothesis led to the construction of an experimental apparatus through which data was collected to validate or invalidate the initial theory. The *in silico* experiments are thus presented as imaginary experiments in which hypotheses are tested on facts that are stored in memory. Note that, in addition to hypothesis validation, AI techniques can automatically generate many plausible hypotheses from data sets which can then be tested for facts. This led to the partial automation of scientific discovery. More precisely, being given an ontology, the machine becomes able to generate hypotheses and to test them on data (cf. Kings et al. 2004).

The second role of *in silico* experiments concerns the simulation of natural processes: just as, in any mental experiment, we reproduce real phenomena in our imagination, so, in many *in silico* experiments, the computer mimics material processes by transforming representations. The *in silico* experiment corresponds then to a virtual intervention on a fictitious world.

What's new today is the central role that *in silico* experiments play in contemporary scientific activity. Whereas in the past, many philosophers — including Karl Popper, one of the most famous — have criticized the role of 'thought experiments' in science (cf. Popper 1959), because they did not provide a strong scientific justification, today, *in silico* experiments, which are the computational equivalent of 'thought experiments', are now scientifically defensible, because they provide some tangible results and they are refutable. In other words, and in conclusion, the extensive use of *in silico* experiments in the natural sciences represents an epistemological turn that deserves attention.

3.2 Impacts on the humanities

This revolution in the natural sciences is accompanied by a major transformation in the humanities, i.e., in the disciplines that study human works. In this case, it is no longer a question of extracting general laws from data by induction, but of interpreting individual cases, for example literary works, on the basis of a large variety of data. In literature, we can try to identify markers of influence in the writings of great authors. It will then be possible to validate certain hypotheses, thereby renewing the traditional disciplines of scholarship.

In order to understand this specificity of the epistemological changes of the humanities, by distinguishing them from the transformations that have taken place in the so-called 'hard' sciences, we will draw on the opposition introduced by neo-Kantian philosophers, in this case Heinrich Rickert (cf. Rickert 1921) and Ernst Cassirer (cf. Cassirer 1923;1942), at the beginning of the 20th century, between the 'sciences of the nature', which deal with the world as it appears to us, and the 'sciences of the culture', which study human works. They — and particularly Ernst Cassirer (cf. Cassirer 1942) — show that both the natural sciences and the cultural sciences are empirical sciences, i.e., based on observable facts, but that the logic of each is different. The sciences of nature aim mainly to construct general laws by induction from observations and forgetting individual cases, while the sciences of the culture focus principally on the individual cases to give them meaning by explaining them. In this case, however, it is no longer a matter of extracting general laws by induction from data, but of interpreting individual cases, for example literary works or historical episodes, by using a great variety of data in order to understand them, or, more precisely, to give them meaning. To do this, an approach based on what logicians call abduction must be adopted, that is to say, on the search for explanations in the light of general theories. Thus, in the case of literature, we can try to identify markers of influence in the writings of great authors. From then on, it becomes possible to validate certain hypotheses empirically, which renews the traditional disciplines of scholarship.

Note that, in practice, the distinction between the 'sciences of the nature' and the 'sciences of the culture' is not so abrupt, since there are many cases where 'sciences of the nature' are also, at least in part, 'sciences of the culture' and vice versa. For example, medicine and health sciences are obviously 'sciences of the nature', while the nomenclature reflects medical traditions that depend on culture. Similarly, geography, which is clearly a 'science of the cul-

ture' is also, and in part, a 'science of nature', since it is based on many hard sciences.

Moreover, many epistemologists note that the logic of most of the 'sciences of nature' is not strictly inductive, since the process of discovery has sometimes been seen as mainly abductive, and that the deduction obviously has a place in any scientific reasoning. Symmetrically, the logic of the 'sciences of the culture' is not strictly abductive; deduction plays a role and it may happen that induction be used in some disciplines. This may be the case in literary studies when characterizing the style of an author (cf. Jockers 2013), or the figure used in a particular genre (cf. Boukhaled/Ganascia 2015) or again the expression of a character in a theater play. Nevertheless, it is clear that abduction plays a major role in 'sciences of culture', while induction is prominent in many 'sciences of nature'. Our goal, here, is to show that AI can be useful both for the 'sciences of nature' by mainly providing tools for automatic induction, and for the 'sciences of culture' by helping to interpret individual cases.

To aid in this search for interpretation, a certain number of tools have been developed and deployed that perform multiple operations, such as comparing textual states (cf. Ganascia 2011) or searching for reuses (cf. Ganascia/Glaudes/Del Lungo 2014), or, in archaeology, reconstructing pottery or buildings in three dimensions. These tools do not simply automate existing tasks. They propose new interpretive operators that completely transform the disciplines of scholarship. To illustrate, in the literary domain, Franco Moretti (cf. Moretti 2005) introduces the notion of distant reading, where he identifies general characteristics on large corpora, such as sentence length or punctuation. Similarly, we can characterize quotations or borrowings on large corpora, still in the literary domain. Note that, in both cases, whatever the size of the corpora may be, the inferences are clearly not inductive, but abductive, since they don't generate knowledge by themselves, but help interpretation.

These new interpretation operators have a double contribution. Some have a purely heuristic function by suggesting new avenues of research that need to be explored. They then help bring to light hitherto hidden phenomena, allowing human works to be seen under new conditions. These lines of research then require more rigorous investigation, with proven methodologies.

Others bring empirical elements of validation or invalidation of working hypotheses, for example, in the literary field, by highlighting certain influences, or on the contrary by showing the absence of explicit and/or implicit references and citations. In the latter case, the very scientific basis of certain

disciplines is strongly modified, since, as in the natural sciences, the very notion of proof evolves with the introduction of AI in the cultural sciences.

4. Conclusion

Finally, let us recall that one of the pioneers of AI, Herbert Simon, wrote a book entitled *The Sciences of the Artificial* (cf. Simon 1969), in which he discusses scientific approaches to modeling and the function of models in science. This could lead to the question of what characterizes AI as a science: is it exclusively a theoretical science, based on mathematics, or is it an empirical science? And in the latter case, is it more akin to the natural sciences or, to use the terminology of the neo-Kantian philosophers mentioned earlier, such as H. Rickert or E. Cassirer, to the ‘sciences of the nature’ or to the HSS, i.e., to the ‘sciences of the culture’? What makes us lean toward the former possibility is that Machine Learning is inherently inductive, aiming to generate general rules from particulars. What makes us lean toward the latter possibility is that AI is largely concerned with the modeling of deliberate individual practices that are the result of conscious activities and thus can be seen as human works. As a study of human works, it is therefore a science of culture, in the sense that the term has been defined above. Undoubtedly, the methods it uses are essentially based on mathematical and statistical approaches. At the same time, from a logical point of view, a large part of the activity of AI consists in calculating for and simulating tasks that are the fruit of some human practices, such as those mentioned here, and that, as such, belong to culture. Thus, the study of the relations between AI and HSS leads not only to showing the historical interest of AI for HSS, to highlight the use of AI by HSS and the modifications of the latter, with AI, or what we call the ‘computational turn’ of the latter, but also to show, in this respect, the proximity between AI and HSS.

List of references

- Anderson, Chris (2008): “The End of Theory: The Data Deluge Makes the Scientific Method Obsolete.” June 23, 2008 (<https://pdodds.w3.uvm.edu/files/papers/others/2008/anderson2008a.pdf>).
- Bachelard, Gaston (1938): *La Formation de l'esprit scientifique*, Paris: Presses Universitaires de France.

- Bobrow, Daniel G./Winograd, Terry (1976): *An Overview of KRL, A Knowledge Representation Language*, Stanford Artificial Intelligence Laboratory Memo AIM-293, Stanford: Stanford University.
- Boukhaled, Mohamed Amine/Ganascia, Jean-Gabriel (2015): "Computational Study of Stylistics: a Clustering-based Interestingness Measure for Extracting Relevant Syntactic Patterns." In: *Proceedings of the 16th International Conference on Intelligent Text Processing and Computational Linguistics*, Cairo, Egypt, April 14–20.
- Brachman, Ronald/Fikes, Richard/Levesque, Hector (1983): "KRYPTON: A Functional Approach to Knowledge Representation." In: *IEEE Computer* 16/10, pp. 67–73.
- Cassirer, Ernst (1923): *Substance and Function*, Chicago: Open Court.
- Cassirer, Ernst (1942): "Zur Logik der Kulturwissenschaften." In: *Göteborgs Högskolas Årsskrift 47* [Translated into English under the title *The Logic of the Humanities* (1961)], New Haven, CT: Yale University Press.
- Collins, Allan M./Quillian, Ross (1969): "Retrieval Time From Semantic Memory." In: *Journal of Verbal Learning and Verbal Behavior* 8/2, pp. 240–247.
- Ganascia, Jean-Gabriel (2008): "'In silico' Experiments: Towards a Computerized Epistemology." In: *Newsletter on Philosophy and Computers* 7/2, pp. 11–15.
- Ganascia, Jean-Gabriel (2011): "MEDITE – A Unilingual Text Aligner for Humanities. Application to Textual Genetics and to the Edition of Text Variants." *Supporting Digital Humanities (SDH 2011)*, Copenhagen, Denmark (http://www.poleia.lip6.fr/~ganascia/Medite_Project?action=AttachFile&do=get&target=SDH2011.pdf).
- Ganascia, Jean-Gabriel (2015): "The Logic of the Big Data Turn in Digital Literary Studies." In: *Frontiers in Digital Humanities* 2/7 (<https://doi.org/10.3389/fdigh.2015.00007>).
- Ganascia, Jean-Gabriel/Glaudes, Pierre/Del Lungo, Andrea (2014): "Automatic Detection of Reuses and Citations in Literary Texts." In: *Literary and Linguistic Computing* 29/3, pp. 412–421.
- Goodfellow, Ian/Pouget-Abadie, Jean/Mirza, Mehdi/Xu, Bing/Warde-Farley, David/Ozair, Sherjil/Courville, Aaron/Bengio, Yoshua (2014): "Generative Adversarial Nets." In: *Proceedings of the International Conference on Neural Information Processing Systems (NIPS 2014)*, pp. 2672–2680.
- Haugeland, John (1985): *Artificial Intelligence: The Very Idea*, Cambridge: The MIT Press.

- Jockers, Matthew (2013): *Macroanalysis: Digital Methods and Literary History*, Champaign: University of Illinois Press.
- King, Ross D./Whelan, Kenneth E./Jones, Ffion M./Reiser, Philip G./Bryant, Christopher H./Muggleton, Stephen H./Kell, Douglas B./Oliver, Stephen G. (2004): "Functional Genomic Hypothesis Generation and Experimentation by a Robot Scientist." In: *Nature* 427/6971, pp. 247–252.
- LeCun, Yann/Bengio, Yoshua/Hinton, Geoffrey (2015): "Deep Learning." In: *Nature* 521, pp. 436–444.
- Mach, Ernst (1976): "On Thought Experiments." In: *Knowledge and Error. Sketches on the Psychology of Enquiry*, Dordrecht: D. Reidel Publishing Company.
- McCarthy, John/Minsky, Marvin L./Rochester, Nathaniel/Shannon, Claude E. (1955): "A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence," August 31, 1955 (<http://raysolomonoff.com/dartmouth/boxa/dart564props.pdf>).
- McCulloch, Warren/Pitts, Walter (1943): "A Logical Calculus of the Ideas Immanent in Neuron Activity." In: *Bulletin of Mathematical Biophysics* 5/13, pp. 115–133.
- Minsky, Marvin (1974): *A Framework for Representing Knowledge*, MIT-AI Laboratory Memo 306, Cambridge: MIT.
- Minsky, Marvin/Papert, Seymour (1972): *Perceptrons: An Introduction to Computational Geometry*, Cambridge: The MIT Press.
- Moretti, Franco (2005): *Graphs, Maps, Trees: Abstract Models for a Literary History*, London: Verso.
- Newell, Alan (1982): "The Knowledge Level." In: *Artificial Intelligence* 18, pp. 87–127.
- Newell, Allen/Shaw, J.C./Simon, Herbert (1958): "Elements of a Theory of Human Problem Solving." In: *Psychological Review* 65, pp. 151–166.
- Newell, Alan/Simon, Herbert (1956): *The Logic Theory Machine. A Complex Information Processing System*, P-868, Santa Monica: The Rand Corporation.
- Rickert, Heinrich (1921): *Kulturwissenschaft und Naturwissenschaft*, Tübingen: J.C.B. Mohr.
- Rosenblatt, Frank (1957): *The Perceptron, A Perceiving and Recognizing Automaton*, 85–460-1, Buffalo: Cornell Aeronautical Laboratory.
- Rosenblatt, Frank (1958): "The Perceptron: A Probabilistic Model for Information Storage and Organization in the Brain." In: *Psychological Review* 65, pp. 386–408.

- Rumelhart, David/Hinton, Geoffrey/Williams, Ronald (1986): "Learning Representations by Back-Propagating Errors." In: *Nature* 323/6088, pp. 533–536.
- Samuel, Arthur L. (1959): "Some Studies in Machine Learning Using the Game of Checkers." In: *IBM Journal of Research and Development* 44, pp. 206–226.
- Simon, Herbert (1969): *The Sciences of the Artificial*, Cambridge: The MIT Press.
- Simon, Herbert/Newell, Alan (1958): "Heuristic Problem Solving: The Next Advance in Operations Research." In: *Operations Research* 6, pp. 1–10.
- Sorensen, Roy (1992): *Thought Experiments*, New York: Oxford University Press.
- Turing, Alan (1950): "Computing Machinery and Intelligence." In: *Mind* 59/236, pp. 433–460.
- Valiant, Leslie (1984): "A Theory of the Learnable." In: *Communications of the ACM* 27/11, pp. 1134–1142.
- Vapnik, Vladimir (1999): *The Nature of Statistical Learning Theory*, New York: Springer.
- Vaswani, Ashish/Shazeer, Noam/Parmar, Niki/Uszkoreit, Jakob/Jones, Llion/Gomez, Aidan N./Kaiser, Lukasz/Polosukhin, Illia (2017): "Attention Is All You Need." (<https://doi.org/10.48550/arXiv.1706.03762>).
- Weizenbaum, Joseph (1966): "ELIZA – A Computer Program for the Study of Natural Language Communication Between Man and Machine." In: *Communications of the ACM* 9, pp. 36–35.
- Winograd, Terry (1970): *Procedures as a Representation for Data in a Computer Program for Understanding Natural Language*, MIT AI Technical Report 235, Cambridge: MIT.

