International Journal of Innovative Studies in Sociology and Humanities (IJISSH)

ISSN 2456-4931 (Online) **www.ijissh.org** *Volume: 4 Issue: 10 | October 2019*

Computational Philosophy as Experimental Philosophy

Jordi Vallverdú & David Casacuberta

Philosophy Department, UniversitatAutònoma de Barcelona, 08193 Bellaterra (BCN), CATALONIA- Spain

Abstract: *Since the development of software such as The Logic Theory Machine in 1956 or The General Problem Solver in 1957 the relationship between computers and philosophy has been viewed as a formal, and theoretical one. Computers offered a way to implement and test theoretical philosophical models on language, thought, or methodology.This paper presents experimental philosophy as an alternative methodological framework to understand the relationships between computers and philosophy.*

Keywords: *experimental philosophy, simulations, proofs, computer, epistemology*

1. INTRODUCTION

Since the development of software such as The Logic Theory Machine in 1956 or The General Problem Solver in 1957 (Newell et al. 1958), the relationship between computers and philosophy has been viewed as a formal, and theoretical one. Computers offered a way to implement and test theoretical philosophical models on language, thought, methodology. Computers offered an heuristic (Newell et al. 1957), but the basis was a very well crafted formal model. This paper presents experimental philosophy as an alternative methodological framework to understand the relationships between computers and philosophy.

Experimental philosophy is usually associated with ethics. If we search “experimental philosophy” in Google Scholar we’ll get 16.100 results (all the following searches conducted on February the 24th, 2016). If we add the word “ethics” 5060 results show up, and if we substitute “ethics” by “philosophy of science” the count goes down to 3000, but still a very respectable number. So clearly there is room for an experimental philosophy of science. The number gets remarkably down however if we change “philosophy of science” by “machine learning”. However, if we search “philosophy of science” and “machine learning” we’ll get 4.450 hits. Clearly, machine learning is as relevant to philosophy of science as experimental philosophy is to ethics, but very few people are pursuing what the mixing of experimental philosophy and machine learning may lead to. So, this will be the second main aim of this paper, to show that machine learning can play a very relevant role to develop an experimental philosophy framework for philosophy of science.

In order to do so we need to go beyond the “life in the lab” mindset (Latour, B., & Woolgar 2013) and consider how a computer simulation can be considered also an experiment, following the e-science paradigm. Then we can see machine learning programs not just as formal mathematical systems but as an specific type of simulation that can be very relevant in philosophy of science.

The structure of the paper is as follows: this section presents the aims of the paper, section 2 describes the main protagonists of it: what is experimental philosophy and what is machine learning. Section 3 presents the philosophical background that helps to view experimental philosophy as relevant to develop innovative understanding of science and its methodological assumptions. Section 4 presents some specific working projects that combines machine learning and philosophy of science discussions and the final section is devoted to some prospective of what the future of such association may bring in.

2. UNDERSTANDING EXPERIMENTAL PHILOSOPHY AND MACHINE LEARNING

2.1 What is experimental philosophy

We can find historical uses of the terms “experimental philosophy”, like Henry Power (1664) *Experimental Philosophy in Three Books*, London: T. Roycrof (as well as Isaac Newton, Robert Boyle or John Locke, who were other authors also using this terminology), although this use is not directly related to our contemporary understanding and meaning. When 17th Century authors talked about experimental philosophy they were reacting against previous speculative philosophy, mainly Scholasticism, and requested a new research attention to natural physical events, from an observational perspective (remember the newtonian motto: ‘hypotheses non fingo’). Instead of it, today the increasing spreading and implementation of this idea is deeply connected with the interrelation between experimental approaches and philosophical ideas.

Here we can find ourselves faced to several conceptual challenges: first of all, what does is an experiment?; secondly, what is a philosophical problem? and, thirdly, what means and implies the idea of ‘experimental philosophy’, according to the two previous answers?. Let us to define these preliminary ideas:

a) experiments: philosophers of science have devoted long and deep efforts to define the notion of ‘experiment’. As was precisely conceptualized in Vallverdú (2014a), we can find several types of experiments:

Experiments	Material	Nature	*In vivo*
		Laboratory	*In vitro*
	Non-material	Computational	*In silico*
			In virtuo
		Thought (Gedankenexperiment)	*In mente*
	Hybrid	Computational-material	*In mixtura*

Experiments are, then, a vast range of different empirical practices that try to obtain knowledge from a set of controlled entities under controlled conditions. The notion of ‘empiric’ reminds us that is not just a formal procedure or rules processing activity: even in the case of computational simulations, experiments run into hypothesized real environments which have been quantitatively captured and reproduced into a discretized framework.

Thought experiments (Gedankenexperiment) are perhaps the weak link in this description: despite of the fact that authors like Mach (1905) or Camilleri (2012) have defended the epistemic value of experimental mental activities, their value must be framed into within the walls of propedeutic argumentations. Because sciences do not are based on mere conceptual debates, instead of some external-to-brain evidences (even considering them as mind conceptualizations), thought experiments are the weakest of all the possible experimental forms.

b) philosophical problems: is there any set of conceptual questions which should be considered as directly and intransferable philosophical problems? From a historical perspective, it could be considered a true statement, taking even into account that philosophy and thinking is not an universally unified process (Nisbet, 2003). For our purposes it’s enough to look at the contemporary state-of.the-art at academic researches to identify the outstanding topics in the field (a very rich and diverse, according to the several subspecialities). Anyhow, Philosophy still address some of the oldest problems of several traditions: how do we can know, how does mind works, which is the better ethical/political system, or which is the meaning of the whole existence, among others.

c) experimental philosophy: once clarified in previous sections the possible meanings of ‘experimental’ and ‘philosophy’, we could easily affirm that the range of possible interpretations of the concept are wide and diverse. By the way, we can assume that experimental philosophy is the implementation of philosophical problems into experimental practices or traditions, in a possible interdisciplinary cooperation, but always with the presence of philosophers into the research team. Otherwise, we could talk about the checking of some philosophical questions performed by other disciplines, which adapt and interpret these original problems from their own conceptual background.

Today, we can identify several areas in which experimental philosophy (sometimes short comed as X-Phi, XΦ) is performed: philosophy of mind (Systma, 2014), moral philosophy (Knobe & Nichols, 2008) or epistemology (Alexander & Weinberg, 2007, and Vallverdú, 2014b, in the case of formal logic). Due to its recent existence, several debates are surrounding this new approach, as have recently surveyed by Machery& O'Neill (2014) or Systma & Buckwalter (2016).

2.1.1 What is Experimental Philosophy of Science? (EPS)

As we've seen, there are several approaches to experimental philosophy, being under our current research interest those related with Philosophy of Science. Does exist something like Experimental Philosophy of Science (henceforth EPS)? Our answer is positive and can be rooted into two different but related domains: contemporary sciences, and computational sciences. Contemporary research is computationally embedded (Vallverdú, 2009; Casacuberta & Vallverdú 2014), and increasingly we are entering into the Fourth Paradigm of research (Hey, Tansley & Tolle, 2009). This new research paradigm is modifying the ways by which not only science is performed (Humphreys, 2004), but also affects humanities and, specially, philosophical activities. We will see that the possibility of execute computer simulations as well as big data analysis is radically transforming current and future academic researches.

2.2 What is a simulation.

In previous researches (Vallverdú 2009, 2014a), we've shown that the role of computational simulations has transformed completely the nature of scientific practices. Even in the case of some evident relationships between classic and computational paradigms, simulations are allowing a new epistemological approach to scientific and philosophical knowledge (Hartmann, 1996; Frigg & Reiss, 2009). With the growing and more fundamental use of these kinds of experiments, and with the increase in computer intensive techniques within scientific research, different attitudes towards their nature and scientific meaning have been voiced. Are they true experiments or just pseudo-empirical results used to generate new hypotheses and to further develop a theoretical model? What is their epistemic range and their ontological nature? In recent years several studies on philosophy and simulations have appeared (See Vallverdú, 2014a for more details), covering a broad spectrum of theoretical aspects related to simulations, such as the relation between models and simulations, simulations in physics or their epistemology. First of all, we should point out that the exact definition of Computer Simulations is a complex issue. The definitions vary according to the different areas of specialist expertise. We can find basic and repetitive concepts across all, or at least most, of these definitions: maths, models, computer, process and system. Therefore, our best and simplest definition of a simulation must be: "a mathematical model that describes or recreates computationally a system process". The key aspect of this new approach is to see that due to the computational power of these simulations, as well as to thanks to a conceptual shift towards the acceptation of their epistemic value, they are being implemented into several research domains. This allows us to make new formal and epistemological approaches to research procedures, even in the realm of social sciences (Axelrod, 1987, 1997) or cognitive sciences (Vallverdú et al 2016).

2.3 What are the main relationships between simulations and philosophy of science.

In the previous section we've seen that the computational turn has been widely implemented into a scientific research. Now, our intention is to elucidate whether this shift can also be traced into the practice of philosophers of science. Some decades ago, Paul Thagard (Thagard, 1988) said that that any philosophy science that aimed to have any merit should be realizable as a computer program. This not only affected classic sciences under update but also the philosophical practices as well as the essence of the computer sciences (Korb, 2004). Recent important advances in Physics, like Millennium simulation or the LIGO results on gravitational waves results have been supported by intensive computational simulations. Besides, contemporary philosophers are doing research on experimental philosophy of biology (Stotz, 2006; Callebaut, 201) or cognitive sciences (Fernández, 2003; McClelland, 2009), as well as they are thinking about the role of these new computer environments into the advanced of knowledge (Gustafsson & Vallverdú 2016). The implementation of big data mining has transformed even the practices of study of ancient philosophy, being employed into deep language analysis of data bases (Gentz & Meyer, 2015).

Obviously, there is an important epistemic debate: are simulations changing the way we understand scientific method? The answer must be affirmative, according to the previous ideas we've expressed. And, we need to make an analysis of the relationship between experimental philosophy and computer resources through a special field: machine learning, to be explained into next section.

2.4 What is machine learning

The classical reference for machine learning, Mitchell (1997) states that "A computer program is said to learn from experience E with respect to some class of tasks T and performance measure P, if its performance a tasks in T, as measured by P, improves with experience E". According to more recent and epic definition (El Naqa & Murphy

ISSN 2456-4931 (Online) www.ijissh.org Volume: 4 Issue: 10 | October 2019

2015) Machine learning is "an evolving branch of computational algorithms that are designed to emulate human intelligence by learning from the surrounding environment. They are considered the working horse in the new era of the so-called big data. Techniques based on machine learning have been applied successfully in diverse fields ranging from pattern recognition, computer vision, spacecraft engineering, finance, entertainment, and computational biology to biomedical and medical applications".

We could feel this whole paper with different definitions of what machine learning is, the types of algorithms used and so on, but we find this brief definition from Korb 2004 that simply states: "Machine learning studies inductive strategies as they might be carried out by algorithms." Next he points out that "The philosophy of science studies inductive strategies as they appear in scientific practice." We clearly see here the close connection between two enterprises that, at first sight, might look so distant. All definitions would agree that our main aim is to simulate some sort of learning process, on how an algorithm can get some data and process it to find some regularities. We can therefore view machine learning as a way to simulate the way scientists interrogate the world to get some raw data and develop hypotheses, laws and theories out of it.

Philosophy of science and machine learning share some interests, basicly they both want to know how based on raw data one can put some order over it, to develop some sort of inductive or abductive strategy to build a classification, find common patterns, discover new laws…Another point of connection is they are both interested in the relationship between theory and evidence. Is the model i am building based on this data the best one? What makes a theory a good one? Korb (2001) states: "The two disciplines are, in large measure, one, at least in principle. They are distinct in their histories, research traditions, investigative methodologies; however, the knowledge which they both ultimately aim is in large part indistinguishable." A very similar appreciation can be seen in one of the pioneers in computational philosophy, Paul Thagard, who in Thagard (1988) declared: "The branches of philosophy concerned with reasoning are continuous with psychology and artificial intelligence". Of course, there are several differences between philosophy of science and machine learning. First of all the two disciplines have very different aims: machine learning is mostly practical and want to develop models that can be used for predictions, and solve practical problems, while philosophy of science is both theoretical and normative: we don't just want to describe how scientists do work and model their trains of thoughts, we want to state when scientific research is correct, when it is ok to use an specific scientific method and when is not, and so no.

Second, in general terms, philosophy has some compromise about finding and arguing for truth. Machine Learning does not aim for truth, but to develop helpful instruments. One doesn't care if the rules discovered by a neural network are truth or false, only if they are good enough to get a meaningful task done.

Third, the interests differ as well: philosophy of science has a very broad field of interests (realism, theory versus experiment, measures, creativity and innovation in science, you name it) while machine learning is mostly focused in the problem between induction from data to relevant patterns, which is certainly a relevant scientific problem, but not the only one. Finally, philosophy of science analyses also questions that are not related to the pure logic from data to theory: like socio-political issues under the broad field of Science, Technology and Society studies.Nevertheless, the similarities are worth enough to analyze the role of machine learning as a subset of simulations -ergo a type of experiment as we argued in sections 2.1 and 2.3)- in order to develop an experimental philosophy of science.

3. PHILOSOPHICAL DISCUSSIONS

Although we could find historical connections between the birth of Artificial Intelligence and Philosophy (from logical positivism pioneers to some posterior evolutions within that field), it is a matter of fact that contemporary philosophy has embraced computational approaches, acquiring at the same time some external problems yet incorporated to the philosophical list of problems to be solved (Floridi, 2004). Some of them, like Turing Test, Gettier Problem, the Hard Problem of Consciousness, Automated proofs, or NP problems, among a long list, are still under debate (Shah et al, 2016; Vallverdú, 2011). But let us focus into the most recent debates on Big data related problems.

3.1 History of science as Philosophy's Big Data

Laudan, in his quoted and discussed paper Laudan (1987) argues against the "historical turn" started by Kuhn's *Structure of Scientific Revolutions* and Feyerabend's *Against Method* diatribes -as well as his own philosophical approaches during the 70s. The historical turn wanted to make mandatory to reconstruct rationality in past cases of science, in order to test philosophical models. Instead, Laudan here proposes that methodological rules shouldn't be read as categorical imperatives, but as hypothetical ones. That is, every methodological rule has some inarticulate element (Perry, 1993, 1998) that needs to be articulated in order to understand its rationality and use. So, states Laudan, Popper's methodological rule "avoid *ad hoc* hypothesis" is actually "if one wants to develop theories which are very risky, then one ought to avoid ad hoc hypotheses." (Laudan 1987, p. 24). When these contexts and goals are properly highlighted, we have more elements to study past scientific theories. Based on this concept, Laudan proposes a new role for history. There is no need for a meta-methodology, all that we need to know is to go to history of science, viewed as an inductive database in order to decide which method would be best for an specific scientific inquiry. Does avoid *ad hoc* hypothesis really lead to riskier theories? Having riskier theories does help a specific discipline to progress further? According to Laudan, what makes science a relevant subject for philosophy is not its more rational ways, but the big data framework that it provides: "Hence, the history of science has to be reckoned with, not because scientists are always or more often rational than anyone else (I rather doubt that they are), but rather because the history of science -unlike that of many other disciplines- offers an impressive record of actions and decisions moving closer through time to a realization of ends that most of us hold to be important and worth" (Laudan 1987 p. 28). Of course, Laudan doesn't ever talk about "machine learning" or even "computers" but if philosophy of science is considered as a empirical research to state whether a disciplina progresses or not, based on an inductive analysis of the history of the discipline, then all the observations made in section 2.5 make perfect sense and we can see the relevance of a machine learning experimental approach towards philosophy of science.

3.2 Simplicity versus overfitting

Simplicity is a key category to understand science. Einstein always had simplicity as one of his main heuristic guides in order to develop his theories (Norton 2000), and a similar feeling transpires in Galileo's platonism (DeCaro 1993). Sober attempted to show in Sober (1975) or Thagard (1988) how the simplicity of a scientific theory can actually be measured, Prigogine's description of simplicity in mathematical models as a way to give a novel understanding of order and chaos (Prigogine 1984). There are also several thinkers planning to debunk simplicity since at least Bunge (1965) and Goodman (1972).

Machine learning offer us an alternative view of the problem via the concept of "overfitting". Let's consider for example decision trees, one of the oldest algorithms in machine learning. In decision trees we develop rules to make several branching when two sets cannot be linearly divided. Imagine we want to predict whether an specific person will go to a given day to the gym or not. After some analysis we find that the relevant variables are whether it is holidays or not and whether it is sunny or cloudy. Our protagonist doesn't go to the gym on holidays if it is sunny -she prefers to practice outdoors-, but goes if it is cloudy, and doesn't go to the gym on weekdays if it is cloudy (after working, if it is cloudy she feels a little depressed and prefers to go home, but does go if it is sunny (getting outside and feel the sun gives her enough energy to make her go to the gym. If we analyzed this -very simple- problem using a decision tree, we would get an algorithm performing this:

Consider the same problem, but with a lot less regular distribution of going and not going to the gym, generating a graphical representation like the following:

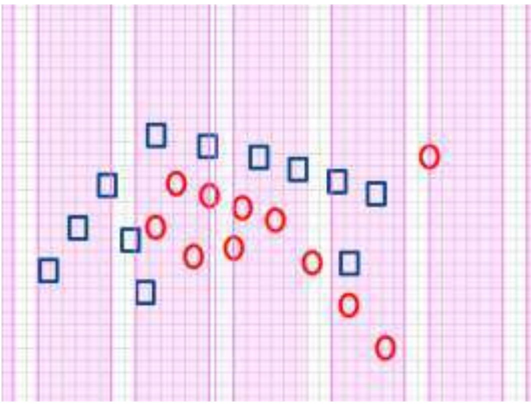

This problem is not linear, there is not a simple line that could clearly cut between the red circles (going to the gym) with the blue squares (not going). The way the algorithm works is dividing the learning set in groups following a decision tree. Let's say we are processing 250 different entries, related to 250 different days of our protagonist, stating some of the relevant characteristics. We divide the 250 examples in a binary tree of 150 and 100, the 100 branch is divided in two new branches of 60 and 40, and so on. In this algorithm, one of the parameters we can play with is called minimum samples split, that is the minimum number of nodes we wish to have in the final level of the decision tree. If we used a relatively big minimum sampling -for example 20- we could get just two lines:

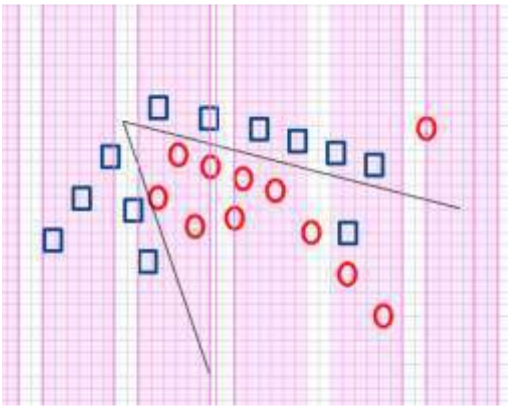

Here, one of the blue rectangles and one of the red circles is out ot the classification, but we have a model that is a lot more generalizable and be used with other input. However if we ask for a low minimum samples split we are in the risk of overfitting, that is, finding a classification that follows so closely the examples that is not capable of generating any useful rule, and therefore the proposed solution is not applicable to any other set. If we overfitted our former example, we could get a graphical solution like this one:

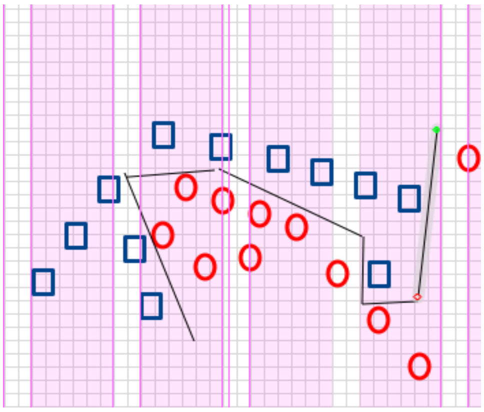

We can see several *ad hoc* movements in order to adapt to the specific distribution of the current examples, even if one of the circles and one of the squares could be just outliers. As we can see there is a relevant relationship between overfitting and simplicity. The simpler the mathematics, the more capability to generalize and therefore find a more useful and relevant solution. To a certain extent, of course, If a solution is too simple it might end up irrelevant as well. If we pay too much attention to the data we might end up developing a very *ad hoc* model that is irrelevant to the problems we want to solve.

So overfitting and simplicity work both as a tradeoff, one in machine learning and the other in science. When looking for a theory to explain a phenomena wants need to get the right amount of simplicity.

3.3 Learning, meta-learning and philosophy of science

Since its very beginning, with the rise of neural networks, one of the main discussions in the field of machine learning is the characteristics of a given method. What are their limits? To which problems does it work best? How can I avoid overfitting? A computer scientist while developing an specific application to detect card fraud, may wonder whether decision trees, naÏve Bayes, deep learning or SVM (Support Vector Machines) is the best selection, and why. This can be based on the intuitive experience of the coder, trusting some frugal heuristic recommended by a former teacher or a rigorous mathematical analysis. But one can also develop meta-learning algorithms in which different algorithms are checked in order to see which one will seem most promising.

It is easy to see the equivalence with philosophy of science. Different philosophers present different “algorithms” describing how scientists are supposed to work, and then they discuss the main benefits and flaws. Induction versus the hypothetico-deductive model; a frequentist analysis of probability versus a bayesian one. All these discussions are similar to the machine learning problem, in the sense that is is comparing and deciding which methods are best fit; in the case of machine learning, to develop a helpful model to make some predictions or calculations. In the case of philosophy of science to state the relative benefits -or even its rationality. The similarities are a lot more relevant when one considers Laudan’s observation described in section 3.2, where Laudan views philosophy of science as analysing the “big data” of history of science and test different scientific methods against different temporal and conceptual backgrounds. A similar position is defended in Giere (2003, 2005) in which whether a scientific theory should be considered true or false is not based on some abstract measure of rationality, but in checking its premises and formulation against the scientific context on what was considered truth in the time the theory was presented. The way Giere discusses the relevance of different understandings of the scientific method while explaining why Wegener’s theory of Pangea was first rejected and some decades later accepted is not that different from the way machine learning engineers discuss the benefits and flaws of a certain algorithm given a context.

3.4 The Statistical Debates, the Computational Era and Philosophy

Once scientific fields started to quantify reality and work with sets of numerical data, emerged the necessity to make statistical approaches to reality, especially in the area of astronomy (Vallverdú, 2016). An intense debate between statistical strategies (mainly frequentist versus Bayesianism) emerged, and with the introduction of computer facilities and the exponential growth of databases, this debate shifted towards a pragmatic approach, although under the general influence of Bayesian practices (Pearl, 2000). Automated discoveries, computational creativity and knowledge discovery fields are deeply influenced by Bayesian methods, which are very far from a subjectivist approach to reality. NARS (Non-Axiomatic Reasoning System) or NAL (Non-Axiomatic Logic) is based on these new statistical approaches, as well as the result of the implementation of non-monotonic logics and reasoning to several domains of academic research, making possible the study or more complex problems (Korb, 1992) and to find solutions to classic ones, like the frame problem. Obviously, some of these simulations are not true experiments, but contribute as a tools for most of them (Beisbart & Norton, 2012). Anyhow, new mixed ways of reasoning (from deductive to abductive) are statistically supported by new computational resources.Although the classic debates about the nature of causality or inference are not solved with general agreement, a new set of computational facilities is making possible the statistical processing of more kinds of data, something that is making possible to find more information amon the vast fields or raw data. DARPA’s ‘Big Mechanism’ project, for example, is possible tanks to a statistical approach to causality, something only possible thanks to new computational scientific environments (Cohen, 2015). The detailed and honest analysis of these real practices should force epistemologists to reconsider the mechanisms by which contemporary sciences obtain knowledge, taking into account the eminent role of computational methods (which include a long array of formal methods and instruments inside computational systems).

4. CASE EXAMPLES

Another indication of the relevance of machine learning for philosophy of science is e-science, or how algorithms are starting to substitute scientists not in mere processes of data collection and brute force processing, but also in accomplishing complex recognition of patterns. like an algorithm that was able to recognize the signature of subatomic particles better than specific software developed by the physicists at CERN even if the authors of the deep learning algorithm didn't know anything about particle physics (Akil, 2015). We can divide these efforts in two main classes, according to its relevance to philosophy: (a) first we have software that tries to mimic scientific discoveries by trying to emulate the reasoning developed by scientists in the past. The aim here is to introduce in the software a simulation on how the scientific method works and see if the program can reach a conclusion similar to those scientists in the flesh.One of the first examples of such a direction is BACON, described in Langley (1979) and Langley et al (1987). BACON, is one of the first examples of software able to discover some natural laws based on observational variables, and achieving results akin to those obtained by scientists in the past. BACON uses several statistical formula akin to what one can see in machine learning, such as linear regression or dimension reduction. Bacon.3 in 1979 was able for example to generate by induction the ideal gas law based on raw data obtained from a lab. It also got pretty close to Kepler's third law, realizing the role that a square in a linear combination plays in the law. The other main class is about (b) scientific creativity; that is the attempt to introduce knowledge about scientific method to a computer program so it can generate by itself new knowledge, that later can be tested and checked by human scientists. Because these proposal of new knowledge can actually be tested by human scientists, these computer simulations also work as empirical tests of philosophical theories about what is the scientific method and how it works. If a computer system is able -using a certain simulation of the scientific method- to present hypothesis that scientists find relevant, then the philosophical proposal embedded in the software has epistemic value. An early case of such direction is Cheeseman et al (1989) which used a Bayesian network analysis system called AUTOCLASS to produce new groupings of stellar objects using the Infrared Astronomical Satellite. The program was able to produce 77 stellar classes which then the human scientists arrange in more complex clusters, and the results were finally published in a referred astronomical journal. A very impressive example of that is ADAM, the robot scientist. Here we don't have just software simulation but a real autonomous agent, a robot that both generates functional genomic hypotheses about a type of yeast, the *Saccharomyces cerevisiae*, and can test those hypotheses using real lab equipment. (King et al 2009).

Even more relevant to our research, and one would say that to science in general, is the automated procedures that allow machine learning algorithms to predict future scientific results -and therefore to give orientation to scientists on where to look while searching for a new model or hypothesis- based on a computational analysis of scientific literature, as described in Nagarajan et al. (2015). Nagarajan and his team have developed KnIT (Knowledge Integration Toolkit). KnIT does not require any pre-processing or labeling of the papers; it has direct access to the documents in digital format and makes a complex pattern matching and knowledge discovery process without any type of human intervention. Its area of expertise is protein interaction, and is able to predict previously unknown protein-to-protein interactions. It is also important to note that we are not talking about mere scholarly knowledge, protein interaction is an applied field of research and it is used to drug design and understanding the precise mechanisms of several illnesses. KnIT detects in the literature direct and indirect references of new discoveries on protein interaction, and then checked that new knowledge with what already is in its database and presents new conjectures to be tested.

Knit combines machine learning with symbolic processing, starts from basic knowledge curated by humans but then starts to build its own database of knowledge, and its results have been tested by human scientists that acknowledged its ability to predict new scientific progress.

5. FUTURE RESEARCH AND END REMARKS

We have presented a comprehensive understanding of experimental philosophy and machine learning, showing how this field could be very relevant to innovate in philosophy of science. After reading the examples and case

ISSN 2456-4931 (Online) www.ijissh.org Volume: 4 Issue: 10 | October 2019

studies presented in section 4, it is normal to wonder if machine learning could be applied to philosophical problems. Could we one day give epistemological problems about the nature of causality or whether induction works to a machine learning system and get some surprising, relevant answers? If ADAM can generate hypothesis and then check them in a real lab, and KnIT is able to predict protein-to-protein interaction that no human scientist have seen so far, could we have something similar in philosophy?

This is certainly an original and exciting research line, and we are sure we'll see some improvements in recent years, however there are also important obstacles.

Both ADAM and KnIT move in a very precise and well defined domain: a specific microorganism or a very specialized research line such as protein interaction. This implies that there are relatively less variables or characteristics to consider, the relevant literature is relatively easy to locate and classify and the expert knowledge used as a basis can be obtained in a systematic fashion from human experts

Philosophy of Science is a much more broad subject, that implies knowledge in lots of different subjects and areas, as well as a general understanding of human nature. If, for example we give now a machine learning algorithm a selection of the main papers of the philosophical discussion on whether there is progress in science, from Popper to Laudan, we would either obtain no generalization at all, because there are too many characteristics implicated, or a very overfitted system in which the system could do little less that recognizing Popper's style of writing versus Lakatos' one.

Nevertheless, machine learning techniques advance and are able to tackle problems that seem impossible a decade ago. Consider for examples the recent advances with deep learning by Google in fields like image or language recognition (LeCun et al 2015) or the game of Go (Silver et at 2016). It is only a question of time in which a whole new paradigm of experimental philosophy will arise: using machine learning techniques to innovate in classical philosophical debates in epistemology, ethics, metaphysics, aesthetics, history of philosophy...

REFERENCES

Akil, A. (2015). Top Tagging by Deep Learning Algorithm Ali Akil. Cern Report. Source: https://cds.cern.ch/record/2054120/files/Report%20CERN%20Akil%20Ali.pdf, accessed, February 26th, 2016.

Alexander, J. & Weinberg, J.M.(2007) "Analytic Epistemology and Experimental Philosophy", *Philosophy Compass*, 2(1): 56–80.

Axelrod, R. (1987) The Evolution of Strategies in the Iterated Prisoner's Dilemma. in G*enetic Algorithms and Simulated Annealing*. London: Pitman, 32–41.

Axelrod, R. (1997) *The Complexity of Cooperation.* Princeton: Princeton University Press.

Beisbart, Claus & John Norton, 2012. "Why Monte Carlo Simulations are Inferences and not Experiments." International Studies in Philosophy of Science 26:403–422.

Bunge, M (1965) "The Myth of Simplicity: Problems of Scientific Philosophy", Philosophical Review 74 (3):402-404 (1965).

Callebaut, W. (2012). Scientific perspectivism: A philosopher of science's response to the challenge of big data biology. *Studies in History and Philosophy of Biological and Biomedical Science 43(1)*, pp. 69-80.

Camilleri, K. (2012) "Toward a Constructivist Epistemology of Thought Experiments in Science", *Synthese*, 191: 1697–1716.

Casacuberta, D. & Vallverdú, J. (2014) "E-Science and the data deluge", *PhilosophicalPsychology*, 27(1): 126-140.

Cheeseman, P., Goebel, J., Self, M., Stutz, M., Volk, K., Taylor, W. & Walker, H. (1989). Automatic Classification of the Spectra from the Infrared Astronomical Satellite (IRAS). Washington, DC: National Aeronautics and Space Administration. (Reference Publication 1217).

Cohen, P.R. (2015). DARPA's Big Mechanism program. *Phys Bio, 12(4):* 0405008.

De Caro, M. (1993). Galileo's mathematical platonism. Philosophy of Mathematics, Part, 1,13-22.

El Naqa, I., & Murphy, M. J. (2015). What Is Machine Learning?. In Machine Learning in Radiation Oncology (pp. 3-11). Springer International Publishing.

Fernández, J. (2003). Explanation by Computer Simulation in Cognitive Science. *Minds and Machines, 13(2),* pp 269-284.

Floridi, Luciano (ed.) (2004) *The Blackwell Guide to the Philosophy of Computing and Information*, Malden (MA): Blackwell Publishing Ltd.

Frigg, Roman & Julian Reiss. 2009. "The Philosophy of Simulation: Hot New Issues or Same Old Stew?" Synthese 169(3):593–613

Hartmann, Stephan, 1996. "The World as a Process: Simulations in the Natural and Social Sciences," in R. Hegselmann, et al. (eds.), Modelling and Simulation in the Social Sciences from the Philosophy of Science Point of View, Dordrecht: Kluwer, 77–100.

Gentz, J. & Meyer, D. (Eds) (2015). *Literary Forms of Argument in Early China*. Boston: Leiden.

Giere, Ronald (2003), The role of computation in scientific cognition. *Journal of Experimental & Theoretical Artificial Intelligence,* 15, pp. 195-202.

Giere, Ronald (2005), "How Models Are Used to Represent Reality". *Philosophy of Science*, 71: 742-752.

Goodman, N. (1972) Problems and Projects. Indianapolis: Bobbs-Merrill.

Gustaffson, C. & Vallverdú, J. (2016) "The best model of a cat is several cats", *Trends in Biotechnology*, DOI:http://dx.doi.org/10.1016/j.tibtech.2015.12.006

Hey, T. Tansley, S. & Tolle, K. (Eds.) (2009). *The Fourth Paradigm: Data-Intensive Scientific Discovery*. USA: Microsoft research.

Humphreys, Paul. 2004. Extending Ourselves. Computational Science, Empiricism, and Scientific Method. Oxford: Oxford University Press.

King, R. D., et al. (2009). "The automation of science." *Science*, 324(5923), 85-89.

Knobe, J. & Nichols, S. (eds.) (2008) *Experimental Philosophy*, Oxford: OUP.

Korb, K.B. (1992). *A Bayesian Platform for automating Scientific Induction*, Ph.D. dissertation. USA: Indiana University.

Korb, K. (2001) "Machine learning as philosophy of science" in K. Korb and H. Bensusan, eds. Proceedings of the ECML-PKDD-01 Workshop on Machine Learning as Experimental Philosophy of Science, Freiburg.

Korb, K (2004). Introduction: Machine learning as philosophy of science.*Minds and Machines, 14(4),* pp. 433-440.

Langley, P. (1979, August). Rediscovering Physics with BACON. 3. In IJCAI(pp. 505-507).

Langley et al (1987) Scientific Discovery: Computational Explorations of the Creative Processes, Cambridge (Mass): MIT Press

Latour, B., & Woolgar, S. (2013). Laboratory life: The construction of scientific facts. Princeton University Press.

Laudan (1987) "Progress or rationality? The prospects for normative naturalism". American Philosophical Quarterly, 24:19-31.

LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. Nature, 521(7553), 436-444.

Mach, E. (1905) "Über Gedankenexperimente", in: Mach, Ernst (1905): Erkenntnis und Irrtum, Leipzig: Verlag von Johann Ambrosius Barth, pp.181-197.

Machery, E. & O'Neill, E. (2014) *Current Controversies in Experimental Philosophy,* Uk: Routledge.

McClelland, J.L. (2009). The Place of Modeling in Cognitive Science. *Topics in Cognitive Science 1,* pp. 11–38

Mitchell, T. Machine Learning. McGraw-HIll, 1997.

Nagarajan, M., et al (2015). Predicting Future Scientific Discoveries Based on a Networked Analysis of the Past Literature. In Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (pp. 2019-2028). ACM.

Newell, A., Shaw, J. C., & Simon, H. A. (1957, February). "Empirical explorations of the logic theory machine: a case study in heuristic." In Papers presented at the February 26-28, 1957, *Western joint computer conference: Techniques for reliability* (pp. 218-230). ACM.

Newell, A., Shaw, J. C., & Simon, H. A.(1958). "Elements of a theory of human problem solving". *Psychological review*, 65(3), 151.

Nisbet, R.E. (2003) *The Geography of Thought*, NY: The Free Press.

Norton, J. D. (2000). Nature is the realisation of the simplest conceivable mathematical ideas': Einstein and the canon of mathematical simplicity. Studies In History and Philosophy of Science Part B: Studies In History and Philosophy of Modern Physics, 31(2), 135-170.

Pearl, Judea. 2000. Causality. Models, Reasoning, and Inference. Cambridge: Cambridge University Press

Perry, S.R. (1993). *The problem of essential indexical and other essays*. Oxford: Oxford University Press

Perry, S.R., "Hart's Methodological Positivism" (1998). Faculty Scholarship. Paper 1136. http://scholarship.law.upenn.edu/faculty_scholarship/1136

Prigogine, I. (1984), Order Out of Chaos: Man's New Dialogue with Nature. Random House.

Shah, H., Warwick, K., Vallverdú, J. & Wu, D. (2016) "Can Machines Talk? Comparison of Eliza with Modern Dialogue Systems", Accepted for publication in *Computers in Human Behavior*: 58: 278–295.

Silver, D., et al. (2016). Mastering the game of Go with deep neural networks and tree search. Nature, 529(7587), 484-489.

Sober, E. (1975) *Simplicity*. USA: Clarendon Press.

Stotz, K. (2006). With 'Genes' Like That, Who Needs an Environment? Postgenomics's Argument for the 'Ontogeny of Information'. Philosophy of Science, 73(5): pp. 905-917.

Systma, J. & Buckwalter, W. (Eds.) (2016) The Blackwell Companion to Experimental Philosophy, USA: Blackwell.

Sytsma, J. (Ed.) *Advances in Experimental Philosophy of Mind*, NY: Bloomsbury

Thagard, P. (1988). *Computational Philosophy of Science*. Cambridge, MA: MIT Press/Bradford Books..

Vallverdú, J. (2009) "Computational Epistemology and e-Science. A New Way of Thinking", *Minds and Machines*, 19(4): 557-567.

Vallverdú, J. (2011) "Patenting Logic, Mathematics or Logarithms? The Case of Computer-Assisted Proofs", *Recent Patents on Computer Science*, 4(1): 66-70.

Vallverdú, J. (2014a) "What are Simulations? An epistemological approach", *Procedia Technology*, 13: 6 – 15.

Vallverdú, J. (2014b) "Logic and Computers. A Sketch of Their Symbiotic Relationships", *Kairos. Journal of Philosophy & Science*, 9: 45-72.

Vallverdú, J. (2016) *Bayesians Versus Frequentists. A Philosophical Debate on Statistical Reasoning*, Germany: Springer. ISBN: 978-3-662-48636-8.

Vallverdu, J. et al(2016) "NEUCOGAR: A Neuromodulating Cognitive Architecture for Biomimetic Emotional AI", *International Journal of Artificial Intelligence (IJAI)*, 14(1): 27-40.

Williamson, J. (2004) "A Dynamic Interaction Between Machine Learning and the Philosophy of Science", *Minds and Machines*, 14: 539-549.