

Automating Science

Philosophical and Social Dimensions

Ross. D. King, Vlad Schuler Costa, Chris Mellingwood, and Larisa N. Soldatova

Clark Glymour argued in 2004 that “despite a lack of public fanfare, there is mounting evidence that we are in the midst of ... a revolution – premised on the automation of scientific discovery” [1]. This paper highlights some of the philosophical and sociological dimensions that have been found empirically in work conducted with robot scientists – that is, with autonomous robotic systems for scientific discovery. Robot scientists do not bring definite answers to the discussed questions, but rather provide “proofs of concept” for various ideas. For example, it is not that robot scientists solve the realist/antirealist philosophical debate, but that when working with robot scientists one has to make a philosophical choice – in this case, to assume a realist view of science. There are still few systems for autonomous scientific discovery in existence, and it is too early to generalize and propose new theories. However, being “in the midst of ... a revolution” it is important for the research community to re-examine views pertinent to scientific discovery. This paper highlights how experience with robot scientists could inform discussions in other disciplines, from philosophy of science to computer creativity research.

Scientific Discovery and Robot Scientists

The branch of Artificial Intelligence (AI) devoted to developing algorithms for acquiring scientific knowledge is known as “scientific discovery.” The pioneering work in scientific discovery was the development of learning algorithms for analysis of mass-spectrometric data [2]. In the subsequent 50 years, much has been achieved and there are now convincing examples in which computer programs have made explicit contributions to scientific knowledge (e.g., [3], [4], [5], [6]). However, the general impact of such programs on science has been limited. This is slowly changing as the expansion of automation in science is making it increasingly possible to couple scientific discovery software to laboratory instrumentation [6], [7], [8], [9].

Science is an excellent testbed for the development of AI discovery systems:

- Scientific problems are abstract, but also involve the real-world level knowledge.

- Scientific problems are restricted in scope – no need to know about “Cabbages and Kings” – and are also extensible.
- Science assumes that Nature is not trying to deceive us, so there is no need to consider malicious agents.
- Scientific knowledge is a public good when it is openly available.
- Science is a worthy object of our study.

A Robot scientist is an example of such an AI discovery system. The Robot scientist is a physically implemented laboratory automation system that exploits AI techniques to execute cycles of scientific experimentation [8], [9], [11]. A robot scientist automatically originates hypotheses to explain observations, devises experiments to test these hypotheses, physically runs the experiments by using laboratory robotics, interprets the results, and then repeats the cycle. The advent of robot scientists is of significant philosophical and social interest. They are also of practical interest as they have the potential to increase the productivity of science: they can work cheaper, faster, more accurately, and longer than humans, and can also be more easily multiplied.

The robot scientist “Adam” was the first machine to autonomously discover scientific knowledge, i.e., to autonomously both form and experimentally confirm novel hypotheses [8]. Adam worked in the domain of yeast functional genomics, and autonomously both generated functional genomics hypotheses about the yeast *S. cerevisiae*, and experimentally tested these hypotheses by using laboratory automation. Adam’s conclusions have been manually confirmed using gold standard experiments.

The robot scientist “Eve” (Figure 1) was designed to make drug discovery more economical, specifically for neglected tropical diseases [9]. Eve integrates and automates library-screening, hit-confirmation, and lead generation through cycles of quantitative structure activity relationship learning and testing. Using econometric modeling, Eve was shown to economically outperform standard drug screening. Eve has repositioned several drugs against specific targets in parasites that cause tropical diseases. One validated discovery is that the anti-cancer compound TNP-470 is a potent inhibitor of dihydrofolate reductase from the malaria-causing parasite *P. vivax*.

The Metaphysics of Robot Scientists

A major motivation for the automation of science is philosophical: if a mechanism can be built that is judged to have discovered some novel scientific knowledge, then this will shed light on the nature of science. To quote Richard Feynman “What I cannot create, I do not understand” (on his blackboard at the time of his death). The advantage of this approach to the philosophy of science, compared to traditional ones, is that it is constructive and objective. In building robot scientists one is confronted with the need to make concrete engineering decisions that relate to a number of important problems in the philosophy of science: the relation between abstract and physical objects, the nature of truth, the relation between observed and theoretical entities, the origin of hypothesis, the problem of induction, etc. This approach to science is analogous to the AI approach to understanding the human mind through the creation of artifacts that can be empirically shown to have some of the attributes of human minds [12].

We argue that the software/hardware isomorphism is the key to bridging the physical/abstract dichotomy in the metaphysics of science (Figure 2). The key to the power of a computer is that computers implement abstract programs in physical devices. This is the insight that distinguished Turing from the other great logicians of his time. Although the idea that a physical object can be isomorphic with an abstract system is at least as old as the abacus, a Universal Turing Machine is a uniquely powerful physical/abstract device.

In a scientific investigation, to relate corresponding abstract and physical entities requires the concept of “truth.” Within philosophy there are a number of competing theories of truth, including: correspondence, pragmatism, verification, and coherence. These theories are associated with different ontologies: correspondence theories with realism, and pragmatism, verification, and coherence with idealism, anti-realism, or relativism [14].

A robot scientist’s physical effectors (laboratory robots) can test the truth or falsehood of an abstract scientific proposition by specific physical experiments: an *Abstract entity* of type *Proposition* is assigned a truth value by a *Physical entity* that participates in a specified *Process*. This is achieved through the designed isomorphism between an abstract *Denotation rule* and a physical *Denotation process* (Figure

2). This operational approach to truth does not discriminate between correspondence, pragmatism, verification, or coherence theories of truth. For a human scientist these different approaches may possibly inspire different ways of doing science, but given the current state of development of robot scientists it is unclear to us whether there is any operational difference between these approaches.

One of the most debated questions in the philosophy of science is that between realism and anti-realism. Realism is “the viewpoint that accords to the objects of knowledge an existence that is independent of whether anyone is perceiving or thinking about them” [14]. The alternative position regards the existence of the real world as a metaphysical question that cannot be answered, and regards scientific theories as instruments of prediction [15], [16]. As physical devices, robot scientists necessarily adopt a realist position as defined above. However, their approach to determining the truth of propositions is also consistent with that of anti-realism. Therefore, with robot scientists there would seem to be no difference in regarding scientific theories as descriptions of reality or as tools for prediction. This approach is related to quietism [17].

The realism/anti-realism debate is closely connected to another area of interest in the philosophy of science that is important in the design of robot scientists: the relationship between observed and theoretical entities. This subject has long been a matter of debate in the philosophy of science, with some philosophers claiming that the distinction is not real and/or important [16]. We argue that the distinction clarifies the robot scientist's reasoning, and that what are observed and theoretical entities is relative to defined instrumentation.

A common view in the philosophy of science is that hypothesis formation necessarily requires human creativity [10]. This view has long been challenged by AI, [e.g., 19]. Most work within scientific discovery has focused on automated hypothesis formation [e.g., 3]. Within the philosophy of science, hypothesis formation has been closely associated with induction. This focus may be due to much philosophy of science being physics focused [16]. In modern biology most hypothesis formation is abductive [11]. What are hypothesized are factual relationships between entities, e.g., that the gene named *YBR060c* encodes for the enzyme named *chorismate mutase*, that the gene named *YPR060c* encodes a protein with a four-helix bundle topology, etc. Such relationships are factual rather than general laws.

Robot scientists follow a hypothetico-deductive methodology. Hypotheses are formed either using abduction or induction. The experimental consequences of these hypotheses are then deductively inferred, and then physical experiments are conducted to observe what causally happens in the real world. Adam used abduction to form hypotheses. A set of models is generated, each with different abducted propositions. With the model (T) these propositions (H) enable the deduction of whether growth is predicted (O) for a particular experiment, i.e., $T \wedge H \models O$. These deductions are monitored by a meta-logical program that determines the truth or falsehood of the abstract *theoretical growth* proposition in the various models [19]. This is then integrated with physical effectors to physically execute an experiment and thereby determine whether actually growth occurs or not, which this can be mapped to the robot scientists abstract model of reality.

Eve uses induction to form hypotheses. To select compounds to test its hypotheses, Eve uses active learning [9], [20]. The active learning task is comparable to that in many other areas of science and engineering: identify or design artifacts that have optimal performance. However, it has an extra ingredient reminiscent of reinforcement learning: balancing the exploration of compound space with the exploitation of regions of highly active compounds.

Robot and Human Scientists

Is Science Solely a Human Activity?

It is easy to find evidence that science should be viewed as solely a human endeavor [e.g., 21]. Within the philosophy of science there are many advocates of a humanistic understanding of science. However, developments in AI question the centrality of human creativity in the creation of scientific knowledge. Although most philosophers of science seem not to have engaged with the possibility of automating science, the views of certain philosophers would appear to infer that science, as a set of practices would be very difficult to automate. Among these are probably the two best known post-war philosophers of science: Karl Popper and Thomas Kuhn (of course, there are many more authors who have relevant writings on the topic, e.g., [22]).

We are unaware if Popper ever directly tackled the question of whether a machine could be engineered to do science. However, it would seem reasonable to infer from his other clearly expressed views – that hypothesis formation requires human creativity, and that induction is a myth – that he would have denied the possibility of mechanizing scientific discovery [10].

In his postscript to *The Structure of Scientific Revolutions* (1970) Thomas Kuhn responds to critics who viewed his work as relativistic [23]. Kuhn uses his account of scientific theories to argue that science is a special case. This is because viewing proponents of competing scientific theories as simply akin to members of different language-communities does not account for scientists as fundamentally puzzle-solvers [23, p. 205]. Kuhn is an advocate of scientific progress and believes theories and paradigms each build upon that which has come before. Similarly, Kuhn was open to the use of computer programs in scientific knowledge making, and uses this fact to display his willingness to accept the importance of rule-following in science [23, p.191]. However, Kuhn's larger point is that it is not quite right to say that scientists do not follow any fixed rules. He is arguing that scientists follow rules based on previous exemplars from their field (see "A Sociological Perspective" section below). Rules cannot be abstracted from exemplars and take their place [23, p.192], and it is prior experience and training that shapes how scientists judge when rules are being followed or not [23, p. 198]. Applied to robot scientists the question remains open as to how prior experience and training become embedded in the rule following of AI programs.

If one accepts that robot scientists can automate many of the steps in the generation of scientific knowledge, then there would appear to be two main "get-outs" that would still enable philosophers to maintain that science cannot be automated. One get-out is that a current robot scientists is not aware that it is doing science, and is therefore the robots are not really doing science [24]. This type of argument would also apply to chess computers [25] – yet we are unaware of any philosopher who has argued that chess computers do not really play chess.

The other get-out is to deny that what was "discovered" was really novel science. For example, it could be argued that the new scientific knowledge was implicit in the formulation of the problem, and is therefore not novel. The argument that computers cannot originate anything and can only do what they are programmed to do is known as "Lady Lovelace's objection" [26]: "The Analytical Engine has no pretensions to *originate* anything. It can do *whatever we know how to order it to perform*" (Lady Lovelace's italics). This argument has real teeth as the robot scientists Adam and Eve are very far from being autonomous agents seeking out their own scientific problems. One counter-argument is that robot scientists are programmed to learn novel scientific knowledge, and as they learn from observations of the physical world the conclusions are not purely deductive.

A variant of the lack of novelty argument is the argument that computers will only ever be able to do "normal science," i.e., within a paradigm, and will never be able to do "revolutionary science" [23], [27]. Certainly existing robot scientists are not capable of doing revolutionary science. However, very few human scientists are either.

There is also an argument that such systems as robot scientists are not truly autonomous, rather that the systems merely are tools of scientists, and that there is always a human-in-the-loop [28]. The existing robot scientists are compliant with the definition of autonomous robots by IEEE P1872TM/D3 standard "a robot performing a given task in which the robot solves the task without human intervention, while adapting to operational and environmental conditions." While one may argue that there are certain shortcomings in Eve's autonomy, the concept of a robot scientist, as originally introduced [11], implies a complete autonomy in scientific discovery.

An Anthropological Perspective

Compared to human scientists, robot scientists have a mixture of super- and sub-human abilities. Laboratory robots have traditionally been used to automate low-level repetitive tasks. Robot scientists inherit this ability and have the super-human capacity to work flawlessly on extremely repetitive tasks for days at a time. In comparison humans perform badly at repetitive tasks, especially those carried out during extended periods [29]. We have confirmed this during our observational studies of human scientists, who routinely make mistakes, particularly when subject to hindrances like stress, time pressure, or distractions. Robot scientists inherit from AI abilities that have traditionally been regarded as high-level for humans, such as a super-human ability to do logical and probabilistic reasoning. However, robot scientists are sub-human in their adaptability and understanding, and human scien-

tists are still unequalled in conditions that require flexibility and dealing with unexpected situations, especially those intuitive functions that might have otherwise been considered low level [30].

Given the mixture of super- and sub-human abilities of robot scientists, it is informative to investigate how human scientists cooperate with their robot counterparts, both to improve the technology by playing to the strengths of human and robot scientists, and to better understand future working relationships between humans and automation. These relationships occur at many levels: from the most profound (deciding on what to investigate, structuring a problem for computational analysis, interpreting unusual experimental results, etc.), to the most mundane (cleaning, replacing consumables, etc.).

One particularly interesting relationship between human and robot scientists relates to the replication of experiments. It has been proposed by Latour [31] that there is a necessary trade-off between the communication of conceptual information and contextual detail: scientists need to report their findings in the most objective and abstract forms possible, so as to create generalizable statements. However, the decision of which detail is conceptual and which is contextual is left to the individual scientist. Anthropological and ethno-methodological research indicates that this is a trait of the human mind: scientists first decide on the experimental design in very broad terms, then later infer how to empirically conduct the experiment, and lastly report their results in mostly the same terms as the experiment was designed [33]. Robot scientists, on the other hand, require a set of definite contextual elements to work properly, and they always log these conditions of the experiments they conduct.

A Sociological Perspective

We argued above that there are convincing examples now of scientists using computer programs to contribute to scientific knowledge. We are particularly interested in the nature of that contribution and how AI-informed scientific knowledge may differ from knowledge gathered without the assistance of computer programs. The earlier description of what a robot scientist does contains two key processes that are critical to understanding scientific knowledge from a sociological perspective: observation and interpretation. We will take each of these concepts in turn. The aim is to offer a sociological perspective on laboratory automation that is informed by some key ideas introduced from sociology of scientific knowledge literature.

When scientists use the term observation they are clearly making a connection between a sensory input (e.g., “seeing”) and the material world around them. However, a great deal of phenomena in the sciences cannot be seen in any literal sense. For example, when a physicist says there are many ways to “observe” the recoil of an atom, they are not referring to a process that can be inspected in the same way that the color red can be seen when a chemistry teacher asks students to observe red fumes in a gas jar [34, pp. 1-2]. It is important therefore to understand exactly what is happening when a robot scientist is said to observe some aspect of the material world.

A system designed to record data must be programmed within a set of parameters, and these parameters need to be established and agreed upon by a group to be validated as correct indicators of the phenomena of interest. For example, an algorithm designed to perform pattern recognition and select objects according to different classes must first be programmed by a group (humans) who have decided which objects should be placed in to each classification, based on examples taken from previous experience of those objects. The objects that seem to generate the most agreement as to the correct class in which they belong can be labelled “exemplars” of that classification. Thomas Kuhn’s *The Structure of Scientific Revolutions* applied the theory of exemplars to scientific knowledge claims, arguing that major advances in scientific fields often amount to a recognition that entities previously thought to be the same actually contain important differences [34].

The language of observation, therefore, is a means by which scientists can make distinctions between what is known and what is not known. There are two important implications of this argument for the use of robot scientists. First, the strength of any knowledge claim is linked to how far observations are agreed upon and shared among authoritative members of a knowledge community. The observations made by the robot scientist were embedded within a well-established and respected research program that employed skilled and knowledgeable scientists. It was those (human) scientists who were recognizably authoritative and credible in their academic discipline. It was mainly because of that credibility that the research team was able to publish the AI-derived data as a contribution to knowledge in their field.

Arguably then, even if robot scientists can be judged to observe phenomena, the validity of the observations takes shape when those observations are discussed and defined through interaction between a social group. Indeed, this process is the very definition of the second concept of interest in this section: “the process... in which individual responses are taken up in to patterns of social interaction ... interpretation” [34, p. 17]. Can a robot scientist be judged to interpret results according to this definition? Here the sociological perspective becomes useful.

Clearly, robot scientists are becoming acceptable as tools for contributing to scientific knowledge among elements of the academic community. The observations made by robot scientists are being taken up into patterns of social interaction via the human designers and operators of those systems, and by the acceptance of AI-derived methods by the wider academic community. Therefore the sociological investigations of robot scientists and their human designers offer a rich seam of possible empirical research: questions that speak to fundamental issues about the role of computer programs for scientific knowledge making in contemporary social life.

Formalization and Reproducibility

One important social and technological aspect of scientific discovery where a robot scientist can offer benefit is reporting and documenting experimental results. This is of particular importance in the context of “the reproducibility crisis”: “The ability to reproduce experiments is at the heart of science, yet failure to do so is a routine part of research” [35]; “More that 70% of researchers have tried and failed to reproduce another scientist’s experiments, and more than half have failed to reproduce their own experiments” [36]. There are many reasons for the non-reproducibility of experimental results: the complexity of experimental and statistical methods, “poor experimental design,” the non-availability of raw data, methods, and code, etc. [35], [36]. The use of natural languages to document experimental procedures makes this worse. For example, consider this routine natural language instruction: “inoculate 4 mL of liquid YPAD and incubate with shaking overnight at 30 C.” This is ambiguous. It does not specify the speed and mode of the shaking – is it 200 rpm, 600 rpm, and orbital or reciprocal? What is the duration of the incubation – is it 8 h, 12 h, or does it not matter? Such information is vital for the reproducibility. The reproducibility of experimental results can be improved through the recording and execution of experiments using robot scientists [37].

Humans are reluctant to record every detail of their cycles of hypotheses – experimentation – interpretations, because it is time consuming and also requires knowledge of reporting standards. Humans are prone to make errors in recording information, and they are biased. Robot scientists are free from such limitations, and can record all the information required for the reproducibility at almost no additional cost and in accordance with the best practices following the recommended standards.

Conclusions and Discussions

Going back to the Glymour paper that opens this article, “Kuhn said that scientific revolutions generally meet fierce resistance – and the automation of discovery in science is no exception” [1]. However, it is to be hoped, that the collaboration between human and robot scientists will produce better science than either can alone – human/computer teams still play better chess than either alone. To understand how best to synergize the strengths and weaknesses of human and robot scientists we need to better understand the anthropological and sociological issues involved in human/machine collaboration. It is also reasonable to hope that developments in robot scientists will contribute to the philosophy of science: compared to traditional approaches the development of robot scientists is constructive and objective.

In chess there is a continuum of ability from novices up to Grandmasters. We argue that this is also true in science, from the simple research of Adam/Eve, through what most human scientists can achieve, up to the ability of a Newton or Einstein. If one accepts this, then just as in chess, it is likely that advances in technology and our understanding of science will drive the development of ever-smarter robot scientists. To encourage research in this area Hiroaki Kitano has called for new grand challenge for AI: to develop an AI system that can make major scientific discoveries in biomedical sciences worthy of a Nobel Prize [6]. This may sound fantastical, yet the Physics Nobel Frank Wilczek is on record as saying that in 100 years’ time the best physicist will be a machine [38]. If this comes to pass, this will not only transform technology, but our understanding of science and the Universe.

Author Information

Ross. D. King is with the Manchester Institute of Biotechnology, University of Manchester, Manchester, U.K.

Vlad Schuler Costa is with the School of Social Sciences, University of Manchester, Manchester, U.K.

Chris Mellingwood is with the School of Social and Political Sciences, University of Edinburgh, Edinburgh, Scotland.

Larisa N. Soldatova is with the Computing Department, Goldsmiths, University of London, London, U.K.

Acknowledgment

The authors of this manuscript have been partially supported by the AdaLab project funded by EPSRC UK (EP/M015661/1).

References

- [1] C. Glymour. "The Automation of Discovery," *Daedalus*, Vol. Winter (2004), 69-77.
- [2] B.G. Buchanan, G.L. Sutherland, & E.A. Feigenbaum. "Heuristic DENDRAL: A program for generating explanatory hypotheses in organic chemistry." In B. Meltzer & D. Michie (eds.), *Machine Intelligence 4*. Edinburgh: Edinburgh University Press, 1969.
- [3] P. Langley et al. *Scientific Discovery: Computational Explorations of the Creative Process*. Cambridge, MA: MIT Press, 1987.
- [4] J.M. Żytkow, J. Zhu, & A. Hussam. "Automated discovery in a chemistry laboratory," in *Proceedings of the 8th National Conference on Artificial Intelligence (AAAI-1990)*, 1990, 889-894.
- [5] R.E. Valdes-Perez. "Principles of human-computer collaboration for knowledge discovery in science," *Artificial Intelligence*, vol. 107, pp. 335-346, Feb. 1999.
- [6] H. Kitano. "Artificial Intelligence to Win the Nobel Prize and Beyond: Creating the Engine for Scientific Discovery," *AI Magazine*, vol. 16, pp. 39-49, Spring 2016.
- [7] P. Langley. "The computational support of scientific discovery," *Intl. J. Human-Comp. Stud.*, vol. 53, pp. 393-410, Sep 2000.
- [8] R.D. King et al. "The Automation of Science," *Science*, vol. 324, pp. 85-89, Apr. 2009.
- [9] K. Williams et al. "Cheaper Faster Drug Development Validated by the Repositioning of Drugs Against Neglected Tropical Diseases," *J. R. Soc. Interface*, vol. 12, pp. 1-9, Feb. 2015.
- [10] K. Popper. *The Logic of Scientific Discovery*. London: Hutchinson, 1972.
- [11] R.D. King et al. "Functional genomic hypothesis generation and experimentation by a robot scientist," *Nature*, vol. 427, pp. 247-252, January 2004.
- [12] A. Sloman. *The Computer Revolution in Philosophy*. Sussex: Harvester Press, 1978.
- [13] M. David. "The Correspondence Theory of Truth." In E.N. Zalta (ed.), *The Stanford Encyclopaedia of Philosophy*, Fall 2016.
- [14] J. Owens. "Realism." In *Encyclopaedia Britannica Deluxe Edition 2004 CD*. London: Encyclopaedia Britannica UK.
- [15] R. Carnap. "Intellectual Autobiography." In: P. Schilpp (ed.), *Philosophy of Rudolph Carnap*. La Salle: Open Court, 1963.
- [16] M. Curd & J.A. Cover. *Philosophy of Science: The Central Issues*. New York: W.W. Norton & Company, 1998.
- [17] A. Miller. "Realism." In E.N. Zalta (ed.), *The Stanford Encyclopaedia of Philosophy*, Winter 2016.
- [18] D. Gillies. *Artificial Intelligence and Scientific Method*. Oxford: Oxford University Press, 1996.
- [19] R. Carnap. *Introduction to symbolic logic and its applications*. Mineola: Dover Publications, 1958.
- [20] D.A. Cohn, Z. Ghahramani, M.I. Jordan. "Active learning with statistical models," *J. Artif. Intell. Res.*, vol. 4, pp. 129-145, Mar. 1996.
- [21] I. Kasavin, T. Rockmore, & E. Blinov. "Social Epistemology, Interdisciplinarity and Context," *Epistemology & Philosophy of Science*, vol. 37, pp. 57-75, Jan. 2013.
- [22] N. Cartwright, N. "Relativism in the Philosophy of Science. In *Relativism: A Contemporary Anthology*." Krausz, M. New York: Columbia University Press, 2010, 86-99.
- [23] T.S. Kuhn. *The Structure of Scientific Revolutions*. Chicago: University of Chicago Press, 1962.
- [24] J. Searle, J. *Minds, Brains and Science*. Cambridge: Harvard University Press, 1986.
- [25] M. Davies. *Engines of Logic: Mathematicians and the origin of the Computer*. New York: W. W. Norton & Company, 2000.
- [26] A. Turing. "Computing Machinery and Intelligence," *Mind*, vol. 59, pp. 433-460, Oct. 1950.
- [27] J. Preston, "Review: 'Artificial intelligence and scientific method' and 'AI: essays at the interface'," *The British Journal for the Philosophy of Science*, vol. 48, pp. 610-612, Dec. 1997.
- [28] D.A. Mindell. *Our robots, ourselves: Robotics and the myths of autonomy*, 2015.
- [29] O'Hanlon, J.F. "Boredom: Practical Consequences and a Theory," *Acta Psychologica*, vol. 49, pp. 53-82, Oct. 1981.
- [30] D. B. Kronenfeld et al (eds). *A Companion to Cognitive Anthropology*. London: Wiley-Blackwell, 2011.
- [31] B. Latour. "Circulating Reference." In B. Latour, *Pandora's Hope: essays on the reality of science studies*. Cambridge: Harvard University Press, 1999.
- [32] M. Lynch. *Art and Artifact in Laboratory Science: a study of shop work and shop talk in a research laboratory*. London: Routledge, 1985.
- [33] S. Schaffer & S. Shapin. *Leviathan and the Air-Pump: Hobbes, Boyle, and the Experimental Life*. Princeton: Princeton University Press, 1985.
- [34] B. Barnes, D. Bloor, & J. Henry. (1996). *Scientific Knowledge: A Sociological Analysis*. London, Athlone.
- [35] Nature Editorial. "Reality check on reproducibility," *Nature*, vol. 533, pp. 437-438, 2016.
- [36] M. Baker & D. Penny. "Is there a reproducibility crisis?," *Nature*, vol. 533, pp. 452-454, 2016.
- [37] L.N. Soldatova et al. "An Ontology for a Robot Scientist," *Bioinformatics*, vol. 22, pp. e464-471, 2006.

[38] F. Wilczek. *Fantastic Realities: 49 Mind Journeys and a Trip to Stockholm*. Singapore: World Scientific Publishing, 2006.

Figures



Figure 1. The Robot Scientist 'Eve'.

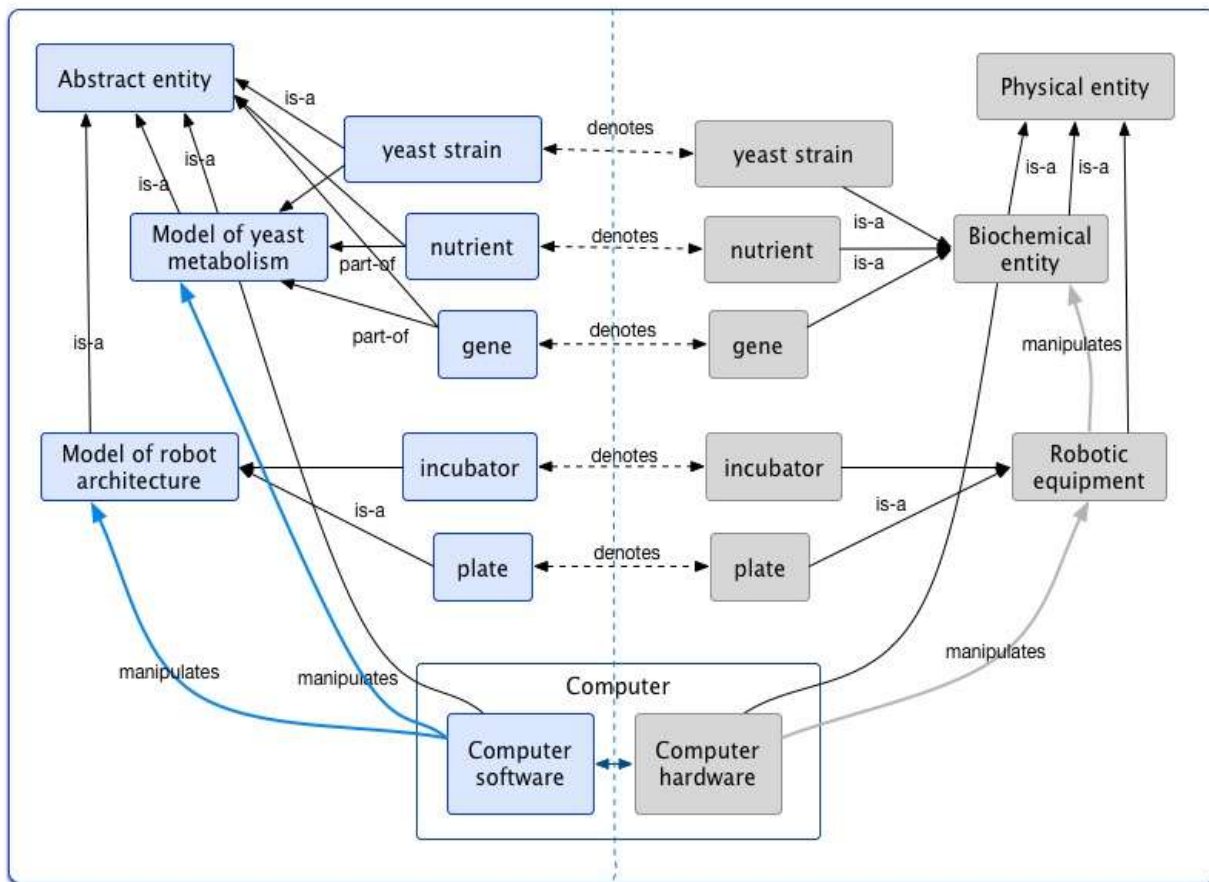


Figure 2. The Robot Scientist Universe (a fragment).