

Large-scale simulations of brain mechanisms: beyond the synthetic method

Edoardo Datteri, Federico Laudisa

Dipartimento di Scienze Umane per la Formazione “R. Massa”, Università degli Studi di Milano-Bicocca

{edoardo.datteri | federico.laudisa}@unimib.it

Abstract. In recent years, a number of research projects have been proposed whose goal is to build large-scale simulations of brain mechanisms at unprecedented levels of biological accuracy. Here it is argued that the roles these simulations are expected to play in neuroscientific research go beyond the “synthetic method” extensively adopted in Artificial Intelligence and biorobotics. In addition we show that, over and above the common goal of simulating brain mechanisms, these projects pursue various modelling ambitions that can be sharply distinguished from one another, and that correspond to conceptually different interpretations of the notion of “biological accuracy”. They include the ambition (i) to reach extremely deep levels in the mechanistic decomposition hierarchy, (ii) to simulate networks composed of extremely large numbers of neural units, (iii) to build systems able to generate rich behavioural repertoires, (iv) to simulate “complex” neuron models, (v) to implement the “best” theories available on brain structure and function. Some questions will be raised concerning the significance of each of these modelling ambitions with respect to the various roles played by simulations in the study of the brain.

Key-words: large-scale brain simulations; simulation methodologies in neuroscience; synthetic method; biological accuracy; models in neuroscience; computational neuroscience.

1. Introduction

As extensively discussed in Roberto Cordeschi’s *The Discovery of the Artificial* (Cordeschi, 2002), the implementation of machines which can be sensibly said to accurately reproduce biological mechanisms has been occasionally pursued in biorobotics and Artificial Intelligence. This ambition is being pushed to unprecedented levels of biological accuracy in a number of contemporary research projects, which aim at building large-scale simulations of brain mechanisms. Notable examples are the “Blue Brain project” (Markram, 2006), the “Cognitive Computing via Synaptronics and Supercomputing” project (Ananthanarayanan *et al.*, 2009), and the “Cognitive Computation Project” (Eliasmith *et al.*, 2012); for a review, see (De Garis *et al.*, 2010). According to Blue Brain Project leader Henry Markram, a “quantum leap” towards the development of extremely accurate artificial models of brain mechanisms is now made possible by the availability of extremely powerful supercomputers, such as the IBM Blue Gene, able to carry out billions of floating-point operations per second.

The broad ambition to «simulate the brains of mammals with a high level of biological accuracy» (Markram, 2006, p. 153) is not always accompanied by a clear statement of the research questions that these projects are expected to address, however. In some cases, purely technological motivations (connected to the development of new-generation biologically inspired supercomputers) seem to prevail over neuroscientific interests. Nor the leading researchers always justify their insistence on biological accuracy and on the importance of simulating extremely large-scale networks with respect to the study of the brain. Moreover, various forms of “biological accuracy” appear to be pursued in these projects. Some of them strive to reach the *number of neural units* existing in the brain of mammals, but use extremely *abstracted neuron models*; other projects aspire to build very *fine-grained models* of neurons, while yet other projects focus on the *number of modelled behaviours*. The purpose of this paper

is to contribute to the rising philosophical debate on these research and development enterprises (see, e.g., Miłkowski, 2015) by making these distinctions and by introducing a number of epistemological and methodological questions on the import of these different modelling ambitions with respect to the roles that simulations may play in neuroscientific research. Notably, as discussed in the next section, these roles are not limited to the discovery of models of brain mechanisms as in the “synthetic method” which has been analysed in a philosophical and historical perspective by Roberto Cordeschi.

2. *The synthetic method, today*

As discussed by Cordeschi (2002, pp. 1-10 and pp. 82-115), one of the first examples of the so-called “synthetic method” for the study of intelligent behaviour and cognition can be found in the reaction of physiologist Jacques Loeb to the construction, in 1912, of a machine able to steer towards sources of light. In his *Comparative Physiology of the Brain and Comparative Psychology*, Loeb (1900) proposed an explanation of moths’ ability to track light sources (a form of behaviour often called *phototropism*) based on a very simple mechanism: light stimulates muscle activity, so that the motor organs located at the side hit by light move faster than the organs located at the opposite side. The robot built by John Hammond Jr. and Benjamin Miessner had two front light sensors and a simple electro-mechanical mechanism steering the robot towards the side receiving higher light stimulation. Miessner stressed that the structure of the robot was «very similar to that given by Jacques Loeb, the biologist, of reasons responsible for the flight of moths into a flame» (cited in Cordeschi, 2002, p. 3-5). Indeed, some years after, Loeb realized that the ability of the robot to track light sources could be taken to support his theory on moth phototropism: a machine reproducing the essential aspects of the theory proved able to generate phototropic behaviours.

This example illustrates the structure of the so-called “synthetic method”, which characterizes important areas of research in Artificial Intelligence and contemporary neuroethology. More schematically (see Figure 1), let BL be a description of a particular aspect of the behaviour of living system L, and let ML be a description of the mechanism which is hypothesized to produce BL in L. To test whether this hypothesis is true, one may build a machine A governed by mechanism ML (more precisely, by a mechanism MA which is similar in the relevant aspects to ML¹) and compare its behaviour with the behaviour of L. If A’s behaviour BA is similar in relevant respects to BL, one may be induced to corroborate the hypothesis that ML produces BL, thus to *explain* BL with reference to ML. Otherwise, one may be induced to reject that hypothesis. Under a variety of epistemological and methodological assumptions, whose analysis is out of the scope of this contribution (see Cordeschi, 2002, 2008; Webb, 2006; Datteri and Tamburrini, 2007; Datteri, 2016), the synthetic method may therefore assist one in identifying the mechanism underlying a particular (observed) behaviour. This may be called a *model-oriented* use of simulations.

¹ As discussed later, to justifiably bring A’s behaviour to bear on the plausibility of ML, one has to assume that A is an accurate implementation of ML. This raises the problem of understanding under what conditions A can be sensibly said to be an accurate implementation of ML. To address this problem, whose analysis is out of the scope of this paper, it is useful to distinguish the target biological hypothesis ML from the blueprint MA describing the technical specifications that artificial system A is expected to satisfy. Indeed, A may fail to be an accurate implementation of ML for two kinds of reasons: 1) A may fail to implement all the technical specifications expressed in blueprint MA (for example, because the required electro-mechanical components were not available, and components differing from them in some non-negligible aspects have been used instead); or, 2) MA may consist in a distorted version of ML. A more detailed discussion of these methodological issues, together with an account of what is for a blueprint of a computer or robotic system to be an accurate translation of a biological mechanism description, is provided in (Datteri, 2016).

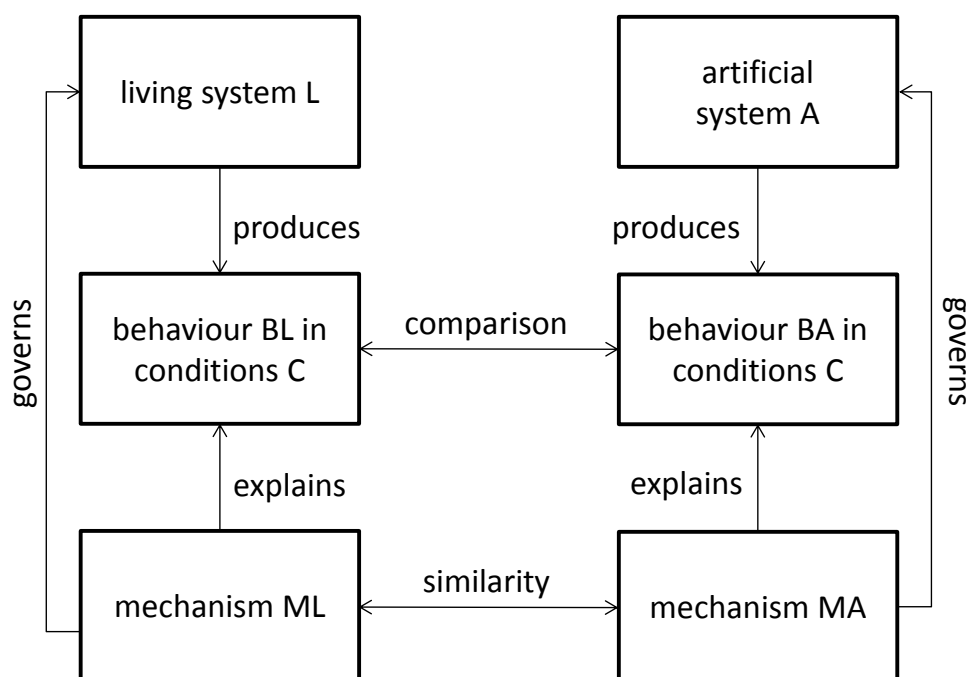


Figure 1 – Schema of the “synthetic method”

Nowadays, however, simulations are also used for a rather different purpose, namely, to obtain behavioural data which are difficult or impossible to obtain through alternative strategies. For example, molecular-level simulations are used as “computational microscopes” (Dror *et al.*, 2012) to predict the behaviour of ion channels under a variety of physiological conditions. Simulations are essential to obtain fine-grained descriptions of ion channel behaviours: there are simply no alternative ways to observe them at the same spatial scale and at the same level of detail. In another branch of simulation-supported biology, called “evolutionary biorobotics”, a research group led by John Long builds robots reproducing the sensory-motor mechanisms and the physical structure of extinct animals to obtain their behaviour under a variety of conditions (Long, 2012). The purpose of these studies is not to discover the mechanism underlying a particular behaviour, but to obtain data on the behaviour of a system which is hard or impossible to observe through more conventional techniques. This may be called a behaviour-oriented or a *data-oriented* use of simulations.

Model-oriented and data-oriented simulation studies crucially differ from one another in the nature of their research goals. This difference also reflects some *methodological* differences between the two, a brief discussion of which will be useful to assess the epistemic value of large scale simulations of the brain. First, note that the model-oriented (synthetic) methodology crucially involves a comparison between the behaviour of the artificial system A and the behaviour of the target system L. Such a comparison is not part of the data-oriented methodology, exactly because there are no living system data to compare artificial behaviours with (recall that data-oriented simulations are used to *obtain* behavioural data which are difficult or impossible to obtain otherwise).

A second difference concerns the degree of corroboration of the mechanism ML simulated – via translation in model MA – in the artificial system. A fundamental methodological requirement of model-oriented simulation studies is that the artificial system A must simulate accurately the model ML under scrutiny – otherwise, there would be no reason to bring A’s behaviour to bear on the plausibility of ML itself (see footnote 1). It is not required, however, that ML be a “good” (that is to say, explanatory and highly corroborated) model of the target system L. Even though, as we will discuss later on, some evidence may already be available in support of ML, the goal of model-oriented simulations is exactly to corroborate it – and for this reason one cannot assume that ML is a good model of the target system *before* carrying out the simulation experiments. Instead, to make a proper *data-oriented* use of a simulation, one has to

assume that ML is a good model of L, otherwise there are no reasons to consider the behaviour of the simulation as the behaviour that the target system would have produced in those conditions. To consider the output of a simulation as informing one on the behaviour of a particular class of ion channels under particular physiological conditions, one has to ensure that the simulation is based on the best theories available on the physical structure and on the molecular interactions governing channels belonging to that class. Otherwise, the simulation would be informative of the behaviour of ion channels characterized by a different structure and different molecular interactions – or of the behaviour of a radically different, possibly non existing, system.² In addition, one has clearly to assume that A is a good simulation of ML, as in the model-oriented case.³ These distinction will be useful in the ensuing discussion of the methodologies adopted in contemporary large-scale brain simulation projects. Indeed, the leaders of these projects are often ambiguous as to whether their goals are on the data-oriented or on the model-oriented side – in the latter case, their insistence on the plausibility of the neuroscientific theories used within their projects is not totally justified.

Note also that, in light of these methodological differences, model-oriented and data-oriented simulation studies give rise to relatively different epistemological issues. In both cases, one may legitimately ask under what conditions one is really entitled to draw theoretical conclusions on the target living system based on the analysis of artificial system behaviours. With reference to the data-oriented method, this amounts to asking under what conditions one is really authorized to consider the behaviour of the simulation system as the behaviour that the target system would have produced in the same experimental circumstances. The question whether simulations are experiments is widely discussed in the epistemological literature (Humphreys, 2004; Parker, 2009; Winsberg, 2003). This question concerns data-oriented simulations – thus, notably, it does *not* concern the use of simulations typically made in AI and biorobotics – and can be rephrased as a question whether the data produced by simulation studies of this kind have the same epistemic value as data acquired through observations and measurements made on the target system. The epistemic conditions under which such a judgment can be authorized notably include, as pointed out before, conditions on the plausibility of the theoretical model of the target system: the methodology of data-oriented simulations thus gives rise to the problem of evaluating the explanatory value of the theoretical models they implement. The availability of criteria to evaluate the plausibility of the underlying theoretical model is not required in the model-oriented methodology, instead, exactly because – as pointed out before – that model need not be plausible for the method to be sound. Rather, this methodology raises only the problem of understanding under what conditions A can be considered an accurate implementation of ML (a problem which, as pointed out before, is raised by data-oriented simulations too).

To sum up, data-oriented and model-oriented simulation studies have different kinds of scientific goals, have different methodological requirements, and give rise to different epistemological issues. It is worth stressing that the “synthetic method” widely discussed in Cordeschi (2002) coincides with the model-oriented methodology. The role of machines in data-oriented simulations, only minimally covered by Cordeschi’s analysis, is therefore rather different from the role assigned to machines in theoretically oriented (rather than in engineering oriented) AI and contemporary biorobotics.

As mentioned before, it is not always clear whether contemporary large-scale brain simulation projects (and calls for projects) are data-oriented, model-oriented, or both. The ninth of the “14 Grand Challenges for Engineering in the 21st Century” proposed by the US National Academy of Engineering is entitled “Reverse-Engineer the Brain”.⁴ The goals of this challenge

² This does not exclude that one can assess the behaviour of the simulation under implausible physiological or environmental conditions, in order to predict how the target system would behave in those cases.

³ Mixed uses of simulations are clearly possible, as the same simulation system can be used in a data-oriented and in a model-oriented way at different times. Consider, for example, weather forecasts. Simulations are used to predict meteorological data based on the best weather models available at the moment; when actual weather data become available, a comparison between them and the results of the simulation is used to refine the underlying theoretical model or to correct some theoretical assumptions accompanying it, as in model-oriented simulation studies.

⁴ <http://www.engineeringchallenges.org/9109.aspx>, visited on 19/09/2015.

are recapitulated in Roysam, Shain, Ascoli (2009): «reverse engineering the brain goes well beyond the idea of mapping its structure, its cellular makeup, and molecular composition, although these are necessary prerequisites. To meet the ninth Grand Challenge, one must take the nontrivial next step, and create a successful computational system (combining appropriate hardware and software components) that algorithmically recapitulates all important brain functions» (p. 2). The goal of this Challenge, as stated here, consists in the simulation of a detailed, large, and *plausible* model of the brain (it must recapitulate “all important brain functions”, which therefore must have been already discovered before building the simulation), and for this reason the Challenge praises satisfaction of the central requirement of data-oriented simulation studies.

Why should the brain be reverse-engineered? According to the text of the Grand Challenge, simulating brain activities may lead «to deeper insights about how and why the brain works and fails», thus to the discovery of models of brain functioning: this is a model-oriented use of simulations (this claim contrasts with the previous statement that the model must recapitulate “all important brain functions”). Immediately after, the authors add: «Such simulations will offer more precise methods for testing potential biotechnology solutions to brain disorders, such as drugs or neural implants». Here it is suggested that a computer model of the brain could assist in discovering how brain behaviour would change in particular conditions, that is to say, under the effect of certain drugs or after connection with additional devices. This goal is closer to the data-oriented side, as it is the goal of obtaining data on the target system (a plausible model of it is required).

Similar ambiguities can be found in the scientific literature on specific simulation projects. According to Eliasmith and Trujillo (2014), for example, one of the reasons to build large-scale simulations is «to understand mysterious brain disorders, from autism to addiction». This objective is closer to the model-oriented side, as it consists in the discovery of theoretical models of brain and cognitive phenomena. Similarly, Kandel *et al.* (2013) point out that «the long-term goal of these highly ambitious projects is to gain a better understanding of the anatomical, molecular and circuit bases for the logical operations carried out by the human brain» (p. 659). In their opinion the Blue Brain Project «aims to understand the human brain by simulating its functions through the use of supercomputers» (p. 659). Another objective of large-scale simulations is, according to Eliasmith and Trujillo (2014), «to develop and test new kinds of medical interventions, be they drugs or stimulation» (p. 3). Here the authors have plausibly in mind the realization of an accurate simulation of the brain and the observation of its behaviour under the effect of particular interventions or medications: this would be a data-oriented use of simulations.

It is worth noting, in addition, that simulations are sometimes expected to play a key role in *integrating* knowledge coming from different studies and possibly from different research disciplines. Markram *et al.* (2011) claim that one of the goals of the Human Brain Project (HBP) is «to integrate everything we know in multilevel brain models» (p. 40). He claims that «the HBP sets academia and industry on a new road to understanding the human brain. On the way, it will *unify existing biological knowledge*, generate new approaches and methods for the brain sciences, and develop new intelligent technologies» (p. 41, *emphasis added*). Eliasmith and Trujillo (2014) also point out that a major purpose of building large-scale brain simulations is «to provide a way to organize and unify the massive amounts of data generated by the neurosciences» (p. 3).

Whether and how simulations can really assist in integrating knowledge on the brain is a question that essentially depends on what we mean with “integrating knowledge”. Let us introduce some possible interpretations, which will be further illustrated and articulated later in connection with large-scale simulations of the brain.

One way simulations can contribute to the integration of knowledge on brain mechanisms is inherent in the model-oriented (synthetic) method. We have claimed that the plausibility of the

mechanism description ML simulated in the artificial system A is not a requirement of “good” model-oriented studies. This is not to say that ML must be implausible, nor that no evidence at all on ML must be available before carrying out the simulation experiments. It may be the case that ML has been previously corroborated through other experimental strategies, and that a model-oriented study is expected to provide further evidence in support of it. Or, it may be the case that ML has been only partially (in a sense to be clarified) corroborated before running the simulation. With reference to the purely notional mechanism described in Figure 2, suppose that strong evidence – coming from different studies made at different times – is already available on the existence of components b_1 , b_2 , and b_3 in L and on their behaviour, but that no evidence is available as to whether their organized interaction can actually produce the behaviour of interest. This is a first sense of ML being only partially corroborated: even though much is known on the mechanistic structure of the target system L, yet one is not in the position to corroborate the hypothesis according to which ML actually produces (that is to say, enables one to explain) the behaviour of interest. The ability of artificial system A to produce the behaviour under investigation may be taken as a basis to accept that hypothesis. Note that in this case the simulation would have contributed to *integrating* what different studies on L have previously discovered about the same system, in the same way as different pieces of a puzzle are assembled together to produce a particular figure.




Figure 2 – A purely notional mechanism description

Or, suppose that strong evidence is available on the existence and behaviour of b_1 and b_2 only, but that the existence and behaviour of b_3 are highly speculative. A’s success in reproducing BL may be taken as evidence to corroborate the theoretical claims made in ML as far as b_3 is concerned. This notional example illustrates another way in which model-based simulation studies can integrate already available knowledge on the brain, the term “integration” here implying the filling of gaps in a mechanism description that, as a result, becomes fully effective in explaining the target behaviour. An example is the biorobotics study on rat navigation described in Burgess *et al.* (2000), in which robotic behaviours have been taken as a basis to believe in the existence of so-called “goal cells”, never discovered in the rat at the time of publication of that work, but whose functional role must be instantiated in the robotic system for the latter to generate the behaviour under investigation (see Datteri and Tamburrini, 2007 for a discussion).

The claim that simulations may enable one «to integrate everything we know in *multilevel* brain models» (Markram *et al.*, 2011, p. 40, emphasis added), however, is likely to refer to forms of integration which are different from those discussed so far: his point is that building a simulation can contribute to identifying relationships or bridges between various levels of analysis at which a given phenomenon can be explained. Whether this is the case or not is a question that depends on what Markram means by “level” of analysis. This point will be

discussed in the next section, in which various concepts of level will be discussed in connection with the notion of “biological accuracy”.⁵

3. *Dimensions of theoretical modelling in contemporary brain simulation projects*

Contemporary large-scale simulation projects aim at simulating the brain at high degrees of biological accuracy. On a closer look, their efforts towards biological accuracy take different forms, corresponding to different views on what makes a biologically accurate model or simulation. Some of these views are discussed below. The distinctions between model-oriented and data-oriented simulations made in the previous section will be useful to assess the role that the different forms of biological accuracy identified here can play in the study of the brain.

3.1. *Levels of functional decomposition*

As mentioned earlier, the Blue Brain Project had the ambitious goal of building «accurate models of the mammalian brain from first principles» (Markram, 2006, p. 155). In his (2006) *Nature Reviews* paper he offers several insights to understand what makes, in his opinion, an accurate theoretical model of the brain, and therefore to understand the nature of the goals of the Blue Brain project. A diagram included in that paper is particularly interesting in this respect, as there he represents what he thinks are «the minimal essential building blocks required to reconstruct a neural microcircuit» (see Figure 2 in Markram, 2006, p. 155). Let us focus on a specific subset of these requirements. He points out that «microcircuits are composed of neurons and synaptic connections» (p. 155), and that neurons must be characterized by their gene expression, electrophysiological, and morphological profiles. Most notably, he adds that «to model neurons, the three-dimensional morphology, *ion channel composition*, and distribution of electrical properties of the different types of neuron are required» (p. 155, emphasis added).

The relationship between a neural microcircuit, the neurons composing it, and the ion channels spanning the membrane of each neuron is a relationship of *mechanistic decomposition*, at least as far as the electrophysiological profile of neurons is concerned. Indeed, the electrical activity of the microcircuit as a whole – for example, the relationship between the “input” and “output” neurons of it – crucially depends on the electrical activity of each neuron of the network and on their mechanistic organization. And the electrical activity of each neuron depends on the number, distribution, and type of the ion channels spanning its membrane. In other terms, neurons are components of the mechanism responsible for the electrophysiological behaviour of the whole circuit, and ion channels are components of the mechanism responsible for the electrophysiological behaviour of individual neurons.

To generalize, the process of iterating mechanistic analysis over the components of a particular mechanism generates a hierarchy of levels of analysis, as schematized in Figure 3. The mechanism description at level $n-1$ mentions the mechanisms governing the behaviour of the components which are the base units of the mechanism at level n . The relationship between levels is based on *mechanistic decomposition*. This recursive process corresponds to the progressive opening of closed boxes in Rosenblueth’s and Wiener’s account of theoretical modelling (Rosenblueth, Wiener, 1945, p. 319).

⁵ In the “horizontal”, non-multilevel cases discussed here, integration is achieved by corroborating a mechanistic model of the target system. Integration is therefore essential to mechanistic explanation. Whether multi-level integration is a requirement of a “good” explanation is a more controversial question, which clearly depends on what is meant with “level”. A detail analysis of this broad problem goes out of the scope of the present article. A closely related question not addressed in this paper is whether horizontal or multi-level integration in simulative studies can contribute to scientific unification in the sense discussed by Friedman (1974) and Kitcher (1981). For a critical analysis of whether scientific unification, as interpreted by the latter authors, is necessary and/or sufficient for a “good” explanation see (Gijssbers, 2007).




Figure 3 – Schema illustrating mechanistic decomposition

Markram's claim discussed above suggests that, in his opinion, mechanistic decomposition down to the level of ion channel behaviour is required in a “good” model of a neural microcircuit. Moreover, he and other researchers advocate the building of supercomputer – possibly more powerful than the IBM Blue Gene, used in the Blue Brain Project – enabling one to iterate mechanistic decomposition at even lower levels. In an introductory paper to the Human Brain Project (Markram *et al.*, 2011) it is suggested that «petascale computers, now available, are potentially powerful enough for cellular-level simulations of the whole rodent brain, or for molecular level simulations of single neurons. Exascale computers, predicted for the end of the decade, could allow cellular level simulations of the complete human brain with dynamic switching to molecular-level simulation of parts of the brain when required» (p. 40).

The realization of simulations spanning many levels of the mechanistic decomposition hierarchy, from neural behaviour down to molecular interactions and possibly beyond, seems to be one of the ambitions pursued in large-scale brain simulation projects. Technological factors, he claims, constitute the main obstacle to the achievement of such an objective.

The main limitations for digital computers in the simulation of biological processes are the extreme temporal and spatial resolution demanded by some biological processes, and the limitations of the algorithms that are used to model biological processes. If each atomic collision is simulated, the most powerful supercomputers still take days to simulate a microsecond of protein folding, so it is, of course, not possible to simulate complex biological systems at the atomic scale. (Markram, 2006, p. 158).

However, this claim leaves open the problem of understanding why this very computationally demanding ambition should be pursued. Are there really good reasons to praise mechanistic decomposition down to the level of ion channels, or possibly even beyond, in the construction of large-scale brain simulation projects? This question takes different forms, and possibly admits of different answers, depending on whether the simulation has a model-oriented or a data-oriented goal, or on whether it is expected to provide the basis for multilevel integration of knowledge on the brain.

With reference to a model-oriented use of simulations, it can be rephrased as a question on the characteristics of a “good” theoretical model of the brain. Recall that the objective of a model-oriented simulation study is to test whether the mechanism description ML implemented in the machine is a good basis to explain why system L produces behaviour BL: therefore, by the question above one is asking whether the explanatory value of ML increases with the number of mechanistic decomposition levels covered by it. According to some influential theories on mechanistic explanation (Craver, 2007; Woodward, 2002), explanatory mechanism descriptions mention all and only those factors that make the difference with respect to the phenomena under investigation. And, by going downward the decomposition hierarchy and recursively “looking inside” previously closed boxes, one progressively identifies more and more difference-making factors. However, as pointed out by Eliasmith and Trujillo (2014), the choice of the level at which mechanistic analysis should bottom out is likely to depend on the characteristics of the phenomenon to be explained. Pushing mechanistic analysis down to the level of ion channels is likely to be required, for example, if one intends to explain why certain

forms of behaviour are produced under particular kinds of physiological conditions affecting ion channel activity.

As far as the data-oriented use of simulations is concerned, the question above points to the benefits of going deeper and deeper in the mechanistic decomposition hierarchy with respect to the goal of *predicting* the behaviour of the target system. In principle, if the mechanism at level n correctly describes the behaviour of the various components and their interconnections, one of the two key requirements of data-simulation studies is satisfied: the machine implements an (at least predictively) adequate theoretical model of the target system. By moving down to level $n-1$ one identifies the mechanisms governing the behaviour of level- n components. But it is not clear why this move should increase the predictive value of the simulation, given that (by assumption) the behaviour of level- n components has been already accurately identified. In other terms, this move would not improve one's knowledge of the behaviour and organization of level- n components – it would only explain why their behaviour is as it is. Penetrating into deeper levels of the mechanistic decomposition hierarchy might be useful, however, to understand how the behaviour of the target system would change under interventions on components of those levels. For example, simulating brain processes all the way down to ion channel behaviour might enable one to predict the effects of blocking or perturbing the activity of particular kinds of ion channels on the overall brain behaviour – a possibility which would be impossible to obtain in more “shallow” simulations.

Finally, it is worth noting that the development of a mechanistic decomposition hierarchy that spans many levels of analysis may represent a step towards what Markram calls “multilevel integration” of knowledge of the brain. The ability to simulate mechanisms at various levels of analysis in the hierarchy may contribute to achieving this goal. Suppose, for example, that one has a plausible theory of the behaviour of components at level $n-2$ and a plausible description of the behaviour of the system at level n . Ideally, simulations of these theories may enable one to discover a plausible mechanistic description of the system at level $n-1$, thus to build a bridge between the two. Indeed, one could implement data-oriented simulations of $n-2$ components to predict their behaviour; try and organize them in various ways so as to produce simulations of various putative $n-1$ components; then, iterate the same process on the newly formed $n-1$ components until a mechanism is found which produces exactly the behaviour defined at level n . In this notional example, data-oriented simulations of mechanism components situated at “low” levels of analysis are used as building bricks to progressively walk the decomposition hierarchy upward.

3.2. *Size of the theoretical model*

Another broad goal of contemporary large-scale simulation projects, to be kept distinct from the goal discussed in the previous section, is to extend the “horizontal” rather than the “vertical” size of the model, namely to simulate mechanisms composed of a huge quantity of neurons or base units. One goal of the Blue Brain Project was to build a simulation of a portion of the somatosensory cortex of the rat composed of about 10.000 neurons, while the Blue Gene – the supercomputer used in the experiments – was reported to be able to simulate a 100.000-neuron neural network. Eliasmith's SPAUN model comprises 2.5 million neurons, and Markram welcomed the development of computational techniques able to simulate the entire human brain with its 100 billion neurons. These techniques «provide a strong foundation for taking the next quantum step, to further increase the size of the modelled network to an unprecedented level» (Markram, 2006, p. 154). It is reasonable to believe that, in these authors' view, size is an important dimension of biological realism: their ambition towards the construction of biologically realistic simulations is nearly always accompanied by emphasis on the huge number of base units they can implement.

This dimension is conceptually distinct from the “vertical” axis of mechanistic decomposition discussed before and from the other dimensions of theoretical modelling that will be analysed in the ensuing sections. Surely, the number of base units of a theoretical model is likely to increase as soon as one goes downward along the mechanistic decomposition hierarchy (each neural area has many neurons; each neuron has many ion channels). However, one can increase the number of base units without changing the level of mechanistic decomposition simply by adding other same-level units and mechanisms. One can choose, for example, to increase the number of neurons in an artificial neural network to improve its input-output accuracy or its learning profiles, or to add other same-level networks and mechanisms to increase the behavioural repertoire of the system. Indeed, the path imagined by these authors towards the creation of an artificial brain is a sort of upside-down version of the mechanistic decomposition hierarchy: «A natural progression is ... to simulate neurons embedded in microcircuits, microcircuits in the local circuits of brain regions, and circuits within regions and the whole brain» (Markram, 2006, p. 154).

The construction of supercomputers able to simulate networks composed of billions of neurons surely represents a major technological advancement. It is not obvious, however, that increasing the size of a simulation is of any theoretical interest with respect to the explanation or the prediction of brain behaviour. Let us consider first the model-oriented role sometimes assigned to large-scale simulations of the brain. The question can be reformulated as to whether increasing the size of the underlying mechanism description ML increases its explanatory power, and it is reasonable to believe that this needs not be the case – it depends on the particular *explanandum* addressed in the model-oriented study. Some forms of behaviour may well be produced by complex interactions between many brain areas while other behavioural *explananda* may be suitably addressed by considering lower-scale mechanisms. Cognitive science and neuroscientific behavioural *explananda* are typically defined in a way that legitimates abstraction from the interaction of concurrent mechanisms and boundary conditions. For example, every form of sensory-motor coordination in the everyday life is likely to require massive and widespread brain activation. However, scientific explanations of sensory-motor coordination capacities start with the definition of *explananda* which abstract away from many aspects of their everyday-life form – consider, for example, the goal of understanding why particular kinds of eye movements are produced under well-defined and typically narrow classes of visual stimuli in aseptic and artificial laboratory environments. The typical way scientific *explananda* are carved out of everyday behaviour is exactly meant to authorize abstraction from concurrent mechanisms and possible perturbing factors in the explanation (Bogen and Woodward, 1988; Suppe, 1989; Datteri and Laudisa, 2014).

To sum up, contemporary supercomputers may well contribute to the understanding of brain mechanisms by simulating theoretical models which make reference to a multitude of brain areas and to a huge number of base units (Eliasmith and Trujillo, 2014), and large-scale simulation projects may end up with the development of technologies and computational frameworks that, *in principle*, enable one to test hypotheses of high dimensionality. Nevertheless, “good” mechanistic explanations of behaviour need not be large. They must capture all and only the causally relevant factors for the production of the behaviour to be explained (Woodward, 2002), and whether the number of those factors is high or low is a question that depends, among other factors, from the characteristics of the *explanandum*.

The magnitude of the model may be relevant in a data-oriented perspective, however. As mentioned earlier, increasing the number of units in an artificial neural network may have positive effects on its ability to learn particular input-output functions. Moreover, adding same-level mechanisms may increase the behavioural repertoire of the system and consequently, as discussed in the next section, the usefulness of the simulation to predict the various behavioural effects of particular interventions or stimulations.

3.3. *The number of modelled phenomena*

Markram's Blue Brain and Human Brain projects, and Eliasmith's Cognitive Computation project – the latter aiming at developing a 2.5-million-neuron model of the brain called "Spaun" – are some of the most popular large-scale simulation projects proposed in recent years. Despite the complexity of their internal architecture, some general differences in their modelling approaches have been identified. Eliasmith and Trujillo (2014) and Miłkowski (2015) have pointed out that Markram's projects are not guided by clear-cut *explananda*. Even though, as mentioned above, understanding intelligent behaviour is among their goals, a sufficiently well-defined description of the behaviours to be explained is missing. The main ambition of the Blue Brain and of the Human Brain Project is, rather generically, to «simulate the brain of mammals» (Markram, 2006, p. 153) and "to reconstruct and simulate the human brain and its diseases» (Kandel *et al.*, 2013, p. 659). The lack of well-defined *explananda* is consistent with a view of the Blue Brain and of the Human Brain projects as aspiring to build (technological, computational, and conceptual) *frameworks* for the explanation of brain behaviours, rather than to formulate full-fledged explanations of particular kinds of behaviours. And it is also consistent with a data-oriented interpretation of these projects, according to which their prominent goal is to generate data on the target behaviours rather than to explain them.

As pointed out by the same authors, Eliasmith's Cognitive Computation project has a different nature. Indeed, it has a relatively well-defined and fairly articulated *explanandum* to address. The Spaun model is a huge neural network able to perform eight tasks, including copy drawing, image recognition, a task involving reinforcement learning, a serial working memory task, counting, question answering, rapid variable creation, and fluid reasoning (Eliasmith *et al.*, 2012). Eliasmith is careful to note that «the central purpose of this work is not to explain any one of these tasks, but to propose a unified set of neural mechanisms able to perform them all. In a sense, the complex task solved by Spaun is one of coordination. That is, the rapid flexibility of biological systems is its target of explanation» (p. 1024; see also Donnarumma *et al.*, 2012).

To sum up, while clear descriptions of *explananda* to be addressed are missing in Markram's projects, Eliasmith's simulation comes with a relatively precise definition of the tasks it can perform and of the *explanandum* it may contribute to addressing. And the fact that his neural network model is able to perform a variety of tasks can be considered one of the most distinctive features of the Cognitive Computation project with respect to other model-based AI and biorobotics simulation studies (focused on the production of a much more limited behavioural repertoire) and, especially, to other large-scale simulation projects (lacking a clear description of their *explananda*). As pointed out by Eliasmith.

Although impressive scaling has been achieved, no previous large-scale spiking neuron models have demonstrated how such simulations connect to a variety of specific observable behaviors. The focus of this past work has been on scaling to larger numbers of neurons and more detailed neuron models. Unfortunately, simulating a complex brain alone does not address one of the central challenges for neuroscience: explaining how complex brain activity generates complex behaviour (Eliasmith *et al.*, 2012 p. 1202).

The ambition to model many different behavioural phenomena in one simulation is conceptually distinct from the ambition to go downward through the mechanistic decomposition hierarchy (Section 3.1), as penetrating into lower levels of analysis needs not have effect on the behavioural repertoire of the system, and from the ambition to simulate huge numbers of base units (Section 3.2), even though in some cases – as mentioned above – adding same-level mechanisms may increase the number of capacities displayed by the system. And, similarly to the previous section, it is legitimate to ask what is the explanatory and predictive value of building such behaviourally rich simulations.

This question has an easy answer as far as the data-oriented role of simulations is concerned. Having one system that reproduces many aspects of the behaviour of the target system may

enable one to predict many different behavioural effects of the same intervention or stimulation. For example, it might assist one in identifying the consequences of the delivering of particular drugs on various sensory and motor modalities at the same time. On the model-oriented side, over and above the obvious epistemic value of a unified theory able to encompass a variety of behavioural phenomena, the experimental role of such a simulation is likely to depend on the characteristics of the *explanandum* that is addressed from case to case: if that *explanandum* concerns just one form of behaviour, it is not clear why a simulation succeeding in the reproduction of that behaviour would be less valuable from an epistemic point of view than a simulation able to reproduce that behaviour plus other ones. And one may reasonably ask whether such a striving for unification can be justified even though “good” neuroscientific explanations for many particular forms of behaviours are still awaited.

3.4. *Abstraction in the model of the base units*

A research group led by IBM researcher Dharmendra Modha has built a massively parallel cortical simulator, called C2, «with 1.617×10^9 neurons and 0.887×10^{13} synapses, roughly 643 times slower than real-time per Hertz of average neuronal firing rate. The model used biologically-measured gray matter thalamocortical connectivity from cat visual cortex» (Ananthanarayanan *et al.*, 2009, p. 1). C2 approaches the scale of a cat’s brain in terms of the number of neurons and synapses involved, and is often called a “cat’s scale simulation” in the scientific literature and in the press. Modha claims that large-scale simulators, including C2, «have tremendous potential implications for theoretical and applied neuroscience as well for cognitive computing. The simulator is a modern-day scientific instrument, analogous to a linear accelerator or an electron microscope, that is a significant step towards unraveling the mystery of what the brain computes and towards paving the path to low- power, compact neuromorphic and synaptronic computers of tomorrow» (Ananthanarayanan *et al.*, 2009, p. 10).

Large-scale simulations, according to Modha, are therefore of theoretical *and* technological interest, as they may lead to the development of novel and ever more powerful computational techniques and devices. It is worth noting that the development of Modha’s simulator responds to the “SyNAPSE” call launched by DARPA in 2008, whose goal was to promote the development of «electronic neuromorphic machine technology that scales to biological levels. More simply stated, it is an attempt to build a new kind of computer with similar form and function to the mammalian brain. Such artificial brains would be used to build robots whose intelligence matches that of mice and cats»⁶. And in a press report Modha declared that his Cognitive Computing Project “is the quest to engineer mind-like intelligent business machines by reverse engineering the computational function of the brain and packaging it in a small, low-power chip”⁷. It therefore seems that Modha’s principal interests are technological.

The simulation built by Modha’s team has been vividly criticized by Henry Markram in a letter sent to the IBM Chief Technical Officer in 2009. In the letter, Markram dubbed Modha’s reports on C2 as hoaxes and claimed that it was “shameful and unethical” to call C2 a simulation of the cat’s brain. Markram’s point was that the neuron models used by Modha «are *point neurons* [with] *no branches*; *no detailed ion channels*; *the simplest possible equation you can imagine* to simulate a neuron, *totally trivial synapses*, [...] All these kinds of simulations are trivial and have been around for decades - simply called artificial neural network (ANN) simulations. [...] It is really no big deal to *simulate a billion points interacting* if you have a big enough computer».⁸

⁶ From <http://www.artificialbrains.com/darpa-synapse-program>, visited on 19/09/15.

⁷ From <http://www.technewsworld.com/story/65237.html>, visited on 19/09/15.

⁸ From <http://spectrum.ieee.org/tech-talk/semiconductors/devices/blue-brain-project-leader-angry-about-cat-brain>, visited on 19/09/15.

For some commentators Markram has neglected the fact that the goals of Modha's project are prominently technological, and that for this reason he cannot be criticized for having modelled too "simple" neurons. Over and above the details of this controversy, Markram's reaction illustrates another ambition pursued in large-scale simulation projects, which can be distinguished from those discussed in the previous section. It concerns the "complexity" of the model of the base units of the mechanism, which in those simulations consist in artificial neurons. One thing is to simulate a point neuron with a simple input-output function and few dendrites; another thing is to simulate neurons with a realistic degree of dendritic and axonal arborisation, and with "complex" input-output characteristics. Markram often emphasizes that Blue Brain's neurons are "biologically realistic" in this sense of the term, and that this kind of biological realism is to be praised in a large-scale simulation:

At this point, some may ask, why not use this computing power to simulate cortical circuits with artificial neural networks, in which the entire neuron is represented by one summing node (point neuron), connectivity is simplified to reciprocal interactions between all nodes, and functional properties are simplified as 'integrate and fire' types of activity. Such simulations provide a powerful exploratory tool, but the lack of biological realism severely limits their biological interpretation (Markram, 2006, p. 154).

Note that this dimension of "biological realism" is conceptually independent from the dimensions analysed so far. One can change the model of the base unit without changing level of mechanistic decomposition hierarchy (Section 3.1), the size of the model (Section 3.2), and the number of modelled phenomena (Section 3.3).

Markram's claim that "the lack of biological realism" in terms of the complexity of the base unit model limits the "biological interpretation" of the simulation is not self-explanatory, however. With reference to a data-oriented use of simulations, one may reasonably reply that only those aspects of the neuron which make a detectable difference to the overall system behaviour should be included in the model. There are no general reasons to pursue complexification of the base unit in every case. Similarly, Markram's claim can be reformulated in connection with the model-oriented method as stating that plausibility of theoretical model ML improves with the complexity of the base unit. This epistemological assumption is arguably made in other computer-based and biorobotics simulation studies using very complicate model neurons, such as (Reeke, Sporns and Edelman, 1990). Here again, many philosophical analyses of what makes a "good" mechanistic explanation converge in asserting that all and only the causally relevant factors – that is to say, all and only the factors that make a difference with respect to the phenomenon to be explained – should be mentioned in the theoretical model.

For this reason, one may well turn Markram's above claim upside down and conclude that a hyper-complex simulation may actually "provide a powerful exploratory tool", as it may enable one to finely intervene on various aspects of the model neurons (for example, by changing the degree of arborisation or the characteristics of the input-output function) and to assess whether this makes the difference on the target behaviour. A "good" explanatory model would be then obtained by removing all the supposedly irrelevant factors. The model-oriented simulation study on the cellular mechanisms of neuronal synchronization in epilepsy reported in (Traub and Wong, 1982) and cited in (Markram, 2006) proceeds in a similar way, that is to say, by building a detailed (in the sense analysed here) 100-neuron simulation of a brain circuit and – through finely-grained interventions – by identifying features of the model neurons that are not critical to the production of the behaviour to be explained. Model-oriented computer-based and biorobotic studies often involve processes of explanation of machine behaviour, aimed at discovering what aspects of the system are really relevant with respect to the target behaviour, and what aspects have negligible effects instead (see for example Grasso, 2000; Reeve *et al.*, 2005).

3.5. *The plausibility of the theoretical model*

Finally, let us point out that another goal of contemporary large-scale simulation projects is to simulate *plausible*, that is to say, well corroborated theoretical models of the target system. The leaders of these projects put great emphasis on the fact that their simulations are to be based on “quantitative data” concerning “anatomical and physiological properties of the synaptic connections and the ion channels that support the different types of electrical behaviour” (Markram, 2006, p. 156) obtained in experiments performed on biological tissues. Note that the point here is not (only) the level of the decomposition analysis: ideally, one can build a simulation that extends downward to the level of ion channels and, at the same time, build models of ion channel behaviour which conflict with what is known on the behaviour of biological ion channels. Similarly, one can implement very large (Section 3.2) simulations composed of very complicated (Section 3.4) neuron models which are nevertheless known to be false. The point is whether incorporating plausible theories on the target system is a *desideratum* in large-scale simulation projects. In Markram’s and other leading neuroscientists’ opinion, it is.

We have already discussed the significance of this choice in Section 2 in connection with the distinction between data-oriented and model-oriented simulation studies. Markram and colleagues’ emphasis on the model plausibility requirement suggests that they expect their simulations to play data-oriented roles, insofar as the truth of the underlying model is a key requirement of data-oriented simulations. Model plausibility is also essential for the integrative role of multilevel simulations, insofar as they are expected to assist in building bridges between bodies of accepted knowledge on the brain. However, for the reasons discussed above, the truth of ML is not required in a “good” model-oriented study, and the synthetic method has been often adopted through the XX century to test conjectural models whose plausibility was still largely unknown.

4. *Summary and conclusions*

A century has passed since the construction of Hammond and Miessner’s “electric dog” inspired by Loeb’s theories on moth phototropism. Computational and robotic technologies have undergone dramatic improvements in the last decades, and the use of machines as models of biological and cognitive processes has become widespread in the scientific community. However, several methodological questions – many of which were highly debated in the early stages of Cybernetics and Artificial Intelligence, as extensively discussed by Cordeschi in his philosophical and historical works – are still open. To be sure, contemporary scientific literature shows substantial ambiguities concerning the roles of simulations in neuroscientific research, and the methodological constraints that “good” simulation studies are required to meet.

In this exploratory article we have pointed out that simulations can contribute to neuroscience and cognitive science in at least three different ways, namely by supporting (a) the discovery of mechanisms (model-oriented studies), (b) the acquisition of data on target system behaviours (data-oriented studies), and (c) the integration of knowledge coming from different studies and sources. Leaders of contemporary large-scale brain simulation projects, however, are often vague as to which of these roles their simulations are expected to play. The issue is crucial since these roles are rather different from each other, require different experimental strategies, and place different constraints on the features of the simulation system and on the theoretical model implemented in it. For example, we have argued that a “good” data-oriented study requires one to implement the best theories available on the target system, whereas this strong epistemic requirement can be relaxed in model-oriented studies. A lack of clarity on the role a simulation is expected to play in neuroscientific research implies, to a significant extent, a corresponding lack of clarity on the features that the simulation is expected to display. In one important case – the Blue Brain Project – a precise and clear statement of the research questions addressed in the project is even missing.

What we have called here the data-oriented use of simulations is rather different from the model-oriented use of computers and robots extensively made in Artificial Intelligence and biorobotics. Contemporary large-scale brain simulation projects therefore show what the role of machines might be in the study of brain and cognition over and above the “synthetic” method. They also illustrate the aim of pushing biological accuracy to unprecedented levels, by exploiting the extraordinary computational potentialities of new-generation supercomputers. In this article we have distinguished between some conceptually distinct modelling ambitions pursued in these projects, notably, the ambition of reaching extremely “low” levels in the mechanistic decomposition hierarchy, of simulating huge quantities of base units, of generating multiple behaviours by the same system, of simulating “complex” base units, and of simulating explanatory and highly corroborated models of brain activity. And we have argued that the significance of these efforts with respect to the explanation and prediction of brain activity is not always clear. Moreover we stress that many of the epistemological and methodological questions that these efforts give rise to, and that have been introduced in this article, ultimately reduce to one of the most important and hard conceptual problems in the philosophy of science, that is to say, the problem of understanding what makes a “good” mechanistic explanation.

References

- Ananthanarayanan R., Esser S. K., Simon H. D., Modha D. S. (2009). The cat is out of the bag : cortical simulations with 10^9 neurons, 10^{13} synapses, *Proceedings of the Conference on High Performance Computing Networking, Storage and Analysis*, (c): 1-12.
- Bogen J., Woodward J. (1988). Saving the Phenomena, *The Philosophical Review*, 97, 3: 303.
- Burgess N., Jackson A., Hartley T., O’Keefe J. (2000). Predictions derived from modelling the hippocampal role in navigation, *Biological Cybernetics*, 83, 3: 301-312. Doi: 10.1007/s004220000172.
- Cordeschi R. (2002). *The Discovery of the Artificial. Behavior, Mind and Machines Before and Beyond Cybernetics*, Dordrecht: Springer Netherlands. Doi: 10.1007/978-94-015-9870-5.
- Cordeschi R. (2008). Steps toward the synthetic method: symbolic information processing and self organization systems in early Artificial Intelligence. In :Husbands P., Holland O., Wheeler M. a cura di. *The mechanical Mind in History*, Cambridge MA: The MIT press: 219-258.
- Craver C. (2007). *Explaining the Brain: Mechanisms and the Mosaic Unity of Neuroscience*. New York and Oxford: Clarendon Press. Doi:10.1093/acprof:oso/9780199299317.001.0001.
- Datteri E. (2016). Biorobotics. In: Magnani L., Bertolotti T. W., a cura di., *Springer Handbook of Model Based Science*. Heidelberg/Berlin: Springer.
- Datteri E., Laudisa F. (2014). Box-and-arrow explanations need not be more abstract than neuroscientific mechanism descriptions, *Frontiers in Psychology*, 5, 464. Doi: 10.3389/fpsyg.2014.00464.
- Datteri E., Tamburrini G. (2007). Biorobotic Experiments for the Discovery of Biological Mechanisms, *Philosophy of Science*, 74, 3: 409-430. Doi: 10.1086/522095.
- De Garis H., Shuo C., Goertzel B., Ruiting L. (2010). A world survey of artificial brain projects, Part I: Large-scale brain simulations, *Neurocomputing*, 74, 1-3: 3-29.
- Donnarumma F., Prevete R., Trautteur G. (2012). Programming in the brain: a neural network theoretical framework, *Connection Science*, 24, 2-3: 71-90. Doi: 10.1080/09540091.2012.684670.
- Dror R. O., Dirks R. M., Grossman J. P., Xu H., Shaw D. E. (2012). Biomolecular simulation: a computational microscope for molecular biology, *Annual Review of Biophysics*, 41: 429-52. Doi: 10.1146/annurev-biophys-042910-155245.
- Eliasmith C., Stewart T. C., Choo X., Bekolay T., DeWolf T., Tang C., Rasmussen D. (2012). A Large-Scale Model of the Functioning Brain, *Science*, 338, 6111: 1202-1205. Doi: 10.1126/science.1225266.

- Eliasmith C., Trujillo O. (2014). The use and abuse of large-scale brain models, *Current Opinion in Neurobiology*, 25: 1-6. Doi:10.1016/j.conb.2013.09.009.
- Friedman M. (1974). Explanation and Scientific Understanding, *The Journal of Philosophy*, 71, 1: 5-19.
- Gijssbers V. (2007). Why Unification Is Neither Necessary Nor Sufficient for Explanation. *Philosophy of Science*, 74, 4: 481-500. Doi: 10.1086/524420.
- Grasso F. (2000). Biomimetic robot lobster performs chemo-orientation in turbulence using a pair of spatially separated sensors: Progress and challenges, *Robotics and Autonomous Systems*, 30, 1-2: 115-131.
- Humphreys P. (2004). *Extending ourselves: Computational science, empiricism, and scientific method*. New York: Oxford University Press. Doi: 10.1093/0195158709.001.0001.
- Kandel E. R., Markram H., Matthews P. M., Yuste R., Koch C. (2013). Neuroscience thinks big (and collaboratively). *Nature Reviews. Neuroscience*, 14: 659-64. Doi:10.1038/nrn3578.
- Kitcher P. (1981). Explanatory Unification, *Philosophy of Science*, 48(4): 507-531.
- Loeb J. (1900). *Comparative physiology of the brain and comparative psychology*. New York: G.P. Putman's Sons. Doi: 10.5962/bhl.title.1896.
- Long J. (2012). *Darwin's Devices. What Evolving Robots Can Teach Us About the History of Life and the Future of Technology*, New York: Basic Books.
- Markram H. (2006). The blue brain project. *Nature Reviews. Neuroscience*, 7: 153-160. Doi:10.1038/nrn1848.
- Markram H., Meier K., Lippert T., Grillner S., Frackowiak R., Dehaene S., Knoll A., Sompolinsky H., Verstreken K., DeFelipe J., Grant S., Changeux J.P., Saria A. (2011). Introducing the Human Brain Project. *Procedia Computer Science*, 7, 39-42. Doi: 10.1016/j.procs.2011.12.015.
- Milkowski M. (2015). Explanatory completeness and idealization in large brain simulations: a mechanistic perspective. *Synthese*: 1-22. Doi: 10.1007/s11229-015-0731-3.
- Newell A. (1990). *Unified Theories of Cognition*, Cambridge, MA: Harvard University Press.
- Parker W. S. (2009). Does matter really matter? Computer simulations, experiments, and materiality. *Synthese*, 169, 3: 483-496. Doi: 10.1007/s11229-008-9434-3.
- Reeke G. N., Sporns O., Edelman G. M. (1990). Synthetic neural modeling: the "Darwin" series of recognition automata. *Proceedings of the IEEE*, 78, 9: 1498-1530. Doi: 10.1109/5.58327.
- Reeve R., Webb B., Horchler A., Indiveri G., Quinn R. (2005). New technologies for testing a model of cricket phonotaxis on an outdoor robot. *Robotics and Autonomous Systems*, 51, 1: 41-54. Doi: 10.1016/j.robot.2004.08.010.
- Rosenblueth A., Wiener N. (1945). The Role of Models in Science, *Philosophy of Science*. 12, 4: 316-321. Doi: 10.1086/286874.
- Roysam B., Shain W., Ascoli G. A. (2009). The central role of neuroinformatics in the national academy of engineering's grandest challenge: Reverse engineer the brain. *Neuroinformatics*, 7, 1: 1-5. Doi: 10.1007/s12021-008-9043-9.
- Suppe F. (1989). *The Semantic Conception of Theories and Scientific Realism*, Urbana and Chicago: University of Illinois Press.
- Traub R., Wong R. (1982). Cellular mechanism of neuronal synchronization in epilepsy. *Science*, 216, 4547: 745-747. Doi: 10.1126/science.7079735.
- Webb B. (2006). Validating biorobotic models, *Journal of Neural Engineering*, 3, 3: 25-35. Doi: 10.1088/1741-2560/3/3/R01.
- Winsberg E. (2003). Simulated experiments: Methodology for a virtual world, *Philosophy of Science*, 70: 105-125. Doi: 10.1086/367872.
- Woodward J. (2002). What Is a Mechanism? A Counterfactual Account, *Philosophy of Science*, 69: 366-S377. Doi: 10.1086/341859.