



Why is There No Successful Whole Brain Simulation (Yet)?

Klaus M. Stiefel^{1,3} · Daniel S. Brooks²

Received: 14 June 2018 / Accepted: 1 March 2019
© Konrad Lorenz Institute for Evolution and Cognition Research 2019

Abstract

With the advent of powerful parallel computers, efforts have commenced to simulate complete mammalian brains. However, so far none of these efforts has produced outcomes close to explaining even the behavioral complexities of animals. In this article, we suggest four challenges that ground this shortcoming. First, we discuss the connection between hypothesis testing and simulations. Typically, efforts to simulate complete mammalian brains lack a clear hypothesis. Second, we treat complications related to a lack of parameter constraints for large-scale simulations. To demonstrate the severity of this issue, we review work on two small-scale neural systems, the crustacean stomatogastric ganglion and the *Caenorhabditis elegans* nervous system. Both of these small nervous systems are very thoroughly, but not completely understood, mainly due to issues with variable and plastic parameters. Third, we discuss the hierarchical structure of neural systems as a principled obstacle to whole-brain simulations. Different organizational levels imply qualitative differences not only in structure, but in choice and appropriateness of investigative technique and perspective. The challenge of reconciling different levels also undergirds the challenge of simulating and hypothesis testing, as modeling a system is not the same thing as simulating it. Fourth, we point out that animal brains are information processing systems tailored very specifically for the ecological niches the respective animals live in.

Keywords Ecological niches · Hypothesis testing · Levels of organization · Whole-brain simulation

Introduction

Materialist philosophy dictates that brain function should be explainable in terms of physical processes. Modern neuroscience has made indeed impressive progress in finding such explanations. To show the consistency of these results, and extend their insight into new and emerging questions, computational neuroscience aims to construct quantitative models of neuronal dynamics. Exciting progress has been made in simulating subcellular, cellular, and network models. Naturally progressing from these simulations of parts of nervous systems are efforts to simulate complete animal brains. In the last decade several ambitious projects have been attempting to simulate the whole rat (Markram et al.

2015), cat (Ananthanarayanan et al. 2009; Merolla et al. 2014), and generalized mammalian (Izhikevich and Edelman 2008) brains. A project by Eliasmith used a simulation of 2.5 million neurons to operate a simulated robotic arm performing a variety of tasks (Eliasmith et al. 2012). These projects differ in their aims, with some attempting bottom-up assemblies of data, and others trying to connect sensation to behavior with a brain-like architecture. What they have in common are large numbers of simulated neurons (of varying levels of complexity) and claims that their simulation encompasses the dynamics of a complete or whole brain.

Despite massive investments in computing hardware, these attempts have not (yet?) produced results that are remotely satisfactory for explaining human or animal cognition and behavior on the basis of brain function.

In this article, we outline several principled reasons that we believe contribute to this shortcoming. We start by discussing the role of simulations in hypothesis testing and the creeping, rather than sweeping, nature of progress in neuroscience. We then review the work on two of the most thoroughly understood neural systems, the crab stomatogastric nervous system and the complete *Caenorhabditis elegans*

✉ Klaus M. Stiefel
klaus@neurolinx.org

¹ Neurolinx Research Institute, La Jolla, CA, USA

² Department of Philosophy, University of Cincinnati, Cincinnati, OH, USA

³ Marine Science Institute, University of the Philippines, Dilliman, Quezon City, Philippines

nervous system. We outline the difficulties that the researchers attempting to understand these small systems have faced, especially in regards to parameter constraints for simulations. The issues faced when simulating small nervous systems harbor lessons for modeling the immensely more complex vertebrate nervous systems. Additionally, the multilevel nature of neural computation is likely more important when simulating larger nervous systems. We also discuss how the specific ecological specialization of an animal relates to its brain function, and to the simulation of this brain function.

Interestingly, due to the issues mentioned above the literature lacks a consistent definition of a whole brain simulation, which is usually considered a simulation of a large number of neurons, approaching the number of neurons in an animal's brain, often at a level of detail currently only achievable for smaller simulations (Markram 2006).

Simulations as Quantitative Hypothesis

Testing hypotheses is at the core of the scientific method. A numerical simulation of a neural system corresponds to a set of inferences concerning a theoretical explanation of that system (Winsberg 2001). These inferences concern predictions about the system's behavior over time. This often amounts to testing the quality or consistency of that explanation compared to different sets of experimental data. For example, we can ask if our knowledge of ion channel properties and distributions and cellular morphology is sufficient to explain the synaptic potential waveform observed in a single neuron. Since this empirical knowledge is complex, no one can figure out the answer by simply reading the measurements and thinking about them. Human intuition is not capable of deducing a potential waveform from dozens of nonlinear relationships between components of the neuron. Numerical simulations thus provide a solution to this dilemma by calculating the dynamics of the all-biological components represented in the model (Izhikevich 2010).

However, the usefulness of a simulation depends on a number of factors. Simulations do not provide explanations for a phenomenon per se, but rather a measurable means of imitating a real-world system under a set of idealizing assumptions (Hartmann 1996; Winsberg 2003). When imitation is a goal or requirement of the use of a simulation, criteria for evaluating the fidelity of the simulation in capturing aspects of the target system need to be elaborated. For this, a mediating theoretical model is usually required (Krohs 2008, pp. 283–284; Winsberg 1999). A mediating model is one that coordinates the workings of a phenomenon of interest (given a description derived by the mechanism that the model's workings summarize) with the simulation's dynamics. "Within the context of simulations," Krohs explains, "primarily dynamic models are of interest (Hartmann 1996,

pp. 82–83), especially those that refer to the internal mechanism bringing about the dynamics. Such models may be regarded as not only describing, but also as explaining, the process under consideration" (Krohs 2008, p. 178). That is, a theoretical model is needed to mediate between the simulation and the world, with the model identifying (perhaps ostensibly) the parts of the real-world system that figure in producing the phenomenon to be explained. This coordinates the phenomenon with the simulation, in turn creating standards against which the simulation can be judged as to whether it is faithfully imitating the part of the world it is supposed to mimic.

An example of a mediating model is the Reichardt motion detector model used to represent (and sometimes simulate) motion adaptive responses in the fly brain (Borst and Egelhaaf 1989). The Reichardt model is a correlation-type motion detector that offers a mathematical description of light entering individual ommatidia of the fly's eyes from two input channels, which become correlated at later stages of processing. On the one hand, the Reichardt detector is a proxy for numerous anatomical and electrophysiological studies of the fly brain (for a summary and history see Brooks 2014). These studies provide crucial mechanistic details for the workings of motion adaptation, i.e., the ability of the visual system to modulate its response to continuous stimulation. On the other hand, the Reichardt detector acts as a model upon which to simulate motion adaptive scenarios. For instance, stimulation of the visual system with motion will first provoke a high response from motion-sensitive tangential cells (measured in number of spikes per second), but then will rapidly decay towards a semi-steady state upon continuous stimulation with the same motion velocity. This stereotypical response profile can be approximately predicted by simulating motion stimulation with a Reichardt detector (de Ruyter van Steveninck et al. 1986).

Thus, the Reichardt motion detector mediates between a simulation and the real-world details of the target system of investigation, here the visual system of the fly. It coordinates between explanations of the phenomenon, given by the mechanistic details that inform the overall structure of the model, and the output simulations that seek to track and predict neuronal output when exposed to various motion stimuli. The upshot of this mediating relation is that simulations based on such a theoretical model can be assessed according to their ability to faithfully mimic the real-world counterpart to the simulation. More importantly, this mediating relation works even in the absence of a full mechanistic explanation, as when exploratory use of simulations is meant to help identify or supplement the characterization of the real-world target system provided by a model (Egelhaaf and Borst 1993). Obviously, the extent to which simulations correspond (or don't) to their target systems remains a complicated and contingent question (Parker 2009).

In addition to mediating models, the usefulness of simulations is based on other factors as well. Another important factor is the rule that the hypotheses used with simulations should clearly identify what it is about the target system that is being tested. In our first example, we may ask if our knowledge of synapse location, synaptic glutamate receptor kinetics, dendritic thickness and branching pattern, and voltage-gated ion-channel kinetics and distribution explain the rising slope of the postsynaptic potentials measured at the neuron's soma. Note that we are testing a specific question (about the synaptic potential's rising slope) here, and are not attempting to *first* recreate a universal artificial neuron. Relatedly, the questions asked also determine the degree of abstraction used (ion channels and dendrites, in our example). Besides confirming or falsifying the consistency of our understanding of the system, the simulations provide predictions. These predictions will be at the chosen degree of abstraction, about properties of the cell that have not been used to construct the model, and possibly have not been measured yet. An example would be a prediction about the waveform of a second synaptic potential evoked briefly after the first one. The waveform would be altered due to remaining ion channel activation stemming from the first synaptic potential. We have again asked a very specific question about a biological system at one degree of abstraction, and have gotten an answer, plus answers to closely related questions about the same system. We have not, for instance, gained any answers about the possible rearrangement of the cytoskeleton during synaptic activation (no equations in our model describe the cytoskeleton, since we assumed that the synaptic potentials are not influenced by its dynamics on a fast scale).

The situation is altogether different when trying to simulate the brains of mammals with a high level of biological accuracy. In this case, the numerical simulation is supposed to answer *any* reasonable question about the neural system under investigation—the whole brain. This, in our opinion, puts the cart before the horse. The direction a simulation takes in scientific investigation should be guided by some criteria of usefulness like those we detailed above. The massive increase in computational power needed to simulate such an open-ended and hence much more expansive system (Markram 2006) is a technical issue (albeit a hard one), and will not concern us any further here.

The lack of a clear hypothesis and lack of mediating model when attempting to simulate a complete mammalian brain is one issue, and a fundamental one, which impedes such an undertaking. Ideally, the hypothesis will stand at the beginning of the simulation project and will inform the conceptual work that shapes the simulation code. The elements of the simulation will correspond to the questions raised by the hypothesis; it would be difficult to impose a hypothesis onto a simulation that was constructed with a different aim.

The type of progress historically seen in neuroscience is at odds with the breakthrough type of progress promised by such whole-brain simulations. We will discuss this issue next. Then we will elaborate on the massive complexities of mammalian brains, which also have fundamental consequences for attempted whole-brain simulations. These consequences are the role of parameter constraints and of multiple levels of organization in neurobiology. Finally, we will discuss the specialized roles of animals in their ecological niches, and how that affects efforts to simulate complete brains and their behavioral outputs.

Parameter Constraints

Small Nervous Systems

To understand the difficulties arising from simulating biological systems with multiple parameters at multiple levels let us first look at efforts to simulate much smaller nervous systems than those of mammals. Nervous systems in the animal kingdom span many orders of magnitude in their complexity. Two well-studied small nervous systems are the decapod crustacean (crab and lobster) stomatogastric ganglion (STG) and the nematode worm nervous system.

The following paragraphs are in no way meant to be a comprehensive review of the neurobiology of the STG, but rather to give an idea of the complications faced when trying to understand a ganglion composed of only a small number of neurons. These complications then compose a stepping-stone to estimating the challenges when attempting to simulate much larger complete mammalian brains.

The STG is not a complete nervous system, but a central pattern generator producing the peristaltic rhythms of the crustacean digestive system. While it only constitutes a part of the crab's nervous system, it acts in relative isolation from the sensory and motor components. Its neurons are also relatively large and easy to record from, and the same neurons occur in every individual crustacean. All of this makes it an attractive study object. A total of 30 neurons in six classes generate the two rhythms (pyloric and gastric) produced by the STG. The neurons are connected with chemical and electrical synapses. Some of the neurons are electrically quiet when not stimulated, while others are intrinsic oscillators. These oscillations occur due to an interplay of de- and hyperpolarizing ion channels. The intrinsic oscillators force their rhythmicity onto the intrinsically quiet neurons, and the complete network generates the intestinal rhythms of use to the crab (Selverston 2008).

The ion channel densities of the currents giving rise to these oscillations are curiously not conserved between neurons. While each neuron of one type in every individual crab produces the same output patterns, these patterns are

the result of very different ion channel density combinations. Every neuron has an individual ion channel parameter set, which generates identical output dynamics (Schulz et al. 2006). Interestingly, the STG produces very similar, albeit somewhat faster output patterns when the ambient temperature rises. This is a very important issue since the body temperature of marine invertebrates is identical to the ambient temperature, which in temperate oceans varies greatly over the course of a year (Marder et al. 2015).

A number of neuromodulators also change the properties of the neurons in the STG. These messenger molecules are broadcast from other parts of the crustacean nervous system to the STG, and they modify the behavior of individual neurons by altering some of their ion channel conductances. These neuromodulators cause a reconfiguration of the whole STG, and cause it to output different rhythms (Hamood et al. 2015). Their effect, when looking at it from the point of view of a simulation, amounts to a change of multiple parameters causing altered output patterns.

Hence, even a neural system composed of only 30 neurons has turned out to be surprisingly complex. While 30 neurons is a small number for a neuroscientist, a system of 30 coupled nonlinear oscillators (each describing one neuron, with several dimensions each) is massive for a mathematician. Closed form (analytic) solutions of such systems are realistic only for low dimensions (and impossible even in many low-dimensional systems), and are out of the question here. It's necessary to numerically simulate the STG to predict its output.

Work in nonlinear time series analysis and machine learning provides an estimate for the amount of data necessary for a dynamical system of a certain dimensionality. A variety of methods have been used for this purpose, such as maximum likelihood methods, Kalman filters, Monte Carlo statistics-based methods, and Bayesian methods (Ching et al. 2006; Toni et al. 2009). Naturally, the number of data points increases when the dynamical system becomes more complex. Relationships between dimensionality and the amount of data required for a parameter estimate depend on the exact algorithm, but encompass computational complexities similar to $O(L^3)$ (with L being the state dimension, the computational requirements growing with the third power of the state dimension; van der Merwe and Wan 2001).

Hence, a dimensionality that seems low to moderate to the neuroscientist is in fact massive when it comes to mathematical treatment and parameter estimation for numerical simulation.

One magnitude larger in terms of the number of neurons is the nematode *C. elegans*, a small worm with a total of 302 neurons. As in many other invertebrate nervous systems, these neurons have reproducible identities between individual worms (each worm has the same number and types of neurons in the same positions). *C. elegans* has served as a

very popular subject for research in molecular developmental biology (Meister et al. 2010; Varier and Kaiser 2011), and we know a great deal about this small organism, including its complete genome (Gerstein et al. 2010). Several attempts have been made to integrate the knowledge about the *C. elegans* nervous system in quantitative models (Si elegans: Blau et al. 2014; Machado et al. 2014; OpenWorm: Palyanov et al. 2012; Szigeti et al. 2014), but none of these projects has produced a complete working model yet. In contrast to the STG, the neurons of *C. elegans* are small, and it is difficult to record from them. This difficulty has only recently been mastered with the development of imaging methods that allow multi-neuron recordings in freely moving worms (Faumont et al. 2011; Schrödel et al. 2013). The data to begin simulating a whole-worm nervous system has thus not been available for very long. The lack of a functioning complete model of this small nervous system might therefore be due either to fundamental difficulties or to the short time frame modelers have had for the task.

The recently initiated OpenWorm project (Szigeti et al. 2014) is a new effort at a complete simulation of this small animal. Interestingly, the project stresses an open-source approach and flexibly incorporating new data from a community of researchers for testing a variety of hypotheses, rather than one “correct” master model of the worm. The project includes a simulation engine that can cope with models at different levels of complexity, and a fluid dynamics simulation of the body of the worm. The latter module incorporates a feedback from the physical system (ecosystem, in a simple way) surrounding the worm into the simulation. This integrated, open-ended approach is well in accord with some of the points we make here.

The common theme we find in the attempts to simulate the STG and the *C. elegans* nervous system are nontrivial problems with a lack of parameter constraints. In the case of the more thoroughly explored STG, different parameter regimes (due to the individuality of neurons and different channel kinetics at different temperatures) can lead to surprisingly similar neural output. Conversely, a moderate modification of cellular parameters via neuromodulators can completely change the behavior of the system. In the case of *C. elegans*, despite the small number of neurons and a completely known neuroanatomy, no whole-nervous system approach has come to fruition yet—probably at least partially due to a lack of neurophysiological data until recently.

The issues regarding a lack of parameter constraints are not only due to a lack of experimentation and subsequent parameter fitting of models. Rather, this seems to be a fundamental issue, and involves significant variability between individuals, parameter modification in response to changes in external conditions and internal modulators, and many-to-one mappings of parameter sets to model behaviors. There seems to be no one “correct” set of parameters for

the complete STG. We have hence seen a number of serious difficulties in modeling supposedly simple nervous systems, and can next ask how these issues will translate to mammalian nervous systems.

Large Nervous Systems

The STG with its 30 neurons and the *C. elegans* nervous system with 302 neurons are both much simpler than the nervous systems of the mouse (approximately 7×10^7 neurons), cat (10^9 neurons), or a human (8×10^{10} neurons). How do the aforementioned issues of incomplete parameter constraints and brain-body dependence affect simulations of such much larger nervous systems? There is no reason to assume that they don't weigh proportionally heavier in these much larger and massively more complex (10^5 times the number of neurons, mouse versus *C. elegans*) nervous systems.

An issue that could in fact worsen the problem of parameter constraints is that neurons in vertebrates are not developmentally deterministic; unlike in most invertebrates, where each neuron can be reproducibly identified across individuals, no such determinism exists in vertebrates. Thus, instead of dealing only with variable ionic conductance densities giving rise to the same function in otherwise identical neurons (as in the STG), modelers of cortical circuits could be facing local circuits with identical function but comprising variable neuron numbers, types, and connectivity patterns. A computational scientist simulating large nervous systems will likely be facing the issues regarding constraining the parameters seen in small nervous systems, and additional issues peculiar to mammalian nervous systems.

Levels of Organization

Tightly connected with this issue of parameter constraints (or lack thereof) is the aforementioned concept of organizational levels (Sejnowski and Churchland 1994; Craver 2007). Levels of organization comprise compositional layers in nature¹: neural tissue is organized in modules, composed of structured assemblies of modules of the lower level. Brain areas are composed of neural networks, and these networks

are in turn composed of neurons, which are composed of membranes and ion channels. And so on. Indeed, the idea of levels is widespread not only to the neurosciences, but in and across most biological disciplines (Wimsatt 2007). Germane to this view is that there are qualitative differences inherent to each level that demarcate not only natural units (e.g., cells, genes, proteins, and molecules), but also research regimes that make up independent ventures into the brain. This, for one thing, makes comprehensive understanding of even any one level, in isolation, insufficient to explain or understand whole phenomena under consideration. That is, knowledge of cells does not exhaust knowledge of tissues, nor molecules or proteins. Instead, levels of organization in a complex natural system such as the brain indicate that multiple research areas and even disciplines will overlap in their efforts to explain the phenomena engendered by that system.

Though the intent of these multiple research areas may very well converge on explaining or understanding a common phenomenon, reconciling the differences in techniques, aims, and even the vocabulary between the different level-bound perspectives is a daunting challenge for any integrative program. This understandably makes attempts to unify or integrate multiple levels into a singular framework an extremely attractive position to advertise. However, when citing the leveled structure of the brain, one must be wary of invoking more promise than substance (Guttman 1976). It is easily granted that cells are “qualitatively different” than the things composing them; proteins, chemicals, and biomolecules simply do different things than cells do. So, what does it mean for a model, simulation, or explanation to “encompass multiple levels”? A number of specific approaches have been proposed, but until now none have reached consensus (Eronen and Brooks 2018). Minimally, we believe that any attempt to integrate levels together should include demonstrating that different putative levels are arranged in a precise way within the model such that our understanding of the target phenomenon is improved.

This highlights a number of concrete challenges for simulation and modeling. For one thing, it is not a priori clear what inter-level dependencies exist in different neural systems; these must be explored and characterized. Are the parameters on the lower level not related one-to-one to the dynamics on the higher level, as seen in the STG neurons' rhythmic output activity (several parameter combinations can produce the proper rhythmic output)? Or are there correlations between parameters on different levels? For instance, are the connection probabilities (a lower-level property) increased between neurons encoding similar stimuli (a higher-level property) in the mammalian cortex? Though this is sometimes the case (Hubel and Wiesel 1968), myriad other such inter-level rules are conceivable as well, and their existence is rarely investigated.

¹ We understand the hierarchical organization of the nervous system as one mediated primarily by part-whole relations, where what is considered a whole at one level (say, a cell) is a part at another level (say, the tissue level). This is a distinction from other conceptions of hierarchical organization, such as van Essen and Maunsell's (1983) hierarchy of functional streams in the visual cortex, and from David Marr's (1982) trilevel distribution of the algorithmic, computational, and implementational levels. The chief difference here is that these conceptions of level focus on the transmission of information between functional units, rather than track compositional relations between units of nature.

Two frequent, qualitative inter-level relations relevant for integrative brain studies, in addition to composition, include emergence and downward causation. Emergence captures the idea that features or properties located at one level appear as unpredictable or irreducible properties between levels (Stephan 1999). Downward causation, on the other hand, refers to the ability of higher-level properties to influence or change lower-level properties (Campbell 1974). Naturally, lower-level parameters can also influence higher-level parameters (such as behavior), and if the higher-level system is close to a bifurcation (at a point of high parameter sensitivity), they are more likely to do so. But conversely, causal power in systems characterized by indeterminism and/or degeneracy can be most pronounced at a macro level (Hoel et al. 2016).

Moreover, the choices involved in abstraction lie at the core of this issue of the leveled structure of the brain. Simulations and models of neural phenomena, including the whole brain, are premised on, even defined by, abstracting away from the actual details of the systems they represent in their simulating efforts. This contrasts with stronger claims of realism in models and simulations that promise highly “biologically accurate” depictions of the target phenomenon (e.g., Markram 2006; see also Almog and Korngreen 2016). Just because one can construct a simulation does not guarantee *what* one is simulating. Above we noted that successful simulations are accompanied by a mediating theoretical model, which maps the simulation’s elements at least roughly and tentatively onto real-world counterparts. In the absence of such a model, it is entirely unclear *what* is being simulated.

Turning to one simulating venture we have in mind, Markram (2006, p. 153) interestingly does refer (quite prominently) to the organizational levels of the nervous system as a primary motivation for his Blue Brain Project (BBP). It is especially noteworthy that he acknowledges particularly the “qualitative differences” between these levels. He writes:

Atoms are differentially combined to produce a spectrum of molecules, which are *qualitatively very different* from atoms in terms of their properties and the information they contain....Different combinations of proteins produce *qualitatively different* types of cell that can be combined in various ways in the brain to produce distinct brain regions that contain and process *qualitatively different* types of information. (Markram 2006, p. 153; emphasis added)

From the leveled structure of the brain, however, Markram derives an “ultimate question,” which we feel is problematic:

The ultimate question, therefore, is whether the interaction between neurons drives a series of qualitative

leaps in the manner in which information is embodied to represent an organism and its world. *As computers approach petaFLOPS speeds, it might now be possible to retrace these elementary steps in the emergence of biological intelligence using a detailed, biologically accurate model of the brain.* (2006, p. 153; emphasis added)

This, in our opinion, puts the cart before the horse. Procuring a “biologically accurate model of the brain” should represent a *challenge* that levels of organization pose, rather than an opportunity promised, to researchers (see especially our discussion of promise over substance above). When levels are taken seriously, it is irrelevant how fast our computers are (a quantitative point) when we are concerned with the *qualitative* differences that are engendered by moving from one type of (level-bound) data, such as single-cell recordings, to another type of (level-bound) data, such as gene expression. Redescribing these things using one language, i.e., computer programming languages, abstracts away from all but the most rudimentary differences (and similarities) in level-bound empirical structures. This puts pressure on the tantalizing claim that “[a] new approach is now possible that involves a quantum leap in the level of biological accuracy of brain models” (Markram 2006, p. 154).

If we have not emphasized enough, this is all perfectly fine scientific work. The question we are posing concerns rather what kinds of scientific aims are or can be legitimately served with such projects as the BBP. For even when a model is in place, the fidelity of the simulation in mimicking the real-world system itself remains a major undertaking to be explored and evaluated by multitudes of researchers from different research areas. Balancing between realism and abstraction comprises an active array of choices that neuroscientific modelers and simulators face in articulating the findings of their research (Herz et al. 2006; see also Grim et al. 2013).² Here it is clear that we cannot serve all possible research goals with one simulation. Goals like realistically capturing the details of a system or producing general conclusions require trading off between the goals that simulation and modeling may wish to achieve (Levins 1966). That is,

² One important issue that appears here concerns when abstraction (removing detail until a desired grain of description is attained) crosses the line into an idealization (actively distorting factual details of a system). Though both idealization and abstraction play positive, even necessary roles in science, there are distinct issues when one or the other is pursued. Presumably, modelers and simulators of whole brains are interested in abstraction, given their claims of producing “accurate” models of neural systems (see especially the preceding discussion). Nonetheless, we’d like to point out that such aims may pass into the realm of idealization, where a different set of issues crop up (see especially Potochnik 2017). Many thanks to an anonymous reviewer for bringing this to our attention.

realism and abstraction come at a price: realistic simulations may capture actual details of the target system, but the generality of the product's significance will thereby be hindered. Likewise, increasingly abstract simulations may be successfully generalizable, but at the cost of sacrificing realistic details in what they represent. This is simply a consequence of analyzing extremely complex systems whose constituents are either unknown or contain too many variables to be faithfully reproduced.

A further note about realism: too much detail also threatens to decrease the explanatory value of a simulation model. Take as an analogy a map of a given city: tourists visiting and navigating European cities would hardly benefit from a 1:1 recreation of Venice, as they will be inundated with irrelevant information that far exceeds the interests and aims of their recreational visit. Rather, a map or maps containing the relevant information concerning good restaurants or public transportation (and abstracting away from, for instance, personal addresses of the city's inhabitants) would be a more *relevant* description of the city given the tourists' desires in visiting Venice. Likewise, in a neuroscientific context, a description should strike a balance between realism and abstraction (Herz et al. 2006). Getting lost in the details sacrifices this relevance and abstraction to full-blown descriptive realism, which is unable to serve anyone's explanatory goals.

So, combining abstraction with the issue of organizational levels results in the following dilemma: either we wish to emphasize the contributions or findings that are related to one or several levels whose relationships are well specified, or instead we abstract too far away from these differences to find a more general pattern. Since each level of organization presumably deals with its own set of qualitative properties (cells differ fundamentally from tissues, and molecules from cells), abstracting away from these level-bound differences will smear rather than sharpen our focus of the issues arising from the specific structures whose dynamics we seek to uncover by simulation. The issue of levels is of course tightly connected to the choice of the degree of abstraction mentioned above. It's perfectly valid to model the mammalian cortex as a two-dimensional mean-field if that is appropriate for the specific question asked in the study. But a "complete" simulation of a mammalian brain should be able to address an exceptionally wide variety of reasonable questions about that brain, making the choice of the degree of abstraction much more difficult. This strong claim is made, explicitly or implicitly, by the groups developing these simulations. Claims that brain pathologies will emerge from the finished, "complete" simulation are in this category.

Without a specific hypothesis to test, any cut-off above a certain degree of abstraction will necessarily be arbitrary. Novel multi-scale simulation methods might only partially solve this problem—rather they will alleviate its symptoms.

The arbitrary cut-off will likely simply be shifted to lower (smaller and faster) levels, and not abolished. It is to be seen if a lower bound exists where the physical reality is not relevant anymore for the functioning of the nervous system, and if simulations can reach down to that level.

Frequently, the current practice of constructing neural simulations runs in the other direction, from model to question, but we consider this bad practice from an epistemological point of view. The situation is akin to someone who built a randomly connected network of integrate-and-fire neurons, to study a generic case of the propagation of neural excitation. Then, this modeler starts to use this same model to address the spread of epilepsy in one very specific brain area. This will necessarily fail, due to a lack of specificity of the model for the question at hand. The main mistake here is that the model came before the scientific question. When the claim is to build a "complete" simulation of the brain, a vast number of simulations will be like the situation of the modeler in our example. In contrast to the simple example with the network of integrate-and-fire neurons, this shortcoming is not as easy to see in a more complex simulation.

For this reason, although simulation modelers note the leveled structure of the nervous system as a target of their integrative efforts, we believe the challenge of levels to be one in which one poses the problem that one seeks to investigate rather than a promise concerning the goods to be delivered in a full-brain simulation. However, this challenge can also be reformulated as a promising concrete aim. When proper attention is paid to the work bearing on integrating, e.g., the molecular, genetic, and overarching network dynamics of nervous system activity into one simulation, this certainly would offer something more than others can: namely, a more appropriate, relevant, or complete view of the target system of interest.

Animals and Their Ecological Niches

We would like to make a final point. One is unlikely to achieve a satisfactory understanding of the mouse or cat brain without taking the ecological niches of these animals into consideration. A mouse is a nocturnal herbivore that shows an unusually fast reproductive cycle for mammals; a cat is a solitary, nocturnal visually guided ambush predator with an extended period of maternal care, and their brains are specialized for these ecological niches. One of many examples is the tapetum in the back of the eye of the cat. These reflective layers increase the photon yield of vision by reflecting photons back onto the photoreceptor cells in the retina. The gain in photon yield is achieved by a loss in visual resolution, since the reflected photons might not be reflected in a direct 180° angle. This adaptation, a trade-off between a lower light detection threshold and reduced

resolution, is useful for a nocturnal predator. The retina and all subsequent structures in the cat visual system have to cope with this trade-off. Already at the very first stage of the visual system, the feline visual system is significantly different from the human visual system (Oliver et al. 2004).

Neither the cat's nor other animals' brains are scaled-down models of human brains, as which they are often treated. Rather, they are brains of specialized vertebrates that evolutionarily split tens of millions of years ago from a common, shrew-like ancestor. As the late Ted Bullock noted, "Neuroscience is part of biology, more specifically of zoology, and it suffers tunnel vision unless continuous with ethology, ecology, and evolution" (Bullock 1984).

The relationship between an animal's ecological niche and a complete simulation of its brain is twofold. Given the specialized nature of nervous systems, it is of limited use to attempt to simulate "the mammalian brain" or "the cortex," since these structures will widely differ between species. Furthermore, the inputs given to a whole brain simulation will have to be tailored to the percepts experienced by that specific animal in its natural environment; otherwise the dynamics of the nervous system will be incompletely explored. As an example, a simulation of the brain—including the visual system—of the aforementioned cat would not be meaningful if not tested in conditions akin to hunting at low light. Both the ecological relevance of the simulation would be limited without such a test, and, from a dynamical systems point of view, the parameter regime of the simulation would be incompletely explored. Basically, we believe that the environmental niche determines the statistics of the sensory inputs an animal receives in a very fundamental way. This will have influenced the evolution of this nervous system, and the statistics of the neural connections will be determined by these input statistics as determined by the niche. Any successful, realistic whole-brain simulation will have to incorporate such connection statistics. Furthermore, testing of the model only makes sense with inputs akin to those faced by the animal in its natural environment. Otherwise, the response will be somewhat of an *in silico* lab artifact.

Summary

We have argued that efforts to "simulate the brains of mammals" have so far not succeeded not because of a lack of effort or computer capacities, but due to several fundamental limitations. These are the relationship between simulations and hypothesis testing (and a lack of hypotheses), the complications arising from a lack of clear parameter constraints, the multilevel nature of neural computation, and the specific ecological specializations of animals and their brains.

We are under the impression that the multilevel nature of neural computation and the issues regarding parameter constraints are taken seriously by the community engaged in large-scale brain simulations. Issues relating to multi-scale simulations are not yet solved, and there is no consensus if they will eventually prove to be prohibitive for successful whole-brain simulations. A significant amount of work is dedicated to them (Breakspear and Stam 2005; Robinson et al. 2005; Honey et al. 2007; Torben-Nielsen and Stiefel 2009). However, the issues arising from a lack of a clear hypothesis are less well addressed in large-scale full-brain simulations. The issues relating to the very specific ecological specializations of different animals seem to be largely ignored in the course of whole-brain simulation efforts, as indicated by multiple models of generic mammalian brains or generic mammalian cortices.

Acknowledgments We thank the reviewers of the initial version of this manuscript, our colleagues at the Konrad Lorenz Institute, Drs. Greg Cohen, Jonathan Tapson, and the participants of the 2014 Teluride Neuromorphic Cognition Engineering Workshop for helpful discussion.

References

- Almog M, Korngreen A (2016) Is realistic neuronal modeling realistic? *J Neurophysiol*. <https://doi.org/10.1152/jn.00360.2016>
- Ananthanarayanan R, Esser SK, Simon HD, Modha DS (2009) The cat is out of the bag: cortical simulations with 109 neurons, 1013 synapses. In: Proceedings of the conference on high performance computing networking, storage and analysis, SC09, Portland, 14–20 November 2009, pp 1–12
- Blau A, Callaly F, Cawley S, Coffey A, De Mauro A et al (2014) The *Si elegans* project—the challenges and prospects of emulating *Caenorhabditis elegans*. In: Duff A, Lepora NF, Mura A, Prescott TJ, Verschure PFMJ (eds) Biomimetic and biohybrid systems. Living machines 2014, Lecture Notes in Computer Science, vol 8608. Springer, Cham, pp 436–438
- Borst A, Egelhaaf E (1989) Principles of visual motion detection. *Trends Neurosci* 12(8):297–306
- Breakspear M, Stam CJ (2005) Dynamics of a neural system with a multiscale architecture. *Philos Trans R Soc B* 360:1051
- Brooks DS (2014) The role of models in the process of epistemic integration: the case of the Reichardt motion detector. *Hist Philos Life Sci* 36(1):90–113
- Bullock TH (1984) Comparative neuroscience holds promise for quiet revolutions. *Science* 225(4661):473–478
- Campbell DT (1974) Downward causation in hierarchically organised biological systems. In: Ayala FJ, Dobzhansky T (eds) Studies in the philosophy of biology: reduction and related problems. MacMillan, London, pp 179–186
- Ching J, Beck JL, Porter KA (2006) Bayesian state and parameter estimation of uncertain dynamical systems. *Probab Eng Mech* 21:81–96
- Craver C (2007) Explaining the brain. Clarendon Press, Oxford
- De Ruyter van Steveninck RR, Zaagman WH, Mastebroek HAK (1986) Adaptation of transient responses of a movement-sensitive neuron in the visual system of the blowfly *Calliphora erythrocephala*. *Biol Cybern* 54:223–236

- Egelhaaf M, Borst A (1993) A look into the cockpit of the fly: visual orientation, algorithms, and identified neurons. *J Neurosci* 13(11):4563–4574
- Eliasmith C, Stewart TC, Choo X, Bekolay T, DeWolf T et al (2012) A large-scale model of the functioning brain. *Science* 338(6111):1202–1205
- Eronen MI, Brooks DS (2018) Levels of organization in biology. In: Zalta EN (ed) *The Stanford encyclopedia of philosophy*. Stanford University, Stanford. <https://plato.stanford.edu/archives/spr2018/entries/levels-org-biology/>
- Faumont S, Rondeau G, Thiele TR, Lawton KJ, McCormick KE et al (2011) An image-free opto-mechanical system for creating virtual environments and imaging neuronal activity in freely moving *Caenorhabditis elegans*. *PLoS ONE* 6:e24666
- Gerstein MB, Lu ZJ, Van Nostrand EL, Cheng C, Arshinoff BI et al (2010) Integrative analysis of the *Caenorhabditis elegans* genome by the modENCODE project. *Science* 330:1775–1787
- Grim P, Rosenberger R, Rosenfeld A, Anderson B, Eason RE (2013) How simulations fail. *Synthese* 190(12):2367–2390
- Guttman BS (1976) Commentary: is “levels of organization” a useful biological concept? *BioScience* 26:112–113
- Hamood AW, Haddad SA, Otopalik AG, Rosenbaum P, Marder E (2015) Quantitative reevaluation of the effects of short- and long-term removal of descending modulatory inputs on the pyloric rhythm of the crab, *Cancer borealis*. *eNeuro*. <https://doi.org/10.1523/ENEURO.0058-14.2015>
- Hartmann S (1996) The world as a process. Simulations in the natural and social sciences. In: Hegselmann R, Mueller U, Troitzsch KG (eds) *Modelling and simulation in the social sciences from the philosophy of science point of view*. Springer, Dordrecht, pp 77–100
- Herz AVM, Gollisch T, Machens CK, Jaeger D (2006) Modeling single-neuron dynamics and computations: a balance of detail and abstraction. *Science* 314:80–85
- Hoel EP, Albantakis L, Marshall W, Tononi G (2016) Can the macro beat the micro? Integrated information across spatiotemporal scales. *Neurosci Conscious* 1:niw012. <https://doi.org/10.1093/nc/niw012>
- Honey CJ, Kötter R, Breakspear M, Sporns O (2007) Network structure of cerebral cortex shapes functional connectivity on multiple time scales. *PNAS* 104:10240–10245
- Hubel DH, Wiesel TN (1968) Receptive fields and functional architecture of monkey striate cortex. *J Physiol* 195:215–243
- Izhikevich EM (2010) *Dynamical systems in neuroscience: the geometry of excitability and bursting*. MIT Press, Cambridge
- Izhikevich EM, Edelman GM (2008) Large-scale model of mammalian thalamocortical systems. *PNAS* 105:3593–3598
- Krohs U (2008) How digital computer simulations explain real-world processes. *Int Stud Philos Sci* 22(3):277–292
- Levins R (1966) The strategy of model building in population biology. *Am Sci* 54(4):421–431
- Machado P, Wade J, McGinnity TM (2014) Si elegans: FPGA hardware emulation of *C. elegans* nematode nervous system. In: 2014 sixth world congress on nature and biologically inspired computing (NaBIC 2014), pp 65–71
- Marder E, Haddad SA, Goeritz ML, Rosenbaum P, Kispersky T (2015) How can motor systems retain performance over a wide temperature range? Lessons from the crustacean stomatogastric nervous system. *J Comp Physiol A* 201:851–856
- Markram H (2006) The blue brain project. *Nat Rev Neurosci* 7:153–160
- Markram H, Muller E, Ramaswamy S, Reimann MW, Abdellah M et al (2015) Reconstruction and simulation of neocortical microcircuitry. *Cell* 163:456–492
- Marr D (1982) *Vision: a computational approach*. Freeman, San Francisco
- Meister P, Towbin BD, Pike BL, Ponti A, Gasser SM (2010) The spatial dynamics of tissue-specific promoters during *C. elegans* development. *Genes Dev* 24:766–782
- Merolla PA, Arthur JV, Alvarez-Icaza R, Cassidy AS, Sawada J et al (2014) A million spiking-neuron integrated circuit with a scalable communication network and interface. *Science* 345:668–673
- Ollivier FJ, Samuelson DA, Brooks DE, Lewis PA, Kallberg ME, Komáromy AM (2004) Comparative morphology of the tapetum lucidum (among selected species). *Vet Ophthalmol* 7:11–22
- Palyanov A, Khayrulin S, Larson SD, Dibert A (2012) Towards a virtual *C. elegans*: a framework for simulation and visualization of the neuromuscular system in a 3D physical environment. *In Silico Biol* 11:137–147
- Parker WS (2009) Does matter really matter? Computer simulations, experiments, and materiality. *Synthese* 169:483–496
- Potochnik A (2017) *Idealization and the aims of science*. University of Chicago Press, Chicago
- Robinson PA, Rennie CJ, Rowe DL, O’Connor SC, Gordon E (2005) Multiscale brain modelling. *Philos Trans R Soc Lond B* 360:1043–1050
- Schrödel T, Prevedel R, Aumayr K, Zimmer M, Vaziri A (2013) Brain-wide 3D imaging of neuronal activity in *Caenorhabditis elegans* with sculpted light. *Nat Methods* 10:1013–1020
- Schulz DJ, Goillard JM, Marder E (2006) Variable channel expression in identified single and electrically coupled neurons in different animals. *Nat Neurosci* 9(3):356
- Sejnowski TJ, Churchland PS (1994) *The computational brain*. A Bradford Book, MIT Press, Cambridge
- Selverston A (2008) Stomatogastric ganglion. *Scholarpedia* 3:1661
- Stephan A (1999) Varieties of emergence. *Evol Cogn* 5(1):49–59
- Szigeti B, Gleeson P, Vella M, Khayrulin S, Palyanov A et al (2014) OpenWorm: an open-science approach to modeling *Caenorhabditis elegans*. *Front Comput Neurosci*. <https://doi.org/10.3389/fncom.2014.00137>
- Toni T, Welch D, Strelkowa N, Ipsen A, Stumpf MPH (2009) Approximate Bayesian computation scheme for parameter inference and model selection in dynamical systems. *J R Soc Interface* 6:187–202
- Torben-Nielsen B, Stiefel KM (2009) Multiscale modeling of cortical neural networks. *AIP Conf Proc* 1167(1):15–25
- van der Merwe R, Wan EA (2001) The square-root unscented Kalman filter for state and parameter-estimation. In: 2001 IEEE international conference on acoustics, speech, and signal processing. Proceedings (Cat No. 01CH37221), vol 6, pp 3461–3464
- Van Essen DC, Maunsell JH (1983) Hierarchical organization and functional streams in the visual cortex. *Trends Neurosci* 6:370–375
- Varier S, Kaiser M (2011) Neural development features: spatio-temporal development of the *Caenorhabditis elegans* neuronal network. *PLoS Comput Biol* 7:e1001044
- Wimsatt W (2007) *Re-engineering philosophy for limited beings*. Harvard University Press, Cambridge
- Winsberg E (1999) Sanctioning models: the epistemology of simulation. *Sci Context* 12(2):275–292
- Winsberg E (2001) Simulations, models, and theories: complex physical systems and their representations. *Philos Sci* 68:442–454
- Winsberg E (2003) Simulated experiments: methodology for a virtual world. *Philos Sci* 70:105–125