

The Sum of the Parts: Large-Scale Modeling in Systems Biology

Fridolin Gross* and Sara Green†

Systems biologists often distance themselves from reductionist approaches and formulate their aim as understanding living systems “as a whole.” Yet, it is often unclear what kind of reductionism they have in mind, and in what sense their methodologies would offer a superior approach. To address these questions, we distinguish between two types of reductionism which we call “modular reductionism” and “bottom-up reductionism.” Much knowledge in molecular biology has been gained by decomposing living systems into functional modules or through detailed studies of molecular processes. We ask whether systems biology provides novel ways to *recompose* these findings in the context of the system as a whole via computational simulations. As an example of computational integration of modules, we analyze the first whole-cell model of the bacterium *M. genitalium*. Secondly, we examine the attempt to recompose processes across different spatial scales via multi-scale cardiac models. Although these models rely on a number of idealizations and simplifying assumptions as well, we argue that they provide insight into the limitations of reductionist approaches. Whole-cell models can be used to discover properties arising at the interfaces of dynamically coupled processes within a biological system, thereby making more apparent what is lost through decomposition. Similarly, multi-scale modeling highlights the relevance of macroscale parameters and models and challenges the view that living systems can be understood “bottom-up.” Specifically, we point out that system-level properties constrain lower-scale processes. Thus, large-scale modeling reveals how living systems at the same time are *more* and *less* than the sum of the parts.

Keywords

reductionism • simulations • whole-cell modeling • multi-scale cardiac modeling
heuristics • systems biology

Part of a special issue, Ontologies of Living Beings, guest-edited by A. M. Ferner and Thomas Pradeu

*Institut für Philosophie, Universität Kassel, Nora-Platiel-Strasse 1, 34127 Kassel, Germany,
fridolin.gross@uni-kassel.de

†Department of Science Education, Øster Voldgade 3, 1350 København, Denmark,
sara.green@ind.ku.dk

Received 30 September 2016; Accepted 28 February 2017
doi:10.3998/ptb.6959004.0009.010



Editorial introduction: This contribution from Fridolin Gross and Sara Green focuses on the promises and possible pitfalls of large-scale modelling in systems biology, from both a practical and a theoretical point of view. As readers will be aware, molecular biology has often been accused of being “reductionist.” Systems biology has been presented as a potential response to this reductionism. Gross and Green distinguish two types of reductionisms: “modular reductionism” (the claim that living systems consist of individual functional sub-units) and “bottom-up reductionism” (the claim that biological phenomena can and should be studied at the molecular level, understood as the “fundamental” biological level). They show how systems biology questions both forms of reductionism. To this end, they use two examples: whole-cell modelling of a very “simple” bacterium (from a piece published in 2012), and multi-scale computer modelling of the human heart. Without masking the limitations of these models, Gross and Green show the insights they can yield for anyone interested in ontologies of the living world.

As with some of the other contributions in this special issue (e.g. Fagan’s, and Kendig and Eckdahl’s), a virtue of this paper is to illustrate how research strategies—particularly modelling strategies—can influence our ontological categorizations of living things. More specifically, this paper raises the following interesting point: When a scientific approach fails, what does it tell us (if anything) about the world? Reality is often understood as what *resists* our representations. If so, then the examples given by Gross and Green could be particularly revealing. They suggest that purely bottom-up approaches in biology often fail, because they neglect higher-level or “systemic” factors, and must therefore be complemented by systemic approaches. (This is related to the challenge of the “tyranny of scales,” analysed by Robert Batterman and others, which rests on the idea that no single mathematical model can capture behaviours within all spatial scales). Does it suggest that we should believe in a form of ontological emergence and/or in downward causation? Gross and Green also insist—following physiologist Denis Noble, and others—that decades of cardiac modelling have shown that the failures were probably more important than the successes, since they reveal unanticipated links between components previously thought to be independent. Does that tell us something about the existence of “intrinsic” constraints and modularity in living things?

Ultimately, the main question raised by Gross and Green is very similar to that of more metaphysically-inclined contributors to this special issue, such as Rory Madden, Christopher Shields, and Stéphane Chauvier: To what extent is the whole the sum of its parts, and how should we define “wholes” and “parts” in the first place? Gross and Green show that, in biological systems, the whole is both more and less than the sum of the parts: it is *more*, insofar as system-level properties cannot be predicted from properties of components alone; at the same time, it is *less*, insofar as systemic constraints prevent the components from exhibiting some of the properties and degrees of freedom that they can possess in other contexts. There is no doubt systems biology will continue to play a major role at the interface of epistemological and ontological considerations in the near future. —AF/TP

1 Introduction

Over the second half of the last century, molecular biologists have developed ingenious strategies to unravel the mechanisms underlying cellular and organismal behavior by intervening on individual types of molecules or by modifying the structure of DNA in a cell. The approach of molecular biology has often been called “reductionistic” because it attempts to understand phe-

nomena in terms of entities and processes at a lower level, or as subsystems in isolation from the system as a whole. For some people the undeniable success of molecular biology is evidence that living systems are organized in a particular way that lends itself to such reductionistic approaches. There is a risk, however, that our simplifying strategies will underestimate or hide the actual complexity of a system, and that our idea of biological organization is partly an artifact of our methods. Systems biologists have recently pointed to the limitations of reductionist approaches pursued in molecular biology (Kitano 2002; Green and Wolkenhauer 2013). In this paper we examine more closely what is meant by reductionism in this context and what alternatives systems biology may have to offer.

We approach the question of what it means to “study a system as a whole” by examining heuristic strategies in systems biology in comparison to those of molecular biology. By heuristics we understand rules of thumb that are employed to reduce the complexity of a scientific task by making specific assumptions about the system under study (Simon 1962; Wimsatt 2007). Important heuristics in molecular and cell biology are what Bechtel and Richardson ([1993] 2010) identified as the twin strategies of *decomposition* and *localization*. Decomposition starts from the idea that the activity of a system is a product of component functions. Localization then consists in mapping those component functions on structural parts of the system. Drawing on historical case studies, Bechtel and Richardson show that biologists often decompose biological phenomena into component parts and operations and localize these as parts of a mechanism. These heuristics thus rely on a guiding assumption about functional decomposition of biological systems, akin to how “a machine is a composite of interrelated parts, each performing its own functions, that are combined in such a way that each contributes to producing a behavior of the system” (Bechtel and Richardson [1993] 2010, 17). Acknowledging that organisms in many ways are unlike modular machines, Bechtel and Richardson emphasize the biases inherent in such heuristics but also highlight their productivity in the context of the cognitive and technical constraints that research must operate within.

Some of these constraints change over time, however. New methodological strategies can give rise to new ways of triangulating evidence from different sources, enabling researchers to revisit the underlying assumptions of their research practices (Bechtel and Richardson [1993] 2010, xxxvii, see also Wimsatt 2007). In the introduction to the 2010 edition to *Discovering Complexity*, Bechtel and Richardson emphasize the need for strategies of *recomposition* that have become possible with the availability of large datasets and computational models. Recomposition refers to the investigation of whether the postulated component operations, given appropriate information about organizational features and environmental conditions, will yield the systemic behavior. In this context they highlight the increasing application of network models and computational simulations that are central methods in systems biology. Does this expansion of research strategies via computational modeling offer a less reductionistic way to understanding biological complexity? In this paper we aim to throw light on this question through an examination of some of the most advanced models developed in systems biology. We focus on large-scale simulations of a whole cell and the human heart, respectively. We investigate whether and to what extent these modeling projects manage to overcome the limitations of the traditional approach, and whether such approaches have ontological implications for our view of biological organization and complexity.

To clarify what we understand by reductionism in the context of systems biology, we distinguish in Section 2 between “modular reductionism” and “bottom-up reductionism” and examine two types of corresponding recompositional strategies. Section 3 analyzes the first whole-cell model as an example of an effort to recombine all processes within a cell that before had only

been studied as independent modules. Section 4 examines an attempt to recompose processes across different spatial and temporal scales, focusing on multi-scale cardiac modeling. We shall argue that although neither of these strategies offers an unbiased and comprehensive methodology, large-scale modeling provides novel insight into biological complexity. Specifically, it can teach us about the dynamic interfaces that are lost in modular decomposition and reveal the limitations of bottom-up reductionism (Section 5). Section 6 offers concluding remarks.

2 Reductionism and Its Limits in Molecular Biology

Systems biologists often define their approach in contrast to reductionist methodologies (e.g. Kitano 2002). To what extent molecular biology is a purely reductionistic endeavor is debatable, however. Rather than a discipline, it is more of a “technique-based field” (Burian 1993) that contributes to a variety of biological investigations and is often complemented with methods of cell biology or physiology that take into account higher level structures and processes. Yet it is undeniable that the advent of molecular biology has resulted in a tendency to privilege the molecular level and to look for well-delineated mechanisms with a small number of component parts. In what follows we will make this notion of reductionism more precise.

Biological systems are undeniably complex. Organisms, such as ourselves, consist of an incredibly large number of parts—organs, cells, molecules—that are organized and interact in intricate ways. In order to explain living phenomena, biologists have to find ways to cope with this complexity. The strategies adopted are often influenced by the ideas they have about the structure and organization of living systems. With attempts to grasp biological complexity, questions arise as to what extent complex “higher-level” phenomena can be understood through the properties of underlying components.

Biologists nowadays agree that in order to explain the phenomena of life one does not need to invoke forces or substances that are different from those known to physics and chemistry. Thus, a weak form of ontological reductionism, often called “physicalism,” according to which living things consist of the same kind of stuff as all other objects in the material world, is basically uncontroversial (Brigandt and Love 2017). There is disagreement, however, about stronger forms of reductionism that make more specific claims about biological organization or explanation. What seems to be at issue when systems biologists criticize the approach of molecular biology are varieties of part-whole reductionism (Kaiser 2015). Part-whole reductionism assumes that there is a “higher level” of the system as a whole and a “lower” level of the parts of the system, and it is often summarized by saying that the system is “nothing but the sum of its parts.” However, this can be spelled out in different ways, and we think that it is productive to distinguish between two varieties of part-whole reductionism that are often tacitly equated in discussions around systems biology.

One way of understanding part-whole-reductionism is the idea that the behavior of the whole system is decomposable into the activities of its parts. It is closely related to the heuristics of decomposition and localization mentioned in the beginning. A fully decomposable system is one that consists of subsystems, such that the interactions between the subsystems are negligible with respect to the interactions within each of the subsystems. That is, one can study each of the subsystems as if it was isolated from the others (Simon, 1962). Since decomposability presupposes a modular organization of the system, we refer to this idea as *modular reductionism*. Modularity is a rich concept, and different kinds of modules have been discussed in different fields of biology (Callebaut and Rasskin-Gutman, 2005). Winther (2001) distinguishes between structural, physiological, and developmental modules. Modular reductionism,

as we understand it, concerns modularity both in a structural and a physiological sense since it assumes that the behavior of a system is produced by the individual activities of structurally distinct parts.

Modular reductionism thus concerns the relationship *among the parts (or modules)* of a system. It can be distinguished from *bottom-up reductionism* which concerns the relationship *between spatial scales or levels* and can be understood as the claim that higher-level causes are nothing but aggregates of lower-level causes.¹ Strong forms of bottom-up reductionism exclude the possibility that higher-level causes can somehow act on or constrain entities or processes at the lower-level, which is sometimes referred to as “downward causation” (see e.g. Kim 1999). If top-down effects can be neglected in the development of biological models and explanation, we should be able to understand the system “bottom-up” from detailed lower-level descriptions.

Modular and bottom-up reductionism both have ontological, explanatory and methodological versions that are logically independent, although the relations between them have at times been controversially discussed in the philosophical literature. For instance, advocating methodological reductionism does not entail a commitment to the corresponding forms of explanatory or ontological reductionism but can be argued for on purely pragmatic grounds. Bechtel and Richardson ([1993] 2010) have shown that the reductionistic strategies of structural decomposition and functional localization can lead to sophisticated mechanistic explanations. These can take into account non-trivial forms of biological organization even if underlying assumptions, such as modularity, are not justified. However, it is debated whether methodological reduction can lead to oversimplified ontological accounts of living systems (cf. Green 2015; Nicholson 2013). In spite of their success in 20th Century biology, it has been argued that reductionist strategies are inherently limited and lead to partial or incomplete accounts of living phenomena, either because some organizational features of biological systems are irreducible to the lower level (e.g. Mitchell 2003), or because they neglect the role of the context in which a mechanism is embedded (e.g. Gilbert and Sarkar 2000).

Developments over the last decades have increased skepticism regarding reductionistic strategies, but at the same time created the hope that their limitations can be overcome. Massive amounts of data have been accumulated following advances in experimental techniques and provide a glimpse of the actual complexity of biological systems. Moreover, biological researchers today have access to a variety of analytical and computational tools that can better account for the dynamics of complex mechanisms.

Systems biology has been proposed as an approach that provides alternative and superior strategies to deal with biological complexity. For instance, systems biology offers new ways of investigating biological modularity through studies of genes and proteins that are associated in large organized networks (e.g. Hartwell 1999; Ravasz et al. 2002).² In this paper, we focus on the strategy to develop large-scale simulations for computational integration of various orchestrated processes. In what follows we look at two examples to ask whether large-scale modeling can provide alternatives to the two forms of reductionism introduced above: a whole-cell model of the bacterium *M. genitalium*, and multi-scale modeling of the human heart.

¹We shall refer to levels when talking about functional organization into cells, tissues, organs etc. but refer to scales when parsing systems into organizational levels is less clear or when talking about multi-modeling strategies used in both physics and biology.

²For a philosophical discussion of the implications of network modeling for investigating modularity and hierarchical organization in the context of network modeling in systems biology, see Green et al. (2016). For a discussion of modularity in the context of synthetic biology, see Kendig and Eckdahl (this issue).

One important driving force behind the whole-cell model is the aim to go beyond *partial* mechanistic models and to take into account how different mechanisms within a cell are integrated in their cellular context. Can whole-cell models recompose insights about subsystems studied in isolation and thus counteract the biases of modular reductionism? In a similar manner, we investigate whether multi-scale cardiac modeling can overcome the limitations of bottom-up reductionism. In more general terms, we ask whether large-scale simulations provide a way to obtain a more adequate representation of a biological system as a whole.

3 The Whole-Cell Model

As defined above, “modular reductionism” is based on the idea that an organism or a cell is composed of a number of processes or mechanisms that can be studied and understood in isolation. Traditional molecular biology with its focus on individual mechanisms is clearly based on the assumption that biological systems are near-decomposable. However, recent studies of biological networks (e.g. gene regulatory, neural and metabolic networks) show a high level of interconnectivity, questioning to what extent such interactions and the systemic context can be neglected. Thus, the necessity of recomposing insights about specific mechanisms into a whole to study the effects of interconnected pathways and processes is becoming ever more pressing (Bechtel and Abrahamsen 2010). Even if a decomposition strategy successfully illuminates biological phenomena based on isolated representations of parts of a system, recomposition is important to ensure that the assumptions underlying modular reductionism are in fact justified. Furthermore, it might be a fruitful strategy to reveal errors in accounts of component mechanisms and to detect biases introduced by reductionist strategies. Recomposition, however, is not a trivial task. It typically requires that the component mechanisms of the system are fairly well understood, and that strategies are available for integrating data and models into a consistent and intelligible representation of the whole system. A natural aim for recomposition strategies in systems biology is to integrate all cellular processes that are separately well understood and to build a model of a whole cell.

3.1 *The Idea of a Whole-Cell Model*

The dream of a complete whole-cell model is not new (Crick 1973), but it is only recently that the required resources, such as biological information and computational power, have become available to turn it into a serious research project. This idea has been central for many proponents of systems biology from the outset. Around the turn of the millennium the Japanese biologist Masaru Tomita considered whole-cell models as one of the “grand challenges of the 21st century” (Tomita 2001). He argues that in order to understand cellular behavior and to answer the most pressing questions of molecular biology and medicine, one needs a complete simulation of a living cell. Around the same time, his colleague Hiroaki Kitano advocated that such a project should be funded at the same scale as projects like the Human Genome Project. He called this the “Systeome” project:

The goal of the human systeome project, if it is realized at all, shall be defined as to complete a detailed and comprehensive simulation model of human cells at an estimated error margin of 20% by the year 2020 and to finish the identification of the system profile for all genetic variations, drug responses, and environmental stimuli by the year 2030. (Kitano 2002, 9)

However, such bold claims have also been met with skepticism, and the purpose of such projects has been put into question. Philosophers and scientists have typically accepted a tradeoff between the virtues of a plurality of different models because all aspects of a complex system cannot be modeled and analyzed simultaneously (Levins 1966). Ulrich Krohs and Werner Callebaut lump Kitano's Systeme proposal together with other "omic" or data-mining projects that collect data without providing strategies to turn the information into explanatory models. Specifically, they write:

We [criticize] the project of a "realistic" representation of all metabolic processes in a 1:1 manner as lacking explanatory power and, more generally, as being epistemologically misguided. It regresses to the "omic" approach by once again offering "complete" but physiologically uninterpreted data sets. (Krohs and Callebaut 2007, 209)

They argue that models have to be simplified or idealized in order to be informative. If one simply reproduces everything that happens inside the cell on a computer, then the model will be as intractable as the system itself. Such projects have been criticized also by scientists, including systems biologists themselves. They argue that it is doubtful whether one could incorporate, in practice or even in principle, the "astronomical" number of individual interactions, given the computational demands this would entail (Noble 2012). To get a rough idea, Bassingthwaite et al. (2009) estimate that if determined via quantum mechanical calculations, the process of protein folding inside a cell would by itself require months of computation on the fastest parallel computers currently available.

In order to better assess the potential merits and limitations of whole-cell modeling, we believe it is useful to look at concrete examples of projects pursuing the goal of completeness. Recently, in July 2012, the first complete model of a living cell was published in the scientific journal *Cell* (Karr et al. 2012). Taking into account the current state of the art of whole-cell-modeling might allow for a more grounded perspective on the potential role of such models in the life sciences. Based on our discussion we want to argue that such models can play a very specific epistemic role, even if they may not solve all the riddles of life. Although being the first of its kind, the example points to several general problems that every such modeling project will have to face, and the conclusions we wish to draw in the end have implications beyond the particular case at hand.

3.2 *A Model of Mycoplasma genitalium*

The first whole-cell model (WCM) describes the life-cycle of the pathogen *Mycoplasma genitalium* and includes all known molecular components and interactions (Karr et al. 2012). *M. genitalium* is a small parasitic bacterium residing in the urogenital and respiratory tracts of primates. The choice of this organism is not arbitrary: Its genome, which consists of only 525 genes, is the smallest of any living organism found in nature. Already in the late 1950s, a group around the biophysicist Harold Morowitz considered Mycoplasma as the paradigm organism to study the logic of life in a minimal system. As Morowitz recalls:

Just as the hydrogen atom, the smallest and simplest member of the periodic table, had served to sharpen many of the fundamental questions of spectroscopy and quantum mechanics, so, we reasoned, would a minimum biological system play an analogous role. (Morowitz 1984, 750)

The geneticist Craig Venter, several decades later, followed Morowitz's path using *Mycoplasma* bacteria as model organisms for his "minimal genome project" (Glass et al. 2006) and for his experiments of synthesizing and transplanting entire genomes (Gibson et al. 2010). Systematic efforts by a consortium of European research groups have recently accumulated vast amounts of system-wide data about the molecular constituents of *Mycoplasma* (Ochman and Raghavan 2009). The information necessary for attempting to build a complete model was therefore available.

Differently from what some might perhaps expect from a "complete" model, the researchers did not simply lump all the molecular components together to create one big system of equations. There is no single modeling framework that can encompass the variety of biological processes from transcription to metabolism to protein folding. Moreover, the data from experimental measurement of the different processes have different degrees of specificity and uncertainty stemming from the different experimental techniques used. The researchers had to develop strategies for integrating in one system the outputs of Boolean, probabilistic, and constraint-based submodels, among others. Thus, for modeling the diversity of biological processes there must be a division of labor among different model types, and a key challenge is to integrate these while accounting for interdependencies between the processes.

Karr et al. constructed the model by first putting together sub-assemblies, each comprising a substantially smaller number of parts in comparison with the whole system. They functionally decomposed the system into 28 modules, corresponding to different cellular functions (right column of Figure 1). The submodels correspond to processes that describe six areas of cell biology: transport and metabolism, replication and maintenance of DNA, synthesis and maturation of RNA molecules, synthesis and maturation of proteins, cytokinesis (the physical division of the cytoplasm at the end of the cell cycle), and interaction with the host organism. Essentially, each process can be represented as a set of chemical reactions that convert chemical substrates (inputs) into products (outputs) using enzymatic catalysts.

The modular approach allowed the researchers to construct the model by choosing for each module the most adequate style of mathematical representation and to build, parameterize, and test the modules independently. But despite being involved in different functions, most of the processes are interconnected by having common metabolites as inputs or outputs. For example, the *replication* process, which describes the duplication of the genome, requires nucleotides (corresponding to the letters A, T, C, G of the genetic code) that are provided by the *metabolism* process. Such interactions leading to interdependency between modules are represented by updating shared state variables (cellular variables in the left column on Figure 1). These state variables hold the information about the different kinds of entities inside the cell and their configurations.

In the process of model integration, Karr et al. drew on the assumption that the processes by which different functional modules interact can be described on a longer time scale than the processes occurring within each module. Thus, the modeling strategy is based on a temporal decomposition of functional modules but only for short time-scales:

We began with the assumption that the submodels are approximately independent on short timescales (less than 1 s). Simulations are then performed by running through a loop in which the submodels are run independently at each time step but depend on the values of variables determined by the other submodels at the previous time step. (Karr et al. 2012, 390)

In the supplementary material to the article, the authors compare their method to the numerical algorithms that are used to solve systems of ordinary differential equations (ODEs). The




Figure 1: Illustration of the principles behind the whole-cell model of *M. genitalium*. The left column lists the 16 types of cell variables according to 28 submodules representing specific processes assumed to be independent on short timescales. The number of genes associated with the submodels is given in parentheses. Reprinted by kind permission of Elsevier, from Karr et al. 2012.

28 cellular processes can be considered as “meta-equations” that are solved independently for each time step, while the 16 state variables figure in different processes, like variables in a set of equations, and therefore represent interfaces between these processes. The state variables are not simply real numbers as in the case of “ordinary” ODEs. The *Chromosome* state, for instance, “represents the polymerization, winding, modification, and protein occupancy of each nucleotide of each strand of each copy of the *M. genitalium* chromosome, and the (de)catenation status of the two sister chromosomes following replication” (Karr et al. 2012, S10). Mathematically speaking, this object is a set of 12 tensors (multi-dimensional arrays of numbers), each of which stores specific information about every nucleotide of the *M. genitalium* genome. Most of the other states, such as the RNA, Metabolite, or Polypeptide states, are of similar complexity.

3.3 *The Role and Limits of Whole-Cell Modeling*

Karr et al. consulted over 900 primary sources, reviews, and databases in order to gather as much information as possible about their model system. More than 1900 observed parameters were incorporated to specify the organization of the *M. genitalium* chromosome, the structure and function of each gene product, metabolite, and their interactions and reactions. The sheer amount of detail should, however, not lead to the impression that the model is complete with respect to molecular detail. “Completeness” in the case of the whole-cell model should not be understood as fulfilling the reductionist ideal of deriving cellular behavior from the laws of physics and chemistry. It should rather be understood as “functional completeness,” that is, as the attempt to take into account and integrate all known biological activities of the cell in one model.

Importantly, the way in which the individual modules are represented is not necessarily more advanced than the partial models of other systems biologists, and the model recapitulates biological processes only to the extent that they are currently understood. Many processes in the whole-cell model are simply “black-boxed” or represented in very coarse-grained ways. To illustrate, the *Protein Folding* process represents the three-dimensional configuration of each protein as a two-state Boolean variable: “folded” or “unfolded.” The folding rate is a probabilistic function that increments the copy number of folded proteins depending on the amount of unfolded protein, of metabolites, and of chaperones that assist the folding. Thus the physical process of folding that ultimately determines the function of a protein is not simulated at all. A virtual protein’s function does not derive from its virtual molecular configuration but from a set of rules that are explicitly formulated in the model. Processes that are better understood, such as chromosome replication, are modeled in considerable detail. Every single process is implemented according to the best available modeling strategy, but all of them heavily rely on idealizations, and many gaps remain in the overall picture.

While this might come as a disappointment to those who dream of virtual organisms as perfect *in silico* replicas, a functionally complete model may still have the potential to overcome some of the reductionist biases of decomposition and localization. Models in molecular biology typically depict a mechanism as an individual module and treat the rest of the organism as at most providing an input for, and receiving an output from it. In this way, the epistemic task is considerably simplified because any complexity that results from the communication between the modules of the system is ignored. In Karr et al.’s model, by contrast, the inter-module communication is explicitly taken into account. It allows biologists to see what happens when they connect all the pieces in the way they currently understand them, and thereby provides a consistency check of their idea of cellular organization. This is similar to what already Morowitz had in mind:

At 600 steps [his estimation of the number of genes in the *Mycoplasma* genome], a computer model is feasible, and every experiment that can be carried out in the laboratory can also be carried out on the computer. The extent to which these match, measures the completeness of the paradigm of molecular biology. (Morowitz 1984, 752)

Another way to put this is to say that an integrated representation of a whole organism imposes additional constraints on the included models of individual processes. The synthesis of enzymes, for instance, consists of several steps each of which involves a number of chemical reactions that require the presence of particular metabolites. These metabolites themselves have to be produced by other processes that require the presence of particular enzymes. The organism as a whole can sustain itself only if all of the different processes occur in a coordinated fashion such that the output of each process matches the demand of those processes that depend on it. It is precisely this coordination and closure of processes for which a whole-cell model can account.

With regard to integrated behavior, the model makes some interesting predictions that the authors refer to as model-driven discoveries. They noticed, for instance, that the overall length of the cell cycle in the simulation showed considerably less variability than the single stages of the cycle alone. Thus cell cycle length appears to be regulated in some way, even though no regulation has explicitly been incorporated in the model. By analyzing the output of their simulations, Karr et al. found that the availability of single DNA nucleotides seemed to be responsible for the phenomenon. They observed that the lengths of two stages of the cell cycle, replication initiation and replication, are inversely related to each other. If replication initiation is slow, a large pool of nucleotides builds up which in turn speeds up the subsequent replication process. In this context the British systems biologist Mark Isalan notes:

So perhaps the most exciting thing about a whole-cell model is that it may allow us to look beyond the direct molecular “cogs and wheels” that drive biology and into the emergent properties of biological systems. (Isalan 2012, 41)

Note that in this context Isalan does not seem to understand “emergent” in a strong sense as something that cannot be explained or predicted on the basis of underlying molecular processes. Instead, he seems to refer to the fact that the cell cycle control can be understood only when different modules of the system are integrated. This suggests that what systems biologists mean by “emergent” might refer to those behaviors that are left out of the traditional picture of molecular biology due to the biases of decomposition and localization. While predictions of this kind will have to be further investigated experimentally, they point to a way in which whole-cell modeling might not only correct scientists’ ideas about specific mechanisms, but lead to a revised picture of biological organization. Even though the model is based on strong assumptions of modularity itself, it can reveal forms of “distributed functionality” that do not coincide with decompositions into structural modules (Krohs, 2009). It is conceivable that such integrated behavior is ubiquitous in living systems and has up to now largely slipped through the cracks of the framework of modular reductionism. Whole-cell models can be highly valuable by offering educated guesses on where to look for it.

Karr et al. further tested their model by comparing its “phenotypic” behavior against direct experimental observations. The most impressive result in this regard seems at first glance to be the model’s ability to predict the essentiality of genes with 79% accuracy. However, this result has to be put into perspective: given that the large majority of genes in *M. genitalium* are essential, it is maybe not so surprising that the model does a good job in this respect.³ This is not

³Among the experimentally tested genes, about 85% turned out to be essential for the bacterium,

to say that Karr et al.'s result isn't highly statistically significant (that is, it cannot be explained by chance alone), but it is probably not as striking as it might seem at first. Rather than to celebrate this as a big predictive success of the model, it might be more useful to focus on the reasons for deviations between the model and the real system that can be used as tools for further discoveries. The researchers, for instance, looked more closely at three genes whose disruption resulted in discrepancies between model prediction and observation. In one of the cases this prompted them to consider an additional enzymatic reaction that had not been included in the model before, while the other cases suggested slight parameter changes, consistent with the rest of the model's performance. However, such discrepancies can usually be solved only by performing additional experiments. Thus, having a whole-cell model does not dispense with traditional experimental biology.

In summary, the real strength of whole-cell modeling might lie in its ability to accelerate biological discovery by including additional constraints that are invisible when looking at individual chunks of a system. However, it also holds the potential to change our ideas of biological organization. Karr et al.'s model provides one of the few examples of an attempt to completely recompose the information generated in molecular biology, biochemistry and systems biology. Such models carry the potential for uncovering aspects of systems that are hidden by partial representations, or what some systems biologists call "dynamic interfaces" between processes. Important aspects of such interfaces are often discovered through model failure, i.e. through failed attempts to directly integrate different processes. It is thus important to consider the value of such models not as a representational end product but as an epistemic tool with which biologists can arrive at new discoveries and probe their ideas of biological organization.

4 Multi-Scale Modeling

In parallel with the attempt of building whole-cell models of bacteria, several large-scale modeling projects take up the astonishing challenges of modeling human cells, organs or even the whole human body in detail. Cardiac modeling benefits from several decades of efforts to combine experimental, mathematical and computational strategies, starting with Noble's first simulation of the pacemaker in the early 60s (Noble 1960, 1962; Kohl and Noble 2009). Since the heart is currently the most developed "virtual organ," cardiac modeling will be the focus in this section.

To better understand the implications of the strategies used in such models, we first clarify a challenge for multi-scale modeling in general, known as the "tyranny of scales" problem in the context of physics (Oden 2006; Batterman 2012). The problem refers to the observation that no single mathematical model is sufficient to capture behaviors at all spatial scales. Many physical properties and the concepts used to describe them vary with scale. For instance, whereas surface properties are typically negligible when developing macroscale chemical models due to their minor impact at this scale, they dominate the dynamics of materials and particles at the nanoscale (Bursten 2015). Because the significance and conceptual stability of aspects such as surface tension is *multi-valued* across scales (Wilson 2012), the modeler must combine mathematical models that rely on different theoretical assumptions about the target system. The problem has so far not received much attention in the philosophy of biology (see however Lesne 2013; Green and Batterman 2017), but it has important implications for biological research as well as

compared to 71% in the model. If one randomly assigned the genes in the model to the two groups "essential"/"non-essential," while keeping group sizes constant, one would obtain an accuracy of 65% purely by chance!

for philosophical accounts of reductionism.

Bottom-up reductionism, as introduced in Section 1, assumes that it is possible to account for macroscale features through modeling of lower-scale processes. Interestingly, however, this strategy meets limitations even when modeling apparently simple physical system such as a steel beam. Batterman (2012) highlights how the structure and behavior of steel is scale-dependent. Accounting for the regular lattice structure at the atomic scale requires structural models, but at higher scales the material exhibits elastic behavior that is best accounted for through continuum models that ignore the micro-scale structure. At the mesoscale more diverse structures can be observed such as cracks and grain boundaries. Importantly, there is no manageable way to capture all scale-dependent behaviors in one model. A complete understanding of the bending properties of steel thus requires that the gap between the scales can be bridged through a combination of models that account for different aspects of the system.

Life scientists must also account for scale-dependency when modeling structures like the human heart in which the relevant processes span several spatial scales, from nanometers to centimeters, and which is functionally integrated within the human body as a whole. Like in the steel example, some aspects of biological systems require continuum models whereas other aspects require that the structural diversity of the system is accounted for. Continuum models are typically used at higher scales where the effects of many constituents average out, and the models aggregate discrete entities in a continuous variable (typically via partial differential equations). Even though this is clearly an idealization of the system, such idealizations are often required for solving the equations, and they are justified when the macroscopic behaviors are relatively independent of micro-scale properties (Batterman 2017). Sometimes, however, the system behavior is sensitive to microscopic processes. In such cases, continuum models have to be combined with lower-scale discrete models. A major challenge is therefore to connect models that make different assumptions about the system because they target processes at a specific scale.⁴

The discrepancy between continuum and discrete models is coupled to the additional challenge that processes modelled at different scales usually operate on different time-scales. Modeling of processes at the subcellular scale, such as fluxes of ions across the cell membrane, typically requires stochastic models because the dynamics at this scale is “dominated by random and short time fluctuations” (Qu et al. 2011, 22). At higher scales, coarser-grained deterministic models typically give more robust predictions. Because different aspects of the system dominate the behavior at characteristic scales, multi-scale modeling requires the combination of models that rely on conflicting idealizations and often make different predictions about what will happen with the same system over time. Bridging the gap between continuum and structural models, and between deterministic and stochastic models, is a hard challenge. At the same time, however, the requirement for and the success of coarser-grained macroscale models suggests the relative independence of some macroscale properties on lower-scale details (see also Batterman 2017). Accordingly, the tyranny of scales problem presents a severe challenge to bottom-up reductionism by offering resistance to the idea that models of processes at the lowest scale are sufficient for understanding the system as a whole.

In the following, we shall draw on recent collaborative efforts to develop multi-scale sim-

⁴Another important challenge is that the functioning of the heart involves multiple orchestrated processes that require what researchers call “multi-physics” models. This challenge is faced already in whole-cell modeling but the challenge is greater when the modeling task involves the combination of highly diverse models such as biomechanical models of tissue-scale deformations and gene regulation. Clarifying this challenge in detail is, however, beyond the scope of this paper.

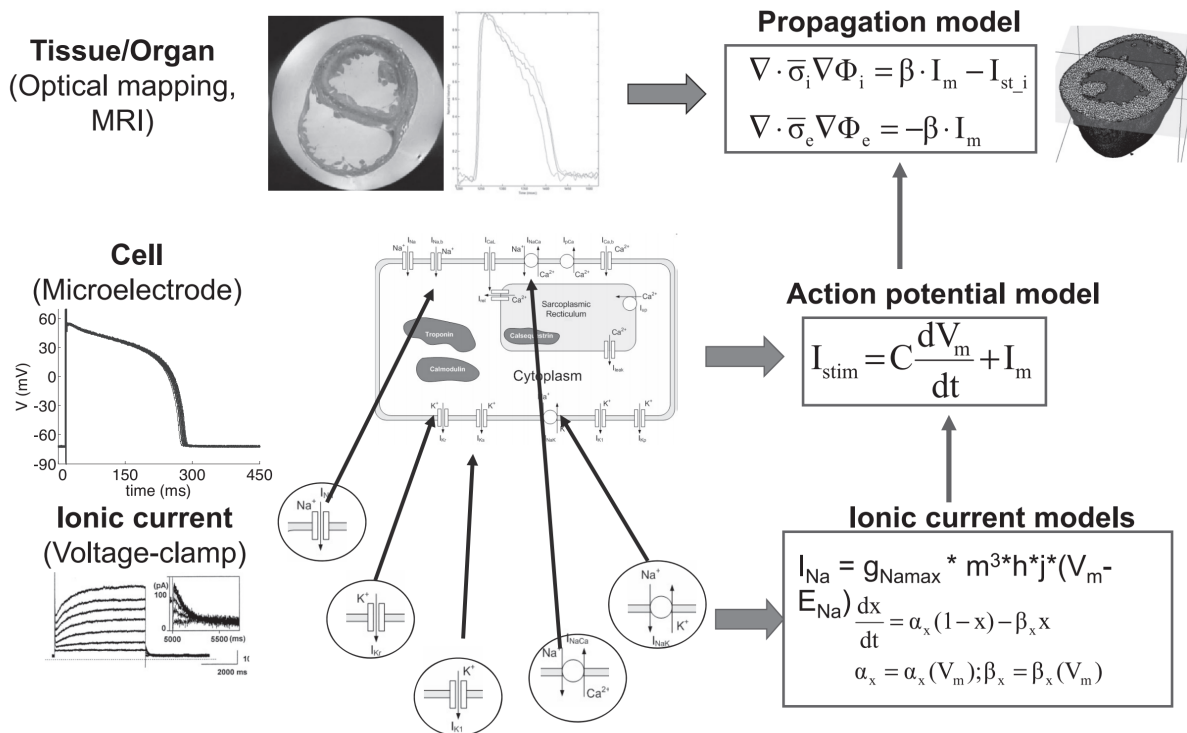


Figure 2: Cardiac modeling at different scales. Reprinted by kind permission from the American Physiological Society, from Carusi et al. 2012.

ulations of the human heart. Developing a detailed and predictive model of a human heart is a challenging task because of the multitude of processes occurring at different scales that are coupled via complex feedback mechanisms between scales. We draw on the example to show how these computational models are inherently multi-scale (temporally and spatially), integrating different models of processes at the subcellular, cellular, tissue, and organ level into 3D anatomical models in organ-level simulations. Sections 4.2 and 5 further discuss the epistemic and ontological implications of scale-dependency for the debate on reductionism.

4.1 Modeling the Human Heart

Multi-scale cardiac models are to some extent constructed in a bottom-up fashion. Ionic current models (molecular level) provide inputs to action potential models (cell-level), and these in turn provide inputs to propagation models (tissues or whole organ) (Carusi et al. 2012; Southern et al. 2008). Figure 2 illustrates the direction of such inputs in a simplified representation of the relations between selected models at different scales.⁵ Furthermore, it illustrates the requirement for meso- and macroscale parameters and models. A closer examination of the modeling procedure reveals that the models are intertwined in iterative loops with important feedback from higher-scale models to lower scales.

The reductionist may assume that detailed knowledge of all micro-level parameters and variables is sufficient. Some have even argued that including higher-level or top-down effects leads to causal overdetermination and possibly a violation of physical laws governing the lowest level

⁵For a comprehensive overview of models used at different scales, see (Southern et al. 2008; Chabiniok et al. 2016).

(e.g., Kim 1998). However, the central issue here is that some lower-scale processes are impossible to model and understand without including constraints imposed by the higher-scale system or subsystem in which these are embedded. Modeling a multi-scale system bottom-up is not possible because microscale processes operate differently than they would outside the context of the system (see below). Importantly, understanding top-down effects as constraining relations on the organization and degree of freedom of lower-scale processes does not lead to the problems associated with downward causation understood as efficient causation between ontologically different domains (Kim 1998; see also Emmeche et al. 2000).

The notion of downward causation we have in mind does not violate physicalism as outlined in the introduction. What is rejected instead is the idea that the multi-scale system can be modeled and explained purely in a bottom-up fashion (see also Love and Hütteman 2012). In the case examined here, the heart rhythm is generated by electrical potential in nerve cells and is *constituted* by the so-called Hodgkin cycle of oscillating ionic current across a membrane. Oscillations occur via gating of protein channels. However, the gating of ion channels is also determined by the cell voltage (a cell-level parameter) that is influenced by intercellular coupling and the dynamics of other processes across the membrane. Accordingly, the behavior of the system cannot be understood from an analysis of the constituents in isolation or at the lowest scale. Patch-clamp techniques allow for measurement of the stochastic behavior of single ion channels, but modeling of these processes must account for the constraints imposed by the coupled cells and biases associated with the isolation procedure (Carusi et al. 2011). To account for such biases, modelers incorporate measurements of the cell voltage via microelectrodes or draw on indirect estimations of current conductance in tissue preparations where cell membranes and coupling are intact, by measuring the effects of blocking of ionic currents. Noble (2012) highlights that such difficulties are not only practical challenges for bottom-up modeling but are also evidence for top-down constraints in biological systems. We clarify his argument below.

Ionic current models are typically ordinary differential equations capturing the kinetics of each component (e.g. individual ion channels). Intuitively, it should be possible to model the heart rhythm bottom-up, given enough information on the protein mechanisms involved in the gating of ion channels. However, the equations cannot be solved without defining the state of the components (the initial conditions) and the boundary conditions. Boundary conditions are mathematically defined restrictions that specify the domains and conditions under which a given mathematical model holds. Imposing boundaries on the domain of the model or specifying the value of the solution are often required for solving mathematical equations. However, these are not just imposed to make models more tractable. Attention to boundary conditions is often important for understanding boundary behaviors, e.g. for understanding which aspects can be ignored when modeling processes at characteristic scales. In this context, boundary conditions are used to represent physical constraints that limit the degree of freedom of a lower scale process (e.g. cell voltage, tissue stiffness etc.). Noble (2012) argues that without such boundary conditions, biological functions like the heart rhythm would not exist. In support of this view, he refers to the result of a computer experiment with a simple model of the heart rhythm. The simulation allowed him to computationally remove the feedback from the cell voltage with the result that the oscillations in the flow of ions ceased.

The same lesson can be drawn when we move up to the scale of cells, where action potential models require inputs from models of the dynamics and constraints of tissue-structures, e.g. via partial differential equations representing propagation of electrical currents through tissues (Figure 2; see also Carusi et al. 2012). To model how electrical currents propagate through the heart as a whole, the researchers must determine the relevant tissue boundaries and biomech-




Figure 3: Illustration of the feedback between differential equations and initial and boundary conditions. Adapted from Noble (2012).

anical properties that influence the electrophysiological properties of the organ. The parameters required for this type of model include fiber conductivity and fiber architecture (the local direction of the conductivity tensor), tissue stiffness etc. Such parameters cannot be determined in a bottom-up fashion but involve optimal mapping techniques (e.g. diffusion tensor-MRI) and biomechanical modeling of the complete tissue structures or even the whole organ.

Figure 3 offers a simple illustration of the relations between differential equations, initial and boundary conditions. Noble (2012) represents initial conditions as operating at the same level as the differential equations because they describe the inputs to a kinetic model in terms of the state of the system at time t . Boundary conditions, on the other hand, define the constraints imposed on the system from the environment (e.g. the extracellular matrix) or from the systemic context of the mechanism or multi-scale system (e.g. the context of the cell or tissue structure). Boundary and initial conditions are, however, not fixed once and for all but can change during the dynamic process. For instance, the stiffness of a cell or tissue structure can change in response to expression of specific proteins. There is therefore a complex and continuous feedback between processes at different scales.

Importantly, Noble considers the need for higher-scale boundary conditions and parameters as a form of downward causation (Noble 2012; see also Emmeche et al. 2000). In the following, we clarify the central role of boundary conditions in connecting levels and show why higher-level and top-down relations are indispensable.

4.2 Connecting Levels

Modeling of a multi-scale system requires that processes at different levels are connected. Re-composing models at different scales while accounting for the nonlinear feedback across scales is rendered difficult by the tyranny of scales problem and the diversity and complexity of the processes involved. As mentioned, no single experimental or mathematical strategy is sufficient,

and modelers must integrate data acquired at different levels, drawing on different preparations and techniques and mathematical models relying on different and often conflicting assumptions. Continuum models and discrete models, as well as stochastic and deterministic models, must be combined through careful attention and decision-making on which aspects are dominant (or can be neglected) at different scales. Boundary conditions play a crucial role for this purpose. As described in the literature on multi-scale modeling in physics and nanoscience, models need to be combined at the point where boundary conditions for one model can no longer be ignored (e.g. Batterman 2012). Drawing on examples from nanoscience, Bursten (2015) describes such relations as non-reductive model interactions.

To understand this concept better, we must elaborate further on the tyranny of scales problem. As Noble himself acknowledges, accounting for the influence of higher scale processes via initial and boundary conditions for differential equations is not enough. Differential equations are useful for describing the kinetics of component parts or pathways, but they face limitations when the aim is to model the electrophysiological or biomechanical properties of tissues. As we have seen, tissue-scale models typically consist of continuum models (partial differential equations) that are not straightforwardly integrated with the discrete cellular and subcellular models (Figure 2). Moreover, different physical and biological processes at intermediate scales require a variety of modeling procedures that cannot easily be combined. Merging different modeling styles by incorporating all parameters into one unified framework is typically not possible. Instead, modelers combine models by using the results of one model as input for another while taking into account which models are most appropriate for specific scale-dependent processes and their relative levels of uncertainty. For instance, modelers may formulate initial and boundary conditions for one spatial scale via the results of another model at a higher or lower scale (Southern et al. 2008). Ultimately, the aim is to develop an integrated *system of models* into a whole heart simulation.

In practice, simulations of the heart are performed using finite element methods on an anatomical 3D mesh based on MRI imaging.⁶ The mesh on one hand integrates and solves equations for the component models but also incorporates anatomical information about fiber architecture and tissue conductance at separate mesh points. A common strategy in cardiac modeling is to develop a bidomain model that relates the discrete models of ion channels and cells to a continuum model that models cardiac cells as a continuous network of resistors. Partial differential equations in the bidomain model, together with physical laws of conductance, resistance and charge conservation, are used to integrate the electrical potentials and current flow in the different tissue nodes. The bidomain model relates the electrical potential to current flow in two domains (the intracellular and extracellular media) by using the discrete subcellular and cell models as inputs to a continuous model (Carusi et al. 2012; Southern et al. 2008). Via the mesh, the bidomain equations are solved numerically and integrated as discrete models (finite elements) using a simulation software. Thus, the mesh mediates between models and

⁶Finite element methods divide complex structures into finite element subunits. They are commonly used in engineering and biomechanical analysis to find approximate solutions for partial differential equations representing spatial structures. 3D anatomical meshes used in cardiac modeling can consist of over thirty million finite elements (Carusi et al. 2012).

simulations.⁷

As mentioned above, the integration cannot happen in a straightforward bottom-up fashion because of the feedback across scales. Carusi et al. (2012) describe the construction of a multi-scale model as a continuous process. The dependence on parameters that cannot be accounted for in one type of measurement or by one modeling framework makes it is very difficult to validate models at each scale separately. Moreover, they add that even if this were possible “there is inevitably a further level of complexity in the integration of levels, and it is difficult, if not impossible, to draw the boundaries of this complexity in advance” (Carusi et al. 2012, H150). Despite these difficulties, the approach is justified because it is the only way to account for the ways in which the whole exerts constraints on the behavior of the parts. Multi-scale modeling thus reveals how biological systems are at the same time *more* and *less* than the sum of the parts. They are *more* in the sense that system-level properties cannot be predicted from properties of component parts alone. But they are *less* in the sense that the organization and constraints of the system prevent the components from exhibiting some of the properties and degrees of freedom that they may have in different contexts (Morin 2008). Without such constraints, biological functions may not be possible at all (Noble 2012).⁸

5 Can Large-Scale Modeling Conquer Biological Complexity?

Large-scale modeling can be seen as a response to biological complexity and as an attempt to account for the biases of traditional reductionistic strategies. Karr et al.’s whole-cell model challenges the assumption of modular reductionism according to which different cellular activities can be studied separately. The example of cardiac modeling, by contrast, illustrates a way of recomposing insights about processes at different scales. In this case the indispensable roles of meso- and macroscale parameters and models question the assumption of bottom-up reductionism that multi-scale systems can be understood from information at the molecular level alone. Both case studies suggest that important features of biological systems may be lost at the interfaces of different modules or scales, and they provide novel insights into biological organization.

At the same time, the examples show that even the most advanced models in systems biology rely on a number of simplifying strategies and are far from unbiased themselves. The WCM makes important progress towards building a more integrated representation of a living cell but itself relies on an assumption of modularity. Is this assumption justified? While some degree of modularity seems necessary for the evolution of complex organisms (e.g. Simon 1962; Wagner and Altenberg 1996), it has also been argued that modularity comes with a cost and that in certain contexts evolution may favor more integrated architectures (Wimsatt 2007; Krohs 2009). Karr et al. do not provide any further justification for their particular interpretation of biological modularity and confine themselves in this regard to the following statement:

⁷It is interesting to note that modeling the heart involves the construction of continuum models via inputs from discrete models, but computational simulation requires a final discretization procedure to solve the complex equations of the bidomain model. While these aspects may be considered mainly as methodological strategies to deal with biological complexity, they call for modesty with respect to one unifying or prioritized level of causation and explanation.

⁸As mentioned in Section 4.a., Noble (2012) highlights that functions like the heart rhythm are constituted by cellular and tissue-level constraints on the pulse-generating oscillations of ionic flow. The idea that organisms are less than the sum of the parts has also been expressed by other systems biologists. For instance, Hofmeyr (2017) has used this expression in arguing that enzyme activity is dependent on constraints given by the cellular and environmental context.

Because biological systems are modular, cells can be modeled by the following: (1) dividing cells into functional processes; (2) independently modeling each process on a short timescale; and (3) integrating process submodels at longer timescales. (Karr et al. 2012, 399)

It is not obvious whether modularity necessarily implies the separation of time scales that this approach is based on. One might argue, however, that in analogy with the numerical solution of a differential equation approximating the true analytical solution, the WCM will approximate the true integrated nature of the target system. But it is probably difficult to assess at what time resolution this approximation will be accurate enough and whether the kind of computational power needed is realistically available. The accuracy will be limited by the degree of resolution of the individual sub-models which, as we have seen, can be of very different types. Moreover, even if it is true that a biological system is modular, it is an altogether different question of whether a particular decomposition into modules is appropriate. To construct their model, Karr et al. had to build on a particular cellular decomposition stemming from the results of previous biological research. For this reason, the model may have inherited some of the biases of modular reductionism. It is difficult to imagine how one could make the model sensitive enough to detect all of these biases, unless all the molecular properties in their cellular context were known with very high precision. As we have seen, this is not the case, and many processes are represented in very crude ways or even completely black-boxed.

Another way in which the whole-cell model is clearly still reductionistic is that it neglects or at least simplifies the environmental context of the bacterium. Systems biologists have to answer the question of what actually the relevant systems are that they should study (cf. Dupré and O'Malley 2005), which can be understood as a higher level question about modularity and system boundaries. Is the most relevant system necessarily the cell or are interactions between cells and organisms equally relevant for a functionally complete model?

Similarly, an important aspect of multi-scale cardiac modeling is to define the relevant boundaries of a system. To what extent is the predictive and explanatory power of multi-scale cardiac models dependent on inclusion of feedback beyond the organ system—such as the whole cardiac system, the whole human body, the environment and other organisms? The direct and easily measurable impact of other systems and environmental influences on the heart rate is a powerful reminder of the integrated and responsive nature of living systems. A natural progression from whole organ models is therefore the construction of a “virtual physiological human” (VPH) aiming for an integrated and systematic understanding of biological processes spanning all relevant levels and time-scales through the use of mathematical modeling (Hunter et al. 2013; Southern et al. 2008).⁹ The goal of such models is, however, not to build the most detailed model possible, but to investigate how much knowledge we need to integrate in order to develop more predictive models. As a review article on multi-scale cardiac modeling states, “it could be by incorporating more detail that we will eventually find it possible to derive the most useful mathematical reductions. This might be a case of exploring more to focus on less” (Davies et al. 2016, 936)

⁹The VPH was launched in October 2011 by the European Commission as a project to set out a roadmap for the “digital patient,” a computer simulation that will generate a virtual version or “3D avatar” of individual patients with the aim of simulating disease processes and drug responses to make medically relevant predictions at the level of the individual patient. The European Virtual Physiological Human project is an integral part of the international Physiome Project (www.physiome.org.nz), initiated in 1997 by the first international efforts to provide database resources for integrating experimental data and computational models.

Increasing the size and complexity of models inevitably leads to an increase in the sources of error, variability, and artifacts, and to complications for model validation (Carusi 2014). Despite the amount of empirical information that went into building the WCM, the parameter values for many processes are still not known well enough. These parameters, therefore, have to be fit or adjusted in order to fulfill certain basic observational constraints and to be consistent with the other processes. For example, the *metabolism process*, which describes the import of nutrients and their conversion into building blocks for macromolecules, was fit to match the observed mass doubling time of *M. genitalium* by means of a method called flux balance analysis (FBA). When models with free parameters reach a certain degree of complexity, there is the general risk that one will always find changes in some of the parameter values to obtain a fit with the empirical data, even if the actual cause of a deviation lies in the structure of the model. One might expect in this regard that the in-built modularity of Karr et al.'s model can facilitate the "debugging," but as we have just shown modularity itself remains an assumption in the model that needs more work in order to be justified. It is thus crucial to determine the relevant level of detail in order to minimize the problem of overparameterization.

Whereas some systems biologists have argued that large-scale modeling breaks with the principle of Occam's razor (Kolodkin and Westerhoff, 2011), others working on whole-organ models stress that: "Occam's principle must be applied at each level to avoid parameter accumulation" (Carusi et al. 2012, H149). Similarly, Kohl and Noble (2009, 4) draw the lesson that "theoretical models of biological behavior are most efficient when they are *as complex as necessary, yet as simple as possible*." It might be tempting to say that large models are simply better and more powerful than smaller models, but there are in fact important trade-offs between the level of complexity that is captured by a model and its ability to lead to successful inferences. A large and very detailed model of a system can be useless or even misleading if we do not have the right kind of data to specify its free parameters. So perhaps most importantly, the level of model complexity has to be balanced against the amount and quality of available empirical information (Gross and MacLeod 2017). Occam's razor remains valid as a methodological principle for modeling, even if we do not expect the underlying systems to be simple.

While validation clearly is a problem for all computational models, large-scale modeling projects face this problem in a more aggravated way since it involves both the validity of the individual sub-models and the validity of their integration. Moreover, both the whole-cell and the multi-scale simulation consist of several models drawing on data from multiple target systems (and different experimental preparations of these). For this reason the boundaries between target and model, and between experiment and simulation, often become blurred (see also Carusi 2012).

Despite these challenges, large-scale modeling may be worthwhile because there may be no other way to gain insight into the integrated nature of living systems and to understand complex diseases. Moreover, the challenges just described can be seen in a more positive light, as providing new opportunities to grasp biological organization and complexity. As we have seen in the case of the WCM, important insights often result from model failure when there are deviations between predictions of different models, or between models and experimental results. Similarly, challenges faced in cardiac modeling have led to increasing acknowledgement of how complex the orchestrated feedback loops between scales are. Looking back at nearly 50 years of cardiac modeling, Kohl and Noble (2009, 3) state that the accumulating body of insights "derived as much from the 'failures' as from the 'successes' of theoretical prediction and experimentation validation." They regard the contradiction of predictions as often being more instructive than their confirmation in leading to new advances of the field. Recent years

of detailed modeling efforts have for instance revealed that entities previously thought to be linked in a simple one-directional way exhibit significant cross-talk, e.g. between ion channels and tissue behavior, electrophysiology, and mechanics (see Carusi et al. 2012, H150, for further examples). Obstacles to model integration in large-scale simulations may therefore uncover limitations in our assumptions about functional isolation of modular components or processes.

As argued in Section 4, the strategies used in the construction of multi-scale cardiac models reveal problems for the reductionist ideal of bottom-up construction. It is simply not possible to model a complex system like the human heart bottom-up because macroscale parameters and models are required to account for the higher-level constraints such as cell voltage, tissue stiffness and geometry etc. Thus, multi-scale modeling reveals the limitations of reductionism and the explanatory relevance of top-down effects.

Insofar as the limitations arise because “the principal physics governing events often change with scale” (Oden 2006, 2930), the requirement for different mathematical models has deep ontological implications. Scientists interested in constraints on form and function in biology have long been aware of the implications of scale-dependency.¹⁰ As we study physical processes across multiple scales, the significance of gravity, surface tension, inertia, electrical charges, the viscosity of the medium etc. changes (Purcell 1977; Thompson 1917/1992; Vogel 2009). As we have seen, the applicability of deterministic and stochastic modeling also depends on spatial and temporal scaling (Qu et al. 2011). Accordingly, there seems to be no privileged ontological level from which all relevant aspects of a multi-scale system can be studied (Noble 2012).

The most fruitful way of defining the system boundaries and the relevant level of detail depends on the specific biological question and the biological complexity that continuously forces modelers and experimentalists to draw on strategies of simplification. As Woodward emphasizes, this view “contrasts with the common philosophical tendency to think there is a single, universal level of causal description that is most appropriate” (Woodward 2010, 297). This does not preclude that some approaches are more appropriate given a specific question or a specific system. For instance, some cardiac conditions may be well understood through a genetic analysis whereas others require a multi-scale analysis (Noble 2012). Thus, the relevant scale for studying living systems can typically not be predicted in advance but depends on the research question, the causal structure of the specific system, as well as on the methodological constraints.

6 Conclusion

Systems biologists often highlight the need to overcome the reductionist approach of molecular biology. In this paper we have discussed two different ways in which reductionism can be understood in this context. We have distinguished between *modular reductionism*, which is the position that living systems consist of or can be understood and studied by looking at individual functional sub-units, and *bottom-up reductionism*, which refers to the idea that biological phenomena can be understood and studied by focusing on the molecular level. We have looked at two detailed case studies in systems biology in order to investigate the strategies that are proposed to go beyond the reductionistic approach.

We have discussed how the WCM manages to take into account the interfaces between the 28 cellular processes modeled for the simple bacterium. It reveals that certain processes, such

¹⁰Already Galilei (Discorsi, 1638) recognized that the relation between the bone structure and animal size is disproportionate because the weight increases with a higher magnitude than the strength of the material.

as cell cycle regulation, might be more integrated into the overall machinery of the cell than previously assumed and suggests a way to understand one sense in which systems biologists speak of “emergent” properties. Our second example, multi-scale cardiac modeling, shows how scale-dependency forces modelers to adopt scale-specific modeling strategies and makes clear why a single level of analysis is insufficient for systems that span multiple spatial scales. Since macro-scale parameters are needed to account for the ways in which the environmental and systemic contexts constrain lower-scale processes, multi-scale modeling questions the reductionist ideal of bottom-up modeling. Moreover, multi-scale modeling can reveal which parts of the system influence the state of the system as a whole. Large-scale modeling thus illuminates a sense in which living systems are more than the sum of their parts, but also how their functions are a result of the constrained organization of these processes. We have argued that this has important philosophical implications for bottom-up reductionism but also practical implications if the aim is to develop predictive cardiac models with biomedical applications.

We have also highlighted the challenges that are involved in large modeling projects, and we are far from recommending such models as the only true approach to systems biology. A model, even if it has the appearance of completeness, is never an all-purpose tool but is designed and limited to address certain questions. Even though realizing the dream of a complete model is sometimes presented as a goal in itself, like sending a man on the moon, scientists themselves are usually aware of the limitations, and they work on large-scale models with specific goals in mind. For instance, in the report of the STEP Consortium (Structuring the Europhysiome), more than 300 stakeholders involved in the realization of the VPH explicitly state that the roadmap of VPH: “[...] must not be taken as a promise to deliver an all-inclusive mathematical model of the human organism, a goal which is not only unrealistic technically (and will be for the foreseeable future) but also, it can be argued, is unrealisable even in principle—the only complete model is the organism itself.” (STEP Consortium 2007, 34).

Similarly, it can be asked whether the idea of a whole-cell model remains realistic if we go beyond a small parasitic bacterium with a near-minimal genome. It might simply exceed any available computational power in the foreseeable future to scale up from *M. genitalium* to more complex organisms. An obvious next project for whole-cell modelers would be the standard model bacterium *E. coli*.¹¹ Yet, this step already corresponds to a tenfold increase in genome size which gives an impression of the complexity that the corresponding model would have to capture. Importantly, however, the utility of large-scale models should not be evaluated on their ability to capture all processes in detail but on the extent to which they provide novel insights to biological complexity and increase our ability to explain and predict living systems. Complex cardiac models developed so far provide reasons for optimism regarding the use of computer simulations for testing for the effects of drugs and surgical interventions, and may in time be useful also for disease prediction and prevention (Noble 2008; Clermont et al. 2011; Chabinio et al. 2016).

Both case studies indirectly inform us about biological organization and complexity by revealing how difficult it is to find adequate representations of biological processes when going beyond the simplifications of reductionist strategies. They require sophisticated ways of combining and integrating different types of mathematical models and rely on alternative simplifying strategies to arrive at tractable research problems. Furthermore, they both rely on knowledge that has been accumulated in more detailed investigations of individual processes or at specific levels.

¹¹It appears that some of the people who built the *M. genitalium* model are now working on a whole-cell model of *E. coli*. See <https://simtk.org/projects/ecoli>.

We conclude from our analysis that large modeling projects should be seen as specific tools that hold the potential to assess and test our assumptions about biological organization and complexity. They can be used to detect reductionistic biases even though they are not without bias themselves. They do not obviate the need for small scale experiments and partial models, but they can complement more traditional techniques in order to get closer to the way in which biological systems function as wholes.

Acknowledgments

We would like to thank Adam Ferner and Thomas Pradeu for editing this special issue, and Miles MacLeod, Thomas Reydon, Mark Bedau, and an anonymous reviewer for useful feedback.

Literature cited

- Bassingthwaighte, J., Hunter, P., and D. Noble. 2009. "The Cardiac Physiome: Perspectives for the Future." *Experimental Physiology* 94 (5): 597–605.
- Batterman, R. W. 2012. "The Tyranny of Scales." In *Oxford Handbook of Philosophy of Physics*, edited by R. W. Batterman, 255–286. Oxford: Oxford University Press.
- Batterman, R. W. 2017. "Autonomy of Theories. An Explanatory Problem." *Nous*. doi:10.1111/nous.12191.
- Bechtel, W. and A. Abrahamsen. 2010. "Dynamic Mechanistic Explanation: Computational Modeling of Circadian Rhythms as an Exemplar for Cognitive Science." *Studies in History and Philosophy of Science Part A* 41 (3): 321–333.
- Bechtel, W., and R. C. Richardson. (1993) 2010. *Discovering Complexity*. Princeton: Princeton University Press.
- Brigandt, I., and A. Love. 2017. "Reductionism in Biology." In *Stanford Encyclopedia of Philosophy* (Spring 2017 edition), edited by Edward N. Zalta. <https://plato.stanford.edu/archives/spr2017/entries/reduction-biology/>.
- Bursten, J. 2015. *Surfaces, Scales and Synthesis: Reasoning at the Nanoscale*. PhD dissertation, University of Pittsburgh.
- Burian, R. M. 1993. "Technique, Task Definition, and the Transition From Genetics to Molecular Genetics: Aspects of the Work on Protein Synthesis in the Laboratories of J. Monod and P. Zamecnik." *Journal of the History of Biology* 26 (3): 387–407.
- Callebaut, W., and D. Rasskin-Gutman, editors. 2005. *Modularity: Understanding the Development and Evolution of Natural Complex Systems*. Cambridge: The MIT Press.
- Carusi, A. 2014. "Validation and Variability: Dual Challenges on the Path from Systems Biology to Systems Medicine." *Studies in History and Philosophy of Biological and Biomedical Sciences* 48: 28–37.
- Carusi, A., K. Burrage, and B. Rodríguez. 2012. "Bridging Experiments, Models and Simulations: An integrative Approach to Validation in Computational Cardiac Electrophysiology." *American Journal of Physiology: Heart and Circulatory Physiology* 303(2): H144–H155.
- Chabiniok, R., V. Y. Wang, M. Hadjicharalambous, L. Asner, J. Lee, M. Sermesant, E. Kuhl, et al. 2016. "Multiphysics and Multiscale Modelling, Data–model Fusion and Integration of Organ Physiology in the Clinic: Ventricular Cardiac Mechanics." *Interface Focus* 6 (2): 20150083.

- Clermont, G., C. Auffray, Y. Moreau, D. M. Rocke, D. Dalevi, D. Dubhashi, D. R. Marshall, et al. 2009. "Bridging the Gap between Systems Biology and Medicine." *Genome Medicine* 1 (9): 88.1–88.6.
- Crick, F. 1973. "Project K: 'The Complete Solution of *E. coli*.'" *Perspectives in Biology and Medicine* 17: 67–70.
- Davies, M. R., K. Wang, G. R. Mirams, et al. 2016. "Recent Developments in Using Mechanistic Cardiac Modelling for Drug Safety Evaluation." *Drug Discovery Today* 21 (6): 924–38.
- Emmeche C., S. Køppe, and F. Stjernfelt. 2000. "Levels, Emergence, and Three Versions of Downward Causation." In *Downward Causation: Minds, Bodies and Matter*, edited by P. Bøgh Andersen, C. Emmeche, N.O. Finnemann, and P. Voetmann Christiansen. Aarhus: Aarhus University Press.
- Gibson, D. G., J. I. Glass, C. Lartigue, et al. 2010. "Creation of a Bacterial Cell Controlled by a Chemically Synthesized Genome." *Science* 329 (5987): 52–56.
- Gilbert, S. F., and S. Sarkar. 2000. "Embracing Complexity: Organicism for the 21st Century." *Developmental Dynamics* 219 (March): 1–9.
- Glass, J. I., N. Assad-Garcia, N. Alperovich, et al. 2006. "Essential Genes of a Minimal Bacterium." *Proceedings of the National Academy of Sciences of the United States of America* 103 (2): 425–430.
- Green, S. 2015. "Can Biological Complexity be Reverse Engineered?" *Studies in History and Philosophy of Biological and Biomedical Sciences* 53: 73–83.
- Green, S., and R. Batterman. 2017. "Biology Meets Physics: Reductionism and Multi-scale Modeling of Morphogenesis." *Studies in History and Philosophy of Biological and Biomedical Sciences* 61: 20–34.
- Gross, F., and M. MacLeod. 2017. "Prospects and Problems for Standardizing Model Validation in Systems Biology." *Progress in Biophysics and Molecular Biology*. doi:10.1016/j.pbiomolbio.2016.11.005.
- Hofmeyr, J. H. S. 2017. "Exploring the Metabolic Marketplace Through the Lens of Systems Biology." In *Philosophy of Systems Biology: Perspectives from Scientists and Philosophers*, edited by S. Green. Springer International Publishing.
- Hunter P., T. Chapman, P.V. Coveney, et al. 2013. "A Vision and Strategy for the Virtual Physiological Human: 2012 Update." *Interface Focus* 3: 1–10.
- Isalan, M. 2012. A Cell in a Computer. *Nature* 488: 9–10.
- Kaiser, M.I. 2015. *Reductive Explanation in the Biological Sciences*. Switzerland: Springer International Publishing.
- Karr, J. R., J. C. Sanghvi, D.N. Macklin, et al. 2012. "A Whole-Cell Computational Model Predicts Phenotype from Genotype." *Cell* 150 (2): 389–401.
- Kendig, C., and T. Eckdahl. 2017. "Reengineering Metaphysics: Modularity, Parthood, and Evolvability in Metabolic Engineering." *Philosophy, Theory, and Practice in Biology* 9 (8). doi:10.3998/ptb.6959004.0009.008.
- Kim, J. 1998. "Making sense of emergence." *Philosophical Studies* 95: 3–36.
- Kitano, H. 2002. "Looking beyond the Details: A rise in System-oriented Approaches in Genetics and Molecular Biology." *Current Genetics* 41 (1): 1–10.
- Kohl, P., and D. Noble. 2009. "Systems Biology and the Virtual Physiological Human." *Molecular Systems Biology* 5: 292.

- Kolodkin, A. N., and H.V. Westerhoff. 2011. "Parsimony for Systems Biology: Shaving Occam's Razor Away." *European Communications in Mathematical and Theoretical Biology* 14: 149–152.
- Krohs, U. 2009. "The Cost of Modularity." In *Functions in Biological and Artificial Worlds: Comparative Philosophical Perspectives*, edited by U. Krohs and P. Kroes. Cambridge: MIT Press.
- Krohs, U., and W. Callebaut. 2007. "Data Without Models Merging With Models Without Data." In *Systems Biology: Philosophical Foundations*, edited by F. C. Boogerd, F. J. Bruggeman, J.-H. Hofmeyr, and H.V. Westerhoff. Amsterdam: Elsevier.
- Lesne, A. 2013. "Multiscale Analysis of Biological Systems." *Acta Biotheoretica* 61 (1): 3–19.
- Levins, R. 1966. "The Strategy of Model Building in Population Biology." *American Scientist* 54 (4): 421–431.
- Mitchell, S. 2003. *Biological Complexity and Integrative Pluralism*. Cambridge: Cambridge University Press.
- Morin, E. 2008. *On Complexity*. New York: Hampton Press.
- Morowitz, H. J. 1984. "The Completeness of Molecular Biology." *Israel Journal of Medical Sciences* 20: 759–753.
- Nicholson, D. J. 2013. "Organisms ≠ Machines." *Studies in History and Philosophy of Biological and Biomedical Sciences* 44 (4): 669–678.
- Noble, D. 1960. "Cardiac Action and Pacemaker Potentials based on the Hodgkin-Huxley Equations." *Nature* 188: 495–497.
- Noble, D. 1962. "A Modification of the Hodgkin-Huxley Equations Applicable to Purkinje Fibre Action and Pacemaker Potentials." *The Journal of Physiology* 160 (2): 317–352.
- Noble, D. 2008. "Computational Models of the Heart and their Use in Assessing the Actions of Drugs." *Journal of Pharmacological Sciences* 107 (2): 107–117.
- Noble, D. 2012. "A Theory of Biological Relativity: No Privileged Level of Causation." *Interface Focus* 2: 55–64.
- Ochman, H., and R. Raghavan. 2009. "Excavating the Functional Landscape of Bacterial Cells." *Science* 326: 1200–1201.
- Oden, J. T. 2006. *Finite Elements of Nonlinear Continua*. New York: Dover Publications.
- O'Malley, M. and J. Dupré. 2005. "Fundamental Issues in Systems Biology." *BioEssays* 27 (12): 1270–1276.
- Purcell, E. M. 1977. "Life at Low Reynolds Number." *American Journal of Physics* 45 (1): 3–11.
- Qu, Z., A. Garfinkel, J. N. Weiss, and M. Nivala. 2011. "Multi-Scale Modeling in Biology: How to Bridge the Gaps Between Scales?" *Progress in Biophysics and Molecular Biology* 107 (1): 21–31.
- Sanghvi, J. C., S. Regot, S. Carrasco, J. R. Karr, M. V. Gutschow, Jr., B. Bolival, and M. W. Covert. 2013. "Accelerated Discovery via a Whole-Cell Model." *Nature Methods* 10: 1192–1195.
- Simon, H.A. 1962. "The Architecture of Complexity." *Proceedings of the American Philosophical Society* 106 (6): 467–482.
- Southern, J., J. Pitt-Francis, J. Whiteley, et al. 2008. "Multi-Scale Computational Modelling in Biology and Physiology." *Progress in Biophysics and Molecular Biology* 96 (1): 60–89.
- STEP Consortium. 2007. "Seeding the EuroPhysiome: A Roadmap to the Virtual Physiological Human." <http://www.europhysiome.org/roadmap>.
- Thompson, d'A. W. (1917) 1992. *On Growth and Form*. Cambridge, UK: Cambridge University Press.

- Tomita, M. 2001. "Whole-Cell Simulation: A Grand Challenge of the 21st Century." *Trends in Biotechnology* 19 (6): 205–210.
- Vogel, S. 2009. *Glimpses of Creatures in Their Physical Worlds*. Princeton, NJ: Princeton University Press.
- Wilson, M. 2012. "What is Classical Mechanics Anyway?" In *Oxford Handbook of Philosophy of Physics*, edited by R. Batterman. Oxford: Oxford University Press.
- Wagner, G. P., and L. Altenberg. 1996. "Perspective: Complex Adaptations and the Evolution of Evolvability." *Evolution* 50 (3): 967–976.
- Wimsatt, W. C. 2007. *Re-Engineering Philosophy for Limited Beings: Piecewise Approximations to Reality*. Cambridge, MA: Harvard University Press.
- Winther, R. G. 2001. "Varieties of Modules: Kinds, Levels, Origins, and Behaviors." *Journal of Experimental Zoology* 291 (2): 116–129.
- Wolkenhauer, O., and S. Green. 2013. "The Search for Organizing Principles as a Cure against Reductionism in Systems Medicine." *FEBS Journal Reviews* 280 (23): 5938–5948.
- Woodward, J. 2010. "Causation in Biology: Stability, Specificity, and the Choice of Levels of Explanation." *Biology & Philosophy* 25 (3): 287–318.

© 2017 Author(s)

This is an open-access article distributed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International license, which permits anyone to download, copy, distribute, or display the full text without asking for permission, provided that the creator(s) are given full credit, no derivative works are created, and the work is not used for commercial purposes.

ISSN 2475-3025