



Model-Based Inferences in Modeling of Complex Systems

Miles MacLeod¹

© The Author(s) 2018

Abstract

Modelers are tackling ever more complex systems with the aid of computation. Model-based inferences can play a key role in their ability to handle complexity and produce reliable or informative models. We study here the role of model-based inference in the modern field of computational systems biology. We illustrate how these inferences operate and analyze the material and theoretical bases or conditions underlying their effectiveness. Our investigation reiterates the significance and centrality of model-based reasoning in day-to-day problem-solving practices, and the role that debugging processes of partial or incomplete models can play in scientific inference and scientific discovery, particularly with respect to complex systems. We present several deeper implications such an analysis has for philosophy of science regarding the role of models in scientific practice.

Keywords Model-based reasoning · Model-based inference · Complexity · Systems biology · Computation

1 Introduction

This paper explores the nature of scientific inference in discovery contexts. It aims to illustrate the fundamental importance of *model-based* inference in the construction of mathematical models of complex systems in the field of computational systems biology, emphasizing the central role partially complete and correct models play in solving such problems. This field aims to build detailed mathematical models of biological systems, particularly biochemical systems. Such systems are by nature complex. Systems biology represents to some extent trends in recent scientific practice towards applying relatively fast and available computer technology to complex systems. In certain types of modeling in systems biology, modelers rely heavily on the affordance of partial or incomplete models to draw inferences about their systems. Our primary goals are to (1) illustrate the structure and function of these inferences (their underlying heuristics and procedures) in systems biology—and describe in what sense they are model-based, and (2) analyze the basis or conditions underlying the ability to draw such inferences when handling complex systems.

Our analysis reiterates the significance and centrality of model-based reasoning in day-to-day mundane problem-solving practices, and the genuine role that otherwise humble engineering-type debugging processes of partial or incomplete models can play in scientific inference and scientific discovery (Wimsatt 2007). Importantly we show that the discovery processes we describe are strongly guided through various heuristics (for model manipulation) amounting to a set of procedures for discovery. In this regard we build on recent insights in discussion about scientific discovery, and show how even in complex cases there may well be an underlying logic guiding discovery rather than pure intuition alone (Ippoliti 2017, 2018a, b). Indeed the structure underlying discovery in these cases raise several implications for our understanding of the epistemic roles models play in scientific practice and the methodological choices researchers make. To meet these goals we draw on a relatively fine-tuned ethnographic study of model-building practices in systems biology.¹

The paper is structured accordingly. Section 2 of this paper outlines the problem spaces common to many modeling in systems biology, and gives an account of the

✉ Miles MacLeod
m.a.j.macleod@utwente.nl

¹ Department of Philosophy, University of Twente, Drienerlolaan 5, 7522 NB Enschede, The Netherlands

¹ This study was 4-year ethnographic study of model-building practices in two systems biology labs led by Professor Nancy Nersessian. Our group performed laboratory observations and interviews of lab participants, some over the course of their graduate research projects. This research was funded by the US National Science Foundation which requires following human subjects' regulations. Identities are concealed for this reason.

complexity of these spaces. Section 3 describes the kinds of error-correcting inferences researchers perform to build models of these systems, and how they do so by relying on the partial or incomplete models they have. In Sect. 4 we go beyond this account of the function of these inferences, to explore the particular material, theoretical and heuristic resources modelers rely upon to draw such inferences. We finish by drawing together some of the deeper implications of this analysis for our understanding of the role of models and heuristics in scientific discovery and the methodological decisions surrounding the use of mathematical models within the field.

2 Problem-Spaces in Computational Systems Biology

One central goal of systems biology is the construction of detailed (rather than heavily idealized) models of biological systems. These models it is hoped can capture and predict the dynamic behavior of complex systems involving hundreds of variables or more in ways experimental biology, which only keeps track of a few variables in any experiment, is incapable of doing (Kitano 2002; Westerhof and Kell 2007). Such models stand the best chance, it is hoped, of predicting accurately the effects of drugs while accounting for differences amongst individuals in drug response in terms of say the effects of parametric or structural variations (Hood et al. 2004). In many respects it is a brute force approach. The aim is to capture details (with some caveats, see Sect. 4) by relying on the capacities of modern computation, and use the faithful accounts of biological mechanisms which result, to achieve high levels of predictive fidelity. This approach may be contrasted with an approach which operates by idealizing and abstracting most details away in the hope of filtering out generic behavior and properties (Levins 1966). For many systems biologists, and indeed experimental biologists, such an abstraction-based approach is seen of limited use for producing predictively robust models at least. Biological systems are too variable and too complex.

That said, most current modeling cannot avoid levels of abstraction, idealization and simplification. While accurate models of the kinds of scale and detail systems biologists aspire to seem far off, the affordances of computation over the last 15 years has nonetheless opened up the possibility of producing models of much greater sophistication than has been possible historically for single modelers (modelers without super-computers) working at the desk. Such models represent systems using around 10–50 variables at rough approximation, and aim to capture at least the dominant dynamical behavior of such systems. Indeed computer development, perhaps more than anything else, has precipitated

the development of modern systems biology or computational systems biology as a distinct field, even though many of the conceptual and mathematical techniques the field relies upon stretch back to the first half of the twentieth century (see dynamical systems theory, cybernetics and related theoretical approaches). As a result there has been an influx of engineers, applied mathematicians, and other quantitative specialists into biology. For the most part systems biology is supposed to function collaboratively, whereby modelers work closely with experimental biologists, for advice on how to improve their models, but these relationships are often difficult and have yet to be institutionalized to any large degree. This leaves current modelers having to draw inferences about how to improve their models themselves, as we will see. Remarkably perhaps, even with limited biological support, modelers have proved capable of drawing inferences about missing elements of biological systems which were not anticipated or known to experimenters. That is, modelers have been capable of generating genuine cases of scientific discovery relying principally on their models alone. Such results serve as a powerful demonstration of the inferential power of mathematical models. But as well emphasize here such inferences are for the most part simply extensions of the ordinary day-to-day applications of partial or incomplete models to model-building.

There are a variety of practices and goals associated with modern system biology, and we describe here only a subset. This subset aims at building predictively reliable mathematical models of biochemical systems, such as gene-regulatory or metabolic systems. Such models are dynamic—tracking variables over time—such as regulation states of genes, or concentrations of metabolites, within an individual cell. Most modelers try to capture such systems using ordinary differential equations (ODEs), on the assumption that spatial effects with these cells do not impact heavily affect the dominant behaviors of these systems or the behaviors of specific interest. The mathematical models and tools of analysis and linearization would be familiar to anyone with knowledge of nonlinear dynamics.

With this information we can begin to describe what problem-solving consists of for modelers within this stream of systems biology. We will lean somewhat on Newell and Simon's characterization of problem-solving in terms of problem spaces (Newell and Simon 1972), although without necessarily signing up to a fully computationalist account of problem-solving (Danks and Ippoliti 2018). Modelers, often graduate students, are tasked at the outset with modeling a particular system, thought responsible for certain biological functions and dynamics behaviors. Their goal is to produce a model capable of capturing faithfully a set of those behaviors which also performs well under perturbation. This system is represented first visually as a network of interacting elements, commonly referred to as a “pathway” (see Fig. 1

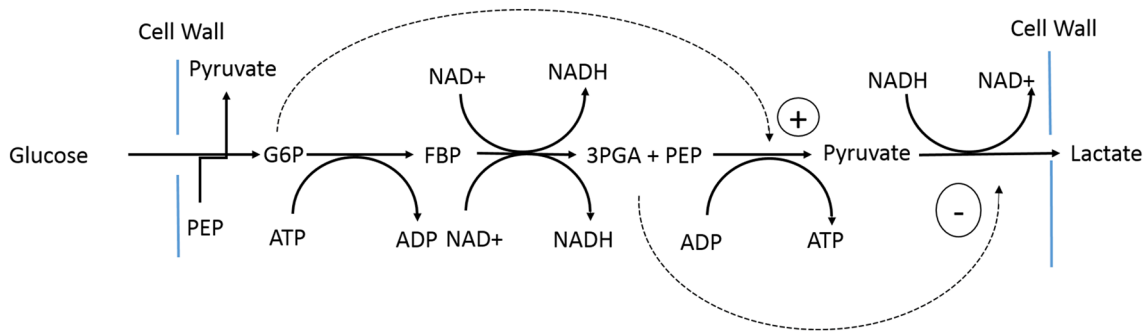


Fig. 1 A pathway model of glycolysis in yeast. This model was the initial model given to G16 (see below). The solid arrows represent main pathway interactions, with multiple elements involved in most interactions. The network has two aspects of complexity. There are feedforward effects, represented by dotted arrows. These are enzy-

matic activation (with a “+”) and inhibition (with a “-”) effects. Secondly multiple elements play multiple roles in the network (note: this representation is not chemically balanced). Reproduced with permission from Voit et al. (2006)

as an example). Nodes in the pathway represent biochemical elements. Arrows represent chemical interactions between those elements. Rates of change of say concentrations are termed fluxes, as biochemicals, through transformation, leave or enter a particular biochemical pool in a cell. For the most part these pathways are not in fact produced by experimental biologists, but need to be assembled by modelers themselves, through intensive study of experimental findings. Experimental biologists may assist in this process, but generally their own research tends to be confined to only pieces of that network, and the overall picture requires systems biologists to bring together results from multiple laboratories themselves.

Once a pathway representation is assembled modelers need to choose an appropriate mathematical framework for modeling the network. In metabolic system modeling, there are many options for representing the interactions between biochemical elements in the system. Michaelis–Menten kinematic representations for instance are often used, although models built from these are sometimes criticized for lacking mathematical analyzability. Other frameworks using more analyzable generic representations such as power laws are also applied (see Sect. 4). These can capture a vast range of behavior between reactants, although without mechanistic detail. In these ways the pathway is translated into a set of coupled ordinary differential equations in which the rates of change of concentrations are related to flux entering a pool and flux leaving. At this point parameters need to be fixed. Every term in the set of equations is modulated by at least two parameters, meaning that these models are generally parameter-intensive. Parameters are derived mostly from the same experimental literature, a process which introduces uncertainties on a number of fronts and has been criticized on a more fundamental mathematical level as ineffective for producing predictively robust or reliable models (Gutenkunst et al. 2007). Generally parameter measurements

of the kinds modelers need are only sparsely measured by experimenters (Westerhof and Kell 2007). Modelers are often forced to extrapolate parameters themselves indirectly from dynamic data, generalizing available measurements over populations, or by applying measurements from different but related systems or cells. All these techniques introduce various uncertainties. Unfitted parameters are then fit generally using a common global parameter optimization or fitting method such as simulated annealing. Some data on system behavior is usually held over from the parameter estimation process to help validate the model. In theory if the network structure is well represented and parameters representative of in vivo systems, the model should prove capable of capturing the network’s dynamics (which is not necessarily surprising given multiple degrees of parameter freedom) and also predicting accurately the response of the model to a variety of perturbations which were not part of the fitting process.

With these observations in mind we can begin to describe in more detail the nature of the problem and problem space modelers are confronted with. Firstly the actual structure of biological network being modeled is rarely accurate or complete enough for the level of accuracy the modelers are trying to achieve. Important elements are missing. The interactions might be modeled using poor formulations or poor parameter choices. Such models are incapable of reproducing the observable behavior of these systems accurately. Hence modelers generally start their research with an inaccurate representation of system structure.

As such one can characterize the problem-solving task facing modelers in Simon’s terms as a search process. Modelers need to search through a set of structural modifications (or manipulations; Magnani 2004a, b) to the original pathway structure they have. Such a space is constrained by the class of modifications that are at least within the realms of biological plausibility (and modelers often try to seek

information out from experiments in this regard). At the same time parameter sets need to be uncovered which fit the model to both training and test data. This adds a layer of complication insofar as a starting model described by a poor parameter set will not be much use for tracking down and debugging structural errors, particularly if parameters are critical to reproducing the nonlinear behaviors of the model. At the same time no amount of parameter searching will help improve a model that is structurally wrong in ways critical to its dynamics.

Further several other aspects contribute to the complexity of the problem solving tasks modelers face. Firstly, these biological systems are typically nonlinear. Nonlinearity manifests itself in two principal ways; through feedback and feedforward loops, and through elements playing multiple roles in a network. The pathway in Fig. 1 above exhibits feedforward effects and has several elements appearing in multiple roles. Both types of complexity typically makes life difficult for researchers trying to predict what causes what in a model, and what the consequence of a structural or parametric change to a model might be. Computation helps dramatically in this regard (see Sect. 4). Secondly the scale at which these systems are modeled, in order to be faithful to the goals of systems biology, tend to be large, requiring that a large set of the interacting elements are represented in the model, rather than abstracted out. The greater the number of interacting terms the harder it can be to isolate pathways in the model and consider them independently of the rest. Thirdly nonlinearity affects the complexity of the parameter spaces, complicating fitting processes. Nonlinear models generate suboptimal global maxima which fitting processes often find. The overall best fit can be very hard to find depending on the complexity of the system.

Factors such as these, and those mentioned before with respect to the difficulties of measuring parameters, mean that modelers face complex model-building tasks. What is remarkable perhaps is that modelers are nonetheless able to use models to draw inferences about their systems in order to improve their models. They do this relying to a large extent on the current partial or complete models they have.

3 Model-Based Inference in Systems Biology

Systems biologists use a collection of different *model-based inferences* to produce models capable of reproducing the dominant dynamical behavior of biochemical networks.²

² Note “model-based inference” is not a concept which readily occurs in the literature, even that on model-based reasoning, even if it explores inferences are derived directly from models. I take it as a sub-type of model-based reasoning.

These inferences help reiterate the importance false-models play in discovery contexts (see Wimsatt 2007; Nersessian 2008), particularly with respect to complex problem solving. Model-based inferences are, as we will see, the life-blood of scientific problem solving in these contexts.

Note at the start however that the term “model-based inference” is not a term which readily occurs in the literature, even that on model-based reasoning. Nersessian however identifies model-based reasoning essentially with inference (Nersessian 2008). In her terms model-based reasoning occurs whenever inferences are made by means of creating, manipulating, evaluating or adapting models (12). In that case there is really no distinction between the two, although in this paper we leave open the fact that model-based reasoning processes other than inferential, such as say explanatory or predictive processes, should also be captured by the phrase “model-based reasoning”, without a specific inference about a scientific phenomenon necessarily flowing from them. In the cases below modelers use models to make relatively discrete inferences about their representations of systems and in turn about the underlying reality of their systems. But there is a natural link nonetheless to what Magnani has discussed as manipulative abduction (Magnani 2004b), and it would be possible to analyze these inferences in those terms, given the extent to which models act as *epistemic mediators* or indeed through notions of distributed cognition (see also Chandrasekharan and Nersessian 2015). We do not pursue any of these deeper levels of philosophical or cognitive analysis in this paper however, preferring to focus on the details of the problem-solving practice.

In the context of modeling these complex systems modelers commonly need to draw at least four types of inferences.

1. Inferences to the location of structural and parametric errors in the model.
2. Inferences about which elements of structure or parameter adjustments could correct those errors.
3. Inferences identifying dominant causal dependencies in the model (and in turn the systems they represent).
4. Inferences about the regions in parameter space where optimal (and biologically plausible) parameter sets are likely to be found.

We will concentrate on the first two in this paper, as most relevant to themes we are discussing, and we will mainly confine ourselves to talking about inferences with respect to structure rather than parameters. The third type of inferences are nonetheless very important, since reducing systems to their dominant causal variables represents a common heuristic applied to simplify model-building tasks, at the cost of course, of giving up a certain degree of accuracy. In Sect. 4 we will see that such inferences help enable inferences of the first two types, by reducing the number of elements

required to build the model, and keeping parameter fitting tasks computationally manageable. However unlike the other first two types of inference, these are for the most part handled purely computationally, rather than in any strong sense, cognitively. Rather, statistical and mathematical procedures, such as types of sensitivity analysis, are used to derive the class of dominant variables. For instance one modeler we studied, G10, used a variance-based analysis in one model to correlate parameter variations of specific variables with effects on the output variable whose performance he was most interested in capturing. He used a Monte Carlo analysis over a large area of his parameter space to build up his statistical data. His analysis instructed him on which variables were most tightly related to model performance and variation. He isolated those. Likewise the fourth type of inference relies on mathematical methods and arguments, to restrict parameter space, but also very often biological arguments. Modelers cite biological factors on the pace at which various reactions can happen to restrict parameter search spaces.

The first two however interest us most here, as examples of the kinds of day-to-day model-based inferences which form the routine of model-building practices. Let us begin to explore then how these inferences commonly operate, and the essential role models play in them. In more formal cognitive terms systems biologists use partially correct models to constrain their task environments to a manageable set of search possibilities. This affords a reasonable chance of locating sources of errors, and inferring elements which will correct those errors, if the model is not too inaccurate. As mentioned the modeler's first task is, somewhat procedurally, to fit their basic pathway representation to standard canonical mathematical forms commonly applied in the field. These it is assumed will hopefully give them a model structure that captures the most important details of the system accurately enough that the model can capture some target behaviors well, at least qualitatively. Simplification processes search out dominant variables involved in these behaviors to simplify this initial model. Once in place and fitted, it is unlikely the model will capture the behavior of the system to the required degree. However with this model in place modelers can begin to explore the relationships between model variables, and more importantly how changes to model structure and parameter values, affect model performance.

The types of inference we have identified above come into play. Firstly internal constraints on the range of behaviors the model can exhibit guide a modeler to the identification of errors or weaknesses within the model. The initial model will have a range of dynamic behavior which can be explored over different initial condition and parameter variations, which can then be compared with data from the actual system. Differences provide clues to the location of problems and missing elements (and poor parameters),

which can be probed further, and narrowed down, by having the computer bring up graphical representations of different parts of the model and/or by having it simulate just parts of the model. Input elements can be “offlined”—set as a constant as per a classical debugging approach—to see if pathways relating to those elements are responsible for models performing poorly. It helps in this regard if experimental evidence is available in which certain variables are kept constant. We call these types of operations “explorations”. Running simulations repeatedly this way help modelers build up information on how certain variables relate to others. Often models refer to these practices in non-technical terms as getting “a feel for the model” (Voit et al. 2012a, b), insofar as explorations build up a modelers’ intuition into how their models operate, and in turn what result different changes might have (see modifications; next paragraph).

Such information will not necessarily completely localize an error or fault but can serve to constrain a set of hypotheses about where that error might be and what is in fact missing. For instance if there is too much flux along a particular channel or section of the model, then this should show up in the numbers, and will be responsible for a model not settling to equilibrium appropriately. Then modelers can begin to pose hypotheses about why flux is coming out too high and what might be wrong. Here a second feature of the role of the models (which we call here “modifications”) come into play; pre-existing models serves as a heuristic means by which modelers can define and structure a task environment for further locating and resolving errors. The model constrains the modifications modelers need to consider, if they trust the model is on the right track and close to reality. In particular a pre-existing model allows them to explore at the outset small localized changes to this model, and expand to more complex changes, as they build awareness of what such modifications do to the model, and in turn how the model operates. The effect of any change can be compared (often through comparisons of dynamical graphs) with experimental output to see if such changes are moving the model in the right direction. The set of potential changes in order somewhat of simplicity include the following, adjusting parameters; adding or removing arrows; adding or removing biochemical elements; changing the mathematical formula for representing given reactions; opening a black-boxed set of reactions (i.e. elaborating the system an arrow might represent); adding new elements outside the model. Hence an important heuristic guiding model-building is to work through those consecutively, starting with the simplest and least disruptive (see the description of G10's research below).

The first three changes can be worked through quickly and further help localize errors if it is possible to do so. If the set of elements currently in the model are well established experimentally then modelers can use this as a clue that other elements might be needed to modulate the

behavior of the original group. At the same time arrows—if interactions do not seem well described or cannot exhibit the range of behaviors required—can be de-black-boxed, and extrapolated in terms of their governing molecular interactions, to produce better representations. One useful aspect of de-black-boxing arrows in this way is that the behavior of the ensuing sub-module is constrained by the performance the original arrow is expected to have, which means the sub-module can be optimized somewhat independently of the rest of the system (Voit 2014). Generally if relevant elements stand in linear relations modifications to these relations can be experimented with first, and will yield more predictable outcomes.

All these aspects have the hallmarks of classical type engineering procedures. Most importantly for our purposes here, in all these cases the initial model provides the structure upon which search decisions are made. Modelers rarely deviate too far from the initial model. If these kinds of relatively small scale procedures fail to produce improvements, the expectation is that there is a fundamental problem with the original pathway model, the experimental data or the mathematical architecture used to represent it. In such cases the modeling task might simply be intractable given the information and data that exists.

We can see how these operations (explorations and modifications) play out in practice in the generation of inferences. In one relatively simple case for instance a modeler, G16, was charged with the responsibility of fixing a model of glycolysis in yeast (see Fig. 1) which failed to capture the right dynamic behavior of Fructose Biophosphate (FBP). The existing model failed to predict how quickly FBP levels declined given a certain glucose input to the system. G16 was lost at first until her PI suggested she “think about it mathematically” by experimenting on interactions upstream of FBP to see if changing the formulae representing those interactions might produce better behavior. This information helped constrain G16’s problem space for her, especially since only certain types of mathematical representation were both considered appropriate for biological systems and capable of dampening FBP production in the right way. Furthermore the model itself was not complicated by feedback relationships, and there were few elements upstream to consider, most of which had a linear causal input to FBP levels (see Fig. 1). The element PEP, which played a multiple role could be kept to last, since playing with its interaction would have produced complex, less predictable, outcomes. By experimenting with different interaction formulae she discovered that a so-called Hill interaction (which has dampening behavior) upstream could in fact produce a sufficiently accurate result.

G16 described her problem solving actions as follows. “I find glitches in the model, and why is it that, for example... And when you look at it there’s no way it can get better

because it depends on two things, and those two other things, for example, are increasing. So you can never get it decreasing for a period of time from those two. Maybe something else has a role that I haven’t taken into account.” G16’s comment captures in particular the way exploration operations work to discover underlying constraints in the existing model and locate potential errors, in this case a missing element which could modulate the effect of the two increasing elements. As described above G16 combined this information and the relatively linear structure of her model with the need for a new interaction formula upstream which could regulate FBP in the right way. This quote is illustrative here just how close her reasoning was structured and dependent upon the model before her eyes, and the use of constraints in the model to guide her thoughts about what is missing from accounts of the system. It also nicely expresses the fact that much of her practice for her is really just an extended process of debugging.

Another more complex example, which has been looked at in some detail, illustrates these kinds of operations on a broader scale. G10 was given the task of modeling lignin synthesis in two plant species (alfalfa and the poplar species; see Chandrasekharan and Nersessian 2015 for the details of this case). Lignin are polymers which help provide structure for plant cells, and often inhibit the use of cells to produce biofuels. The hope for a model of the lignin system is to learn how to control (and ultimately reduce) lignin production. G10 assembled a model from the literature which represented almost all research and knowledge about the lignin system up until that point in time. Most of the limited data he had however could only account for a system at steady-state. He needed a model that could predict what kinds of input could reduce lignin production, meaning, he needed model that could also model the effect of perturbations away from steady state. Hence, for this task, G10 faced a difficult set of conditions. As G10 discovered by running the initial model built using steady-state data, the model failed to predict how the system would fall back into equilibrium correctly (an exploration step), which provided his first clue that the flux in the model was not being regulated effectively. This led him to hypothesize through a sequence of modifications that interactions (arrows) and elements were both potentially missing from the model which could regulate the system back into equilibrium by removing excess flux.

His method was more or less that described above. G10 ran simulations to identify how the model operated, locating how and where flux in the model was accumulating in different channels. With this information he posed a range of hypothetical options for modulating fluxes within the network (see Fig. 2), starting with the simpler options, namely the insertion of new arrows. Of those that would help improve performance a certain collection were biologically plausible. G10 updated his model in this respect

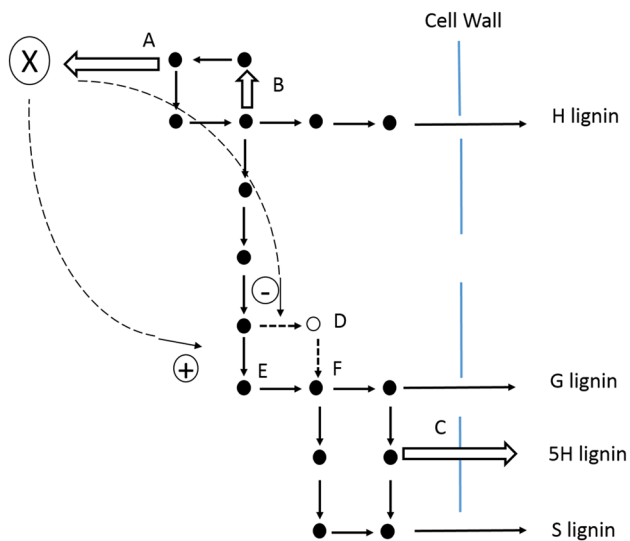


Fig. 2 A schematic representation of G10's pathway. Metabolic elements are represented as nodes for simplicity. In the first step G10 derived arrows A–C to help modulate the system back to equilibrium and handle extra flux. Element D needed to be added in since flow reaching F through E was under-calculated given experimental data. Initially flux through A was hypothesized to leave the cell. However G10 hypothesized a whole new element X was required to down and up-regulate the channels through elements D and F. Reproduced with permission from G10's PhD defense presentation. It is part of the NSF-funded ethnographic study “Becoming a 21st Century Scientist: Cognitive Practices, Identity Formation, and Learning in Integrative Systems Biology” and is reproduced with permission of the author

by including an interaction which had otherwise been left out, adding one new interaction (see D in Fig. 2) and two arrows directing flux out of the system. These moves fixed certain problems, but the model was still over-regulating and under-regulating production of two different types of lignin respectively. However he had exhausted the smaller scale modifications that were reasonable to make (i.e. new interactions and arrows). He had to consider more substantial types of revisions. As such he hypothesized there needed to be an entirely unknown element included in the network correcting regulating factors on both of these, and the only way to achieve this in a mathematically efficient way was to posit an element at one “loose end” of the model which served to regulate both pathways. In this way increases in flux previously leaving the system could be correlated in the right way with regulation of both these pathways. The resulting model performed extremely well. However the element X was not one which could be derived from the current biological literature. It was supported only through these model-based inferences, and indeed the overall effectiveness of model in capturing both the training and testing data. It was a significant discovery. Note that the inference of X was the result of cumulative changes rather than a single change, since the discovery X depended on D having been inserted in the

first place, thus creating two channels of flux in the model (see Fig. 2). This demonstrates how these model results are not one shot affairs but a process of accumulating insights and building up information. G10's model afforded this kind of cumulative step-by-step work, and allowed him to work towards a solution without needing to infer all the modifications required in one step.

In another case C9 was attempting to model the regulatory effect of reduction oxidation on the production of NF- κ B, a chemical important in cell apoptosis. Apoptosis often fails in cancer cells even after treatment from certain cancer cell killing drugs. These drugs trigger the release of reactive oxygen species (ROS). The NF- κ B transcription factor was thought relevant in this respect to the response of cancer cells to redox environments given its hypothetical sensitivity to ROS and its many major cell functions like proliferation and apoptosis.

So our working hypothesis has always been that, some cells are preferentially ...resistant to [cancer drug] doxorubicin because doxorubicin does something that leads to signal transduction with the cell that leads to, you know, anti-apoptotic transcriptions or something like that. And we know in the literature also that there are certain points in the NF- κ B pathway that are ROS you know, um regulated. So then it didn't take too much to say, ok if you have this drug that induces ROS it is a possibility that the ROS that's induced can affect this pathway within this cell that might lead it to be pro-survival, so that's sort of the idea.

Her hypothesis was that certain cells were insensitive to the cancer drug doxorubicin because in those cells the drug stimulated reduction oxidation processes which down-regulated NF- κ B. She had two cell lines, one insensitive (EU1) and one sensitive (EU3), and had previously built a model of redox regulation of NF- κ B. “... *our dynamic simulations highlight the interconnectedness of the redox enzyme network in controlling levels of reactive oxygen species such as H_2O_2 (hydrogen peroxide), and the necessity for inclusion of this network in future efforts to delineate reversible control of protein thiol oxidation.*”

However physical experiments in the lab she was performing herself, which applied doxorubicin to cells, were not in fact bearing out what the original model predicted. NF- κ B was expected to be more active or up-regulated in the case of EU1 cells insensitive to doxorubicin, and down-regulated in the case of the sensitive EU3 cells. This was not happening, suggesting ROS was being used up in some other manner. However her experiments showed that there were in fact more ROS in EU1 cells once doxorubicin was introduced. On the assumption that the original model governing the relation between ROS and NF- κ B was essentially correct, she hypothesized that in fact the surplus of ROS caused by

doxorubicin introduction might itself regulate the amount of active or toxic doxorubicin in the cellular system. She was helped in this regard by a discovery in the literature. “*So I found this one, really instrumental paper, that showed that, the doxorubicin toxicity in cells is something that’s mediated by the level of NADPH that is in the system.*”

As she put it then, “[instead of putting] *all these reactions into one single arrow and then just have an estimate of what the culmination of all of these reactions—would be—we realized that there were areas where there are differences between the EU1 and EU3 cells particularly with their[enzyme] NADHP.*”

C9’s move was to de-black-box a system thought to be diverting flux from the initial NF- κ B system. Her new model of the interactions between doxorubicin, NADPH and ROS, needed to reproduce the overall ROS measurements for both cell types, while at the same time, leaving amounts free that would trigger NF- κ B in the ways measured. This helped constrain this aspect of her model-building. Most importantly however her initial decision to build this sub-model flowed from both the confidence in her initial model, and her conclusions about the capacities of the mechanisms she described, to produce the nonlinear phenomena she was actually observing.

These cases all illustrate the centrality of model-based inference to instances of day-to-day scientific discovery. They illustrate a variety of the modification inferences; namely, interactions, missing elements and de-black-boxing. All their inferences were scaffolded by partially correct or complete models. Models guided discovery in the process of designing better performing models, by facilitating explorations which diagnosed and located errors, and modifications which resolved them. Of course initially these diagnoses of errors and modifications to models had more of the character of hypotheses, rather than inferences, in the sense that inferences are supposedly truth-bearing claims. Indeed initially no modeler likely believes in their truth. Plausibility in these hypotheses is built up through successful overall performance of the models that results, and through their biological plausibility. But nonetheless these hypotheses, and later inferences, were in a very meaningful sense “model-derived”.

4 Supporting Model-Based Inferences for Complex Systems

While such inferences are dependent on models, the structure of the models alone does not fully explain how modelers are able to draw them. As mentioned in Sect. 2 the modeling of biological systems is beset with challenging complexities and uncertainties. Indeed their successful use, despite the complexity of these systems, points to a richer underlying

account of practices in systems biology (and other computational fields), and how various other material, theoretical and heuristic resources are used to support such inferences. To some extent the field of systems biology has assembled various resources which enable these processes of model-building to succeed despite the complexity of the systems involved.

One of these essential resources is computation. Computation enables modelers to both understand how their models operate by manipulating aspects of the model, and in relatively straightforward ways, represent output visually. Computation enables modelers to experiment quickly with alternative hypotheses to see which ones might move the model in the right direction. Importantly for exploration purposes simulation helps modelers discover and get a grip on complex nonlinear behaviors in their models. It allows them to some extent at least bracket nonlinear behaviors. In general due to nonlinear behaviors examination of the mathematical equations alone may not be informative about how a model will behave and about which elements are causing what. Simulation helps modelers bracket domains of parameter and variable variation over which relatively consistent behaviors can be identified and worked with. This is important for modelers who need, to some extent at least, to anticipate through the model what effect a modification will have. It is clear then that computation facilitates the kinds of model-based inferences we have described above, even though the biological systems being modeled have numerous complexities. In fact this analysis suggests a richer story of what it is that computation brings to science, over and above just the ability to solve more complex equations. Computers augment researchers’ own inferential faculties and abilities, and thus increase their ability to handle more complex models, and more complex problems.

In the second place such inferences rely in essential ways on the accuracy of the initial model. Unless the model is within a relatively small number of steps of modification operations of the initial model, this process of model-guided exploration of alternatives and model-modification is unlikely to succeed. The more steps required the less likely an individual step will correlate with a robust improvement in the performance of the model, given the various internal dependencies between variables and parameters these nonlinear models have. Adding variable X may actually lead to a negative performance unless variable Y is also simultaneously added, and finding both simultaneously adds an entire order of difficulty to the discovery process. One of the insights then we can take from cases like these is just how important it is to have initial models which lie within the neighborhood of a robust solution, within the overarching problem space. Models which do not have this quality do not provide a good framework for model-building. In this

regard an “almost” accurate account of pathway structure is important at the outset.

Additionally processes of simplifying both the problem and the models needed to solve it, are necessary for facilitating debugging processes. Modelers often reduce their ambitions when it comes to modeling a system to aim at capturing robustly just particular variable relations, say the input and output variables that matter for their practical goals (e.g. controlling lignin production). This simplification of the representational problem they face, allows them to reduce the complexity of the models to just the dominant dynamical dependencies. These simplifications in turn reduce the scale of the parameter spaces which need to be grappled with, and also constrain models in ways which allow the necessary processes of debugging to function well. Simplifications also operate to some extent to find modularity in their systems (or nearly decomposable elements; Wimsatt 2007), by finding various sections of pathway which can be kept independent of one another. Some of G10’s reasoning for instance depended on being able to single out different flux pathways and improve them independently. This may come at the cost of system functionality. In MacLeod and Nersessian (2015) it was seen that simplification and abstraction procedures are often made at a cost to a system model. Such model cannot handle well anything other than the relations between its target input and output variables. The elasticity models have through the parameter-fitting process allows the discovery of models which fit these relations well, at the cost of finding one which describes all relations in the model well (MacLeod 2016). This problem does not necessarily though affect the ability of such models to reveal genuinely new structural information.

Canonical mathematical frameworks are very important in these regards (i) providing in-built simplification strategies, (ii) providing a reliable system for putting one in the neighborhood of a solution, and (iii) supporting heuristic processes of model modification. Such frameworks provide generic mathematical representations which can cover all types of interactions which occur within a biological network, such that modelers do not need to worry about the type of interaction or its underlying mechanism. Biochemical Systems Theory is one such framework (Savageau 1969). It postulates that all biochemical processes or interactions between two biochemicals, whose concentrations are represented as X and Y , can be modelled using power laws of the form $\alpha X^g Y^h$. In other words, these power laws can capture through appropriate choice of α , g and h the quantitative behavior of any interaction. The law itself is derived through a Taylor linearization of any non-linear form the interaction has. Such formulations are telescopic, meaning in the words of Voit, that “no matter what the level of modeling, the process descriptions are always in the same format. Thus, a complicated process at a low level may be approximated

with a power-law term and included into a higher-order model without destroying its mathematical structure” (Voit 2014, p. 260). Any interaction can easily be substituted out and replaced by a set of interactions (de-black-boxing) or in fact the reverse, replacing multiple interactions with just one. This facilitates some of the simplifications and modifications we have explored above. In general BST models are easy to build (although parameters of course can always be difficult to find), which is especially valuable in the context of systems whose governing processes or mechanisms are ill characterized. (3) The mathematical framework allows systems to be represented reasonably accurately at a “course scale”, which facilitates processes of building models out to improve their accuracy. Voit himself demonstrates his own processes of discovery and de-black-boxing through the use of these model-frameworks, which he calls “mesoscopic”, due to their level of abstractness. Although systems biologists do not use the kind of cognitive description of their practices we have applied, they are acutely aware of the need to provide a framework which assists debugging and the kinds model-building inferences modelers need to rely on.

As such various aspects of practice facilitate the kinds of model-based inferences modelers are required to perform. Such inferences do not happen in a vacuum, but require a framework of support. This framework of support gives us deeper insight into how inferences like these are afforded with respect to complex systems.

5 Conclusion—Models and Discovery

When modelers are often asked what the particular value of modeling is to biology, it is often not that modelers think that their modeling will produce predictively robust models anytime soon at least. Indeed such a belief is hard to sustain given the abstractions and uncertainties which go into these models (see MacLeod and Nersessian 2015). Rather it is that models can contribute insights into the structure of biological pathways inaccessible to traditional experimental methods. Systems biologists often account for the advantages of modeling using the concept of a “systems perspective” or “systems view”, which is somewhat vaguely defined as having an overview or understanding of an entire system, and thus presumably being able to see interactions and relations which occur over much larger scales. The analysis above however suggests that it is not an abstract form of understanding which modelers possess, but rather skills of model exploration and manipulation. They know how to debug models effectively to reveal unknown information about their systems. Models, through trial and error processes, and other modification heuristics, facilitate inferences about relations impossible without them. These inference are often local in the sense

that they do not rely on any broad insight or perspective on how the system works, but are based on manipulations of these models at specific points and through comparisons of specific aspects of them with the data. There need not be a global overview as a systems-view might suggest, although there is no doubt modelers understand their models very well after years of manipulating them.

As such these descriptions of debugging serve somewhat to demystify and reduce complex modeling to mundane but otherwise concrete practices. By doing so we reinforce a number of insights into the nature of scientific discovery. Firstly we reinforce Wimsatt's basic insight into the importance of false models, trial and error, and engineering style tinkering, in our accounts of how modelers build models. In fact the reasoning practices of systems biologists share so much with engineering styles of reasoning that it does challenge the notion that we should treat "engineering" and "science" in this context at least as cognitively distinct. However we reject the distinction here not because we think engineers reason more like scientists, but rather because modelers reason more like engineers in scientific settings.

Secondly, this case strongly rejects the philosophical claim that no rational procedures are operating in discovery contexts. In line with Ippoliti's "heuristic view" of scientific problem-solving, the cases above illustrate that there can be a definite set of rational procedures for discovery, even with respect to complex problems. As we have seen in our case there are more or less agreed upon sets of heuristics for framing the problem using different mathematical frameworks, plus sets of explorations and modifications modelers make are perfectly describable and contained to relatively limited set, as we have seen above. Further the structure of reasoning models the hypothesis-heuristic modification sequences described by Ippoliti (2018a, b). Each new model serves as a new hypothesis on the system. The only diversion from Ippoliti's views suggested by this case, is that there is a degree to which practices in this case do reinforce Simon's highly structured conceptualization of problem-solving and discovery. As such it may be premature to disregard Simon's views as a model of practice entirely, although admittedly Simon offers little regard to problem formulation and conceptualization. However in these cases of modeling in systems biology, the field has somewhat settled on canonical mathematical approaches for framing problems, and what the terms in the mathematical equations will represent. Most work is in fitting those frameworks, and there is a degree to which one might argue in fact that many of the processes we have describe above might be represented well as computational procedures, reminiscent of the BACON program. However whether this would be viable in the case of systems biology is not trivial and it would require further investigation to discover the extent to which human intuition and expertise do play an important role in the ability of

modelers to make the right moves (and employ heuristics fruitfully) in any given scenario (see MacLeod and Nersessian 2018). That said, the extent to which practices like these are repeated across fields in which problem formulations are somewhat pre-established, implies there may well be uses yet for Simon's more mechanistic notions of discovery.

But regardless of whether Simon's notions turn out fruitful in other contexts, the value "friends of discovery" perceive in being able to reveal the heuristics and strategies underlying discovery for scientific practice itself, is strongly reinforced in this case of modeling in systems biology. Most modelers entering systems biology are not aware of the intensity of modeling work, or the extent to which heuristics and trial and errors are the fundamentals of practice, as opposed to more theory-guided modeling work (MacLeod and Nersessian 2013). As such modeling is originally experienced as highly stressful. Modelers struggle at the start to see how they could possibly resolve the problems they are given. If modelers were trained better to recognize the procedural steps available to them—and the rational basis underlying them (which they are currently not trained to do), as well to understand the trial and error nature of how modeling must proceed in the field modeling, they should develop better expectations of the modeling process. There is no doubt an educational value is inherent in this work on discovery in systems biology, which supports overall conjectures on the value of exploring discovery processes for science itself.

In addition to this there are perhaps two important conclusions to be had from this analysis as far as philosophy of science and the importance of discovery processes are concerned.

Firstly, models have a richer life in scientific reasoning beyond their ability to produce explanations and predictions. Cases like these from systems biology reiterate the importance of models to actual model-building—and processes of genuine discovery—over and above the traditional epistemic roles of models philosophers usually investigate. Such traditional investigations typically concentrate on models after they have been constructed, as "finished products". Our assertion here is that model-based inferences are essential to the capacity of modelers to handle the complexity of their modeling tasks, and are intimately tied up from the very beginning with how researchers think about problems and infer properties of their systems. Indeed in practice such models are rarely finished products, always inadequate, always in need of further refinement and development.

Secondly, if a goal of philosophy of science is to understand methodological choices, then we need to pay more attention to discovery processes in order to build the richest explanations we can of why fields use some techniques and tools rather than others. As we have seen the model-based inferences drawn in systems biology rely not just

on the ability of models to constrain problem spaces, and provide task environments, but also on both the use of modern accessible computation to facilitate and examine hypothetical modifications to models, and the choice of modeling frameworks which can simplify modeling tasks and afford the kinds of tasks modelers need to perform. Modelers are explicit when describing how their methodological choices are designed with their capacities for model debugging in mind. In such cases methodological choices, like the use of well-designed canonical mathematical frameworks, are not simply a matter of representational accuracy, or explanatory and predictive value, but pertain to the constraints of problem-solving processes and the needs of inferential practices. There is thus important value, when it comes to understanding methodological choices and decisions, in seeing the function of particular methods and tools from the perspective of the researcher facing a problem-solving task. This view might be opaque to us if we just take perspective of the scientific article writer, who is only keen to justify and rationalize a final result.

Funding This study was funded by a NSF grant, Directorate for Education and Human Resources (Grant No.: DRL097394084).

Compliance with Ethical Standards

Conflict of interest I (Miles MacLeod) declare I have no conflict of interest.

Ethical Approval I (Miles MacLeod) can confirm all procedures performed in studies involving human participants were in accordance with the ethical standards of the institutional review boards (IRB) and with the 1964 Helsinki declaration and its later amendments or comparable ethical standards.

Informed Consent Informed consent was obtained from all individual participants included in the study.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- Chandrasekharan S, Nersessian NJ (2015) Building cognition: the construction of external representations for discovery. *Cogn Sci* 39:1727–1763
- Danks D, Ippoliti E (eds) (2018) Building theories: heuristics and hypotheses in sciences. Springer, New York
- Gutenkunst RN, Waterfall JJ, Casey FP, Brown KS, Myers CR, Sethna JP (2007) Universally sloppy parameter sensitivities in systems biology models. *PLoS Comput Biol* 3(10):e189
- Hood L, Heath JR, Phelps ME, Lin B (2004) Systems biology and new technologies enable predictive and preventative medicine. *Science* 306(5696):640–643
- Ippoliti E (2017) Scientific discovery reloaded. *Topoi*. <https://doi.org/10.1007/s11245-017-9531-3>
- Ippoliti E (2018a) Building theories: the heuristic way. In: Danks D, Ippoliti E (eds) Building theories: heuristics and hypotheses in sciences. Springer, New York, pp 3–20
- Ippoliti E (2018b) Heuristic logic. A kernel. In: Danks D, Ippoliti E (eds) Building theories: heuristics and hypotheses in sciences. Springer, New York, pp 191–212
- Kitano H (2002) Looking beyond the details: a rise in system-oriented approaches in genetics and molecular biology. *Curr Genet* 41(1):1–10
- Levins R (1966) The strategy of model building in population biology. *Am Sci* 54(4):421–431
- MacLeod M (2016) Heuristic approaches to models and modeling in systems biology. *Biol Philos* 31(3):353–372
- MacLeod M, Nersessian NJ (2013) Building simulations from the ground-up: modeling and theory in systems biology. *Philos Sci* 80(4):533–556
- MacLeod M, Nersessian NJ (2015) Modeling systems-level dynamics: understanding without mechanistic explanation in integrative systems biology. *Stud Hist Philos Sci C* 49(1):1–11
- MacLeod M, Nersessian NJ (2018) Cognitive constraints, complexity and model-building. *Hist Philos Life Sci* 40:17
- Magnani L (2004a) Model-based and manipulative abduction in science. *Found Sci* 9(3):219–247
- Magnani L (2004b) Conjectures and manipulations. *Computational modeling and the extra-theoretical dimension of scientific discovery*. *Mind Mach* 14(4):507–538
- Nersessian NJ (2008) Creating scientific concepts. MIT Press, Cambridge
- Newell A, Simon HA (1972) Human problem solving, vol 104, No. 9. Prentice-Hall, Englewood Cliffs
- Savageau MA (1969) Biochemical systems analysis: I. Some mathematical properties of the rate law for the component enzymatic reactions. *J Theor Biol* 25(3):365–369
- Voit EO (2014) Mesoscopic modeling as a starting point for computational analyses of cystic fibrosis as a systemic disease. *Biochim Biophys Acta* 1844(1):258–270
- Voit EO, Almeida J, Marino S, Lall R, Goel G, Neves AR, Santos H (2006) Regulation of glycolysis in *Lactococcus lactis*: an unfinished systems biological case study. *IEEE Proc Syst Biol* 153(4):286–298
- Voit EO, Qi Z, Kikuchi S (2012a) Mesoscopic models of neurotransmission as intermediates between disease simulators and tools for discovering design principles. *Pharmacopsychiatry* 45(1):22
- Voit EO, Newstetter WC, Kemp ML (2012b) A feel for systems. *Mol Syst Biol* 8(1):609
- Westerhoff HV, Kell DB (2007) The methodologies of systems biology. In: Boogerd F, Bruggeman FJ, Jan-Hendrik SH, Westerhoff HV (eds) Systems biology: philosophical foundations. Elsevier, Amsterdam, pp 23–70
- Wimsatt WC (2007) Re-engineering philosophy for limited beings: piecewise approximations to reality. Harvard University Press, Cambridge