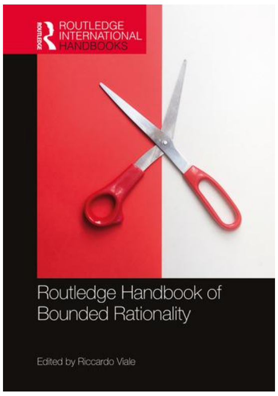

Routledge Handbook of Bounded Rationality

Riccardo Viale

Bounded rationality, distributed cognition, and the computational modeling of complex systems

Publication details
https://test.routledgehandbooks.com/doi/10.4324/9781315658353-7
Miles MacLeod, Nancy J. Nersessian
Published online on: 29 Oct 2020

6

BOUNDED RATIONALITY, DISTRIBUTED COGNITION, AND THE COMPUTATIONAL MODELING OF COMPLEX SYSTEMS

Miles MacLeod and Nancy J. Nersessian

Introduction

Computational modeling and simulation of complex systems are playing a major role in twenty-first-century scientific discovery. Given the rhetoric of "big data" and automated modeling practices, one might expect that computer simulation is a resource for side-stepping human cognitive capacities and constraints. Yet even these practices still require human cognitive engagement with their products for the purposes of validation and interpretation. And, as we have discovered in our investigations of computational systems biology, much problem-solving still involves researchers building their own models in situations of scarce or inadequate data, using laptop computers. Rather than freeing scientists from their cognitive limitations, using computation to extend control and understanding over ever more complex systems requires developing sophisticated problem-solving practices which push them toward these constraints. Several philosophers of science have had the intuition that cognitive capacities and limitations play a significant role in shaping methodological strategies and choices, especially in modeling complex systems (see, for instance, Humphreys 2009). It would seem, then, that cognitive considerations should be factored in with epistemic and practical considerations in advancing justifications for specific methodological choices in such problem-solving processes, but philosophers by and large have not examined these. Unlike most of the research in this area of philosophy, which relies on published papers concerning finished models, our research is based on the premise that it is necessary to examine the model-building practices to develop an understanding of how human cognitive capacities enable and constrain these, as well as the ways and extent to which modelers overcome cognitive limitations. Our analysis is based on a 5-year ethnographic study of model-building practices in computational systems biology. In our investigations of innovative modeling practices in systems biology, we have found the notion of bounded rationality and its mechanisms of search and satisficing particularly useful for explaining and justifying the development of modeling methods and strategies. The overarching goal of contemporary systems biology is to create high-fidelity computational models and simulations of complex biological systems. This goal largely eludes them. Instead modelers rely on strategies of building mid-level models that contain modest details of system composition and organization ("mesoscopic models") and of "building-out" rough initial models and incrementally exploring and modifying these. Thus, modelers build constrained problem spaces, and solution search operates chiefly through generating inferences about erroneous model structure, which requires both a set of search heuristics and constant feedback through computational simulation. The use of rather simplified representations, which are often far from high-fidelity, can be seen on cognitive and other grounds to be rational for the long-term goals of systems biologists, given cognitive and other constraints. This is important, since most models of biological systems currently produced in the field do not achieve the level of fidelity required for prediction and control for medical or other purposes. Nonetheless we will show that the modeling strategies currently employed provide a bounded rational solution to the problem of building such models over the long term. However, ours is a quite different analysis of scientific problem-solving than that advanced by Simon. Specifically, we have been advancing the position that scientific problem-solving requires a distributed cognition analysis. In the case at hand, accounting for scientific discoveries through building computational representations requires a distributed cognition analysis of the *processes* of building these representations, in which the model and modeler comprise a complex distributed model-based reasoning system. Thus, we extend the notion of bounded rationality to distributed cognitive systems. Although not entertained by Simon himself, we see the analytical framework of distributed cognition as a means of further explicating his "parable of the ant," where ant and environment comprise a complex problem-solving system (Simon 1996, p. 63). In the parable, if one simply traces the track of the ant over a two-dimensional space, its problem-solving behavior seems highly complicated. However, once the environment is factored in, each move of the ant can be recognized as a relatively simple localized response to individual obstacles. The ant's behavior is largely a reflection of constraints in the environment and basic internal representations. Simon suggested that "man" could be substituted for "ant." However, while we agree problem-solvers do make local responses based on limited information, we would add that much human problem-solving behavior is not just a *response* to an environment but a more *active use* of affordances in the environment to aid cognitive tasks. Further, humans create affordances in the environments such as external representational artifacts that "off-load" cognitive tasks to the environment, as in the case of the speed bug performing memory functions for the pilot in the process of landing a plane (Hutchins 1995a). In this brief chapter: (1) we outline the modeling tasks in systems biology; (2) we examine their modeling practices; (3) we present an overview of our cognitive analysis; and (4) make a case for their bounded rationality.

Modeling tasks in systems biology

Our ethnographic study of model-building practices tracked modeling in two computational systems biology labs. In each lab the aim is to construct relatively detailed models of metabolic, cell signaling, and gene-regulatory networks. These models are dynamic, tracking the changes in concentration of each metabolite, or the activity of each gene, in a network over time. Philosophically, computational systems biology takes the stance that only high-fidelity mathematical models can capture network behavior sufficiently to gain understanding and reliable control of these networks for medical or technological purposes. However, although the basic mathematical methods of systems biology have a long history, the potential to produce high-fidelity models has only become possible over the last 20 years with the advent of readily available fast computation. It is now possible for a modeler using a laptop computer to create a model which captures 50 or more interacting biochemical elements. For the most part, modelers aim to model these biochemical networks using coupled ordinary differential equations. These networks, however, have a number of complexities which make the modeling task difficult. Structurally, the biochemical networks are highly nonlinear, exhibiting feedback effects, and having elements playing multiple roles at different points in a network (meaning that different points in the network draw on the same chemical reservoir within a cell). Cognitively, non-linearity makes it difficult for modelers to track causal relations in their model and anticipate what modifications will achieve. Additionally, modelers never have a complete biological representation of their network ("pathway diagram") at the beginning of their project and lack sufficient data to build a highly detailed model. Indeed, since experimentalists generally do not study all interactions within a network themselves, modelers (who know little biology) need to search through the experimental literature and databases to build the pathways on their own.

With this information in hand, we can describe the modeling task. First, modelers need to assemble a pathway representation of their network, which is a line diagram connecting elements in a chain according to the order in which they interact. The pathway representation allows them to map variables to a mathematical formulation and choose interaction representations to fill in that formulation. Finally, parameters to fit the model need to be estimated from the available data. Any unknown parameters can then be determined through a parameter-fitting algorithm which chooses a parameter set for the model, according to how well it can reproduce experimental data on the network behavior. Usually some data are held over for testing the validity of the model that is produced. If all this works, then modelers can proceed to make novel predictions about system behavior, such as how to intervene optimally on the systems for medical purposes using target drug combinations.

Although this looks like a relatively straightforward set of procedures, in nearly all cases we have examined, the starting biological pathway representations modelers assemble are never adequate for the network behaviors they want to capture. Biochemical elements are missing or interactions among elements in the network are missing. At the same time, initial representations of those interactions might not be correct. The models they produce thus invariably fail to test well against the available data first time around. The task, then, of the modeler is to correct the models they have produced by deriving or inferring the location of missing elements or poor interaction descriptions in their models and hypothesizing additions or corrections to replace them. Usually a variety of corrections are required to build a model that functions adequately.

Mesoscopic modeling and the building-out strategies

Given the complexity of the systems being modeled, modelers have developed strategies that make the task cognitively tractable. In the first place, they build a lower-fidelity model that bounds the problem space. Such models are called "mesoscopic" (Voit et al. 2012), meaning they are "in between" any one level of description or biological organization in two senses. First, the systems-level is simplified such that only certain relationships or system-level behaviors are represented, rather than the full range of potential behaviors. Second, the network itself is simplified in a variety of ways, often relative to these system-level goals. Quite abstract non-mechanistic representations of biochemical interactions are applied to simplify running simulations, parameter-fitting, and mathematical analysis. Elements thought not to be relevant to target behaviour are left out. Subsystems are typically blackboxed and treated as linear or constant. Note that while such simplifications are typical of any modeling effort, they are problematic to the extent that systems biology aims at high-fidelity models capable of robust predictions.

The mesoscopic model delimits the problem space for searching for solutions to correct the model. Dimensions of the space represent different structural alternatives and biochemical interaction alternatives (in terms of their mathematical description). This space is also constrained by what interactions and structures are biochemically possible, although a modeler (typically from engineering) might not know all of these. The task of the modeler is to locate the alternatives which best capture the functional performance of the target network they wish to capture and also perform best in response to perturbations.[1] It is impossible for modelers to simply hypothesize the needed corrections in one hit. These more limited models still represent many interacting elements, involving many non-linearities. Although it might be possible to predict what one modification might do, understanding the implications of changing two elements is usually impossible. There is thus an imperative on modelers to develop strategies for iteratively localizing problems and fixing them. The strategy our modelers employ has been developed around a central presupposition, namely, that the initial models they formulate are within a relatively small set of changes from a good network representation. This is essential, since for the most part their strategy is constructed almost entirely around making discrete modifications to the original model and using the original model to guide their actions. We call this a "building-out" strategy. In practice, the problem space modelers confront is constructed relative to this initial model. In addition to this initial model representation, model simulation provides information as to how the dynamics of the model differ in behavior from the target. Modelers are thus, at the start, looking for modifications which bring their model into alignment with that behavior. Of course, the nature of the divergences provides clues as to where to look and what to change. With an initial model at hand, then modelers formulate their strategy around two central tasks: *explorations* and *modifications* (MacLeod 2018). Explorations are steps modelers employ to try to understand the causal relationships within the model. Explorations can be done through pen and paper representations, but usually they also require running simulations repeatedly and visualizing the results. To get the information they need, modelers will often temporarily remove elements to try to isolate particular relations. Importantly, they can use the affordances of computation to bracket nonlinear behaviors, through explorations of different parameter and input variations. These repeated processes help modelers develop what they often refer to as having a "feel for a model" (Voit et al. 2012). We have analyzed this notion in detail (Chandrasekharan and Nersessian 2015; MacLeod and Nersessian 2018) and understand it to refer to intuitions modelers build up on how particular variables and parameters in the network causally affect specific others. In more cognitive terms we characterize these intuitions as simulative mental models; internal representations which simulate causal relationships.

Explorations build up families of mental models which themselves represent at least part of the structure of the overall model. These mental models help modelers predict potentially fruitful modifications by allowing them to infer possible effects of modifications. Modifications are hypotheses modelers make which should help move the model toward a sound representation. Working from the original model, modelers can entertain a sequence of increasingly complex heuristics, including: (1) adjusting parameters; (2) adding or removing interactions between elements already in the network; (3) changing interaction formulae; (4) "de-black-boxing" a set of interactions; and (5) adding a new element to the model at different points in the network. This list of options represents discrete, localized changes to models. Each can be hypothesized and then tested according to how well the new model performs and, if heading in the right direction, new modifications can then be hypothesized. In this respect these heuristics follow a hill-climbing logic, with modelers hoping for improvements with each move but not expecting to find any fully definite solution in one hit. Hopefully within a relatively contained sequence of steps the model reaches the performance goals of the modeler. If the model fails to improve, then modelers will likely conclude that there is no cognitively accessible path from their original model to an adequate representation. The model might be incorrect in too many interrelated ways for them to untangle.

We view this building-out strategy as a means by which modelers of complex systems turn problem-solving situations with ill-defined problems, unstructured task environments, and quite general goals into a series of situations approximating those Simon envisaged. For each iteration of the model-building process, there is a relatively well-defined problem space: the modification process has a fixed representation, the goal of "fitting" that model, and a set of heuristics for facilitating what is in essence a search process. However, to understand how such computational model-building processes lead to scientific discoveries we need to embed the notions of problem space, search, and bounded rationality a different kind of cognitive analysis than that advanced by Simon.

Cognitive analysis: distributed model-based reasoning

Computational model-building provides an exemplar of what Hutchins calls creating cognitive powers (1995b; also see Chandrasekharan and Nersessian 2015). Numerous iterations of model-building (including visualizations of behavior), debugging, modifying, and simulation, build the model and modeler into a "coupled" cognitive system. The highly integrated nature of this distributed problem-solving system results in the modeler being able to manipulate and explore a much more complex set of equations that would be possible "in the head" alone—even with pen and paper resources. Scientific problem solving has always been dependent on creating and using external representations, especially various kinds of diagrams. There is a vast literature that examines the cognitive roles of these in both mundane and scientific reasoning. The roles of computational representations, which dominate contemporary science, however, have received scant attention, and what literature there is focuses on the function of model visualizations in mental modeling (Trafton, Trickett, and Mintz 2005; Trickett and Trafton 2002, 2006, 2007) and distributed problem-solving processes (Alač and Hutchins 2004; Becvar et al. 2008). The work of our research group has taken a different approach, focusing on the processes of *building* computational representations and how discoveries emerge from these processes (see, especially, Chandrasekharen and Nersessian 2015). Our analysis starts from Nersessian's account of scientific problem-solving processes in terms of "simulative model-based reasoning" where inferences are drawn in distributed cognitive processes, comprising building and manipulating mental models and external representations such as diagrams and equations (Nersessian 1992, 2002, 2008) or physical models (Nersessian et al. 2003; Nersessian 2009). Inferences are drawn through processing information in the system comprising memory ("internal") and the environment ("external")—what we call *distributed model-based reasoning*—rather than as on the traditional account where all information is abstracted from the external representations (environment) and represented and processed internally, as with Simon's ant (see Greeno 1989, for an earlier similar claim). We argue that the process of building models integrates manipulations in the mental model and in the external representation, creating a *coupled cognitive system* (Osbeck and Nersessian 2006; Nersessian 2008). Here we briefly discuss the extensions required of the framework of distributed cognition to accommodate our account of scientific problem solving, especially as it pertains to computational representations (a detailed account is provided in Chandrasekharan and Nersessian 2015). Much of the research on the function of external representations within the distributed cognition framework has focused on highly defined task environments with well-defined problems and goals, such as piloting ships and planes (Hutchins 1995b) or solving puzzles with fixed representations and rules (Zhang and

Norman 1994). These analyses focus on the way external representations change the nature of the cognitive tasks, especially by reducing cognitive load through off-loading memory to the environment. These external representations are ready-to-hand. Scant attention has been directed toward understanding the role of processes of building the external representations to alter task environments in cognitive processes during problem solving (exceptions include Kirsh 1996 and Hall et al. 2010). Scientific practices provide a prime locus for studying processes of building external representations since building problem-solving environments is a major component of scientific research (Nersessian et al. 2003; Nersessian 2012). When studying science, the focus needs to shift to the building of external structures consonant with how people actively distribute cognition (Hall et al. 2010) through creating problem-solving environments and in turn building cognitive powers (Hutchins 1995b). A research laboratory provides a good example of a problem-solving environment. It contains artifacts salient to the research, people, and socio-cultural practices. A computational model, itself, provides another example of a problem-solving environment. Through building and running the model, the researcher creates cognitive powers that extend her natural ones, such as the ability to synthesize a vast range of data, to visualize complex dynamical processes, and to run through unlimited imaginative (counterfactual) scenarios not possible in her mind alone. Indeed, we have witnessed how building computational models enables systems biology modelers with scant biological knowledge to make fundamental biological discoveries (Chandrasekharan and Nersessian 2015). Our extension of the framework to science has been based on ethnographic studies of the open, ill-formed problem solving tasks of pioneering research and of the processes by which researchers create their cognitive artifacts; in the case at hand, computational representations. The other way is which the framework of distributed cognition is in need of extension is by providing an account of the nature of the internal representations used by the human component of the system. Most research on distributed cognitive processes is silent about mental representations. Zhang and Norman (1994), which explicitly analyzes the interactions among external and internal representations in problem solving, assume the internal representations to be mental models. Nersessian (1992, 2002, 2008) has elaborated a "mental modeling framework," which she argues provides a cognitive basis for the range of model-based reasoning used in scientific practice. This framework draws from the strand of research in the mental modeling literature that examines the processes of constructing and manipulating a working memory model during reasoning and problem solving and is agnostic about the nature of the long-term memory representation. In thinking about scientific reasoning, Nersessian (2008) argues that we need to move beyond the mental models literature per se and create a synthesis of an extensive range of experimental literature spanning over 25 years on discourse and situation modeling (see, e.g., Johnson-Laird 1983; Perrrig and Kintsch 1985; Zwaan 1999), mental animation (see, e.g., Hegarty 1992; Schwartz 1995; Schwartz and Black 1996), mental spatial simulation (see, e.g., Shepard and Cooper 1982; Finke 1989; Kosslyn 1980), and embodied mental representation and perceptual simulation (see, e.g., Glenberg 1997; Bryant and Tversky 1999; Barsalou 1999; Brass, Bekkering, and Prinz 2001). She advances a "minimalist hypothesis" that "in certain problem-solving situations humans reason by constructing a mental model of the situation, events, and processes in working memory that in dynamic cases can be manipulated by simulation" (2002, p. 143). Recent psychological research on scientists and engineers as they try to solve research problems lends support to the hypothesis and further details the nature and role of mental model (or "conceptual") simulations (Trafton, Trickett, and Mintz 2005; Trickett and Trafton 2007; Christensen and Schunn 2009). For instance, Trickett and Trafton (2007) establish the importance of conceptual simulation in the data analysis phase of several different kinds of scientific fields. Most importantly, the use of such simulations increases in cases of inferential uncertainty where scientists appear to be trying to develop a general understanding of the system under investigation. They see their findings as lending support to the extension of the "minimalist hypothesis" to scientific problem-solving. Mental modeling is often carried out in the presence of real-world resources, such as diagrams and the various objects being reasoned about, as Simon has noted for design and other processes (Simon 1996). This is even more so in the case of science, where visual representations, physical models, and computational models are integral to the reasoning in problem-solving processes. Further, the representations and processing required for such sophisticated reasoning are too detailed and complex to be "in the head" alone. In cognitive psychology, the question of the interface between the mental capacity and resources in the world has largely focused on diagrams and other visual representations. Several of these analyses have proposed that the external representations are *coupled* with the mental representations in inferential processing (see, e.g., Greeno 1989; Zhang and Norman 1994; Gorman 1997; Hegarty 2005; Nersessian 2008). Our research extends internal and external representational coupling to dynamical representations: computational simulations and mental simulations. By examining the process through which external and internal representational structures are built, we can begin to understand how these are incorporated into a distributed system that performs cognitive functions such as memory and inference; in the case at hand, distributed model-based reasoning through the coupling of mental modeling and computational simulation processes. From our interview data and the pen and paper sketches the computational modelers in our study make, we have been able to discern some features of the mental representations they use in "debugging" computational models (MacLeod and Nersessian 2018). Broadly, they can be described as simulative models of causal relationships represented in their mathematical networks, which operate in a manner similar to the kinds of mental models captured in investigations of basic causal network reasoning (Hegarty 1992, 2004; Schwartz 1995). These models are qualitative (Roschelle and Greeno 1987), rather than quantitative, and the inferences drawn from them are thus qualitative in nature. Further, as with everyday causal network reasoning, these models tend to be selective and piecemeal representations of the overall systems (Hergarty 1992). In particular, nonlinear relations can be identified and bracketed with the aid of computation into different separate behaviors, allowing mental simulation of each separately. For example, the range over which a feedback structure produces stable oscillations can be separated from the range over which this structure produces more equilibrium-type behavior, and these can handled separately. Incrementally building such mental models in conjuction with processes of computational simulation and visualization creates the coupled cognitive system through which novel inferences about both the computational and real-world biological systems can be made, often leading to important discoveries (Chandrasekharan and Nersessian 2015; MacLeod and Nersessian 2018).

The bounded rationality of model-building practices in systems biology

We think it especially important to consider the rationality of modeling practices in systems biology because the field currently is failing to reach its stated epistemological goals. The practice of mesoscopic model-building and incremental modification through building-out strategies provides an exemplar of what Simon, in a later analysis of bounded rationality, called "procedural rationality" (Simon 1976), that is, the rationality of the processes through which boundedly rational decisions or problem solutions (substantive rationality) are achieved.

To show procedural rationality for a particular strategy, two claims need to be made: first, that a cognitive system is constrained (or bounded), and, second, that the strategy provides an effective pathway given these constraints, toward reaching a set of goals. Building mesoscopic models bounds the problem space of the complex biological system. This facilitates the effectiveness of the building-out strategy as there is a limit to both the complexity of the model (in terms, e.g., of the number of elements reponsible for any particular behavior) and the number of faults which explain why a model departs from the data. Any model which breaches such limits would not be tractable for mental modeling in the ways we have suggested, restricting the possibility of employing the heuristic steps fruitfully. As we have seen, the nature of the building-out strategy, to make discrete incremental modifications to a pre-existing model, suggests the extent to which modelers work to keep their operations within their cognitive limitations. Mental modeling research indicates that among other things there is a critical need to keep modeling practices within the scope of working memory. Inferences are drawn based on causal interactions among elements. As Hegarty has shown in the case of pulley systems, the ability to reason about causal networks is constrained directly by working memory (Hegarty 1992, 2004). The building-out strategy depends on being able to draw inferences as to the sources of errors and about modifications which might eliminate them, which requires starting from a more limited computational model representation. Given these constraints, one can make a good argument that the strategies of systems biology modelers provide a boundedly rational option for their model-building tasks. Producing even a mesoscopic model which captures the systems adequately at the outset is impossible due to incomplete information and the complexity of the systems. The building-out process affords error-correcting inferences. A step-wise approach takes advantage of this while keeping the modeling processes within cognitive control. Modelers resolve the model not by planning several moves ahead, but through selecting among a set of choices which, at any one time, amount to local fixes or improvements. Overall, the rationality of the building-out strategy depends on both initial models being within a relatively small number of steps from a good outcome and on the complexity of interactions within the model being constrained, such that modelers can effectively draw inferences through engagement with computational simulation. One cannot judge *a priori* how well initial models meet the conditions they need to have in order for the building-out strategy to pay-off. But, in our experience, modelers do manage to produce models through these techniques which meet their limited data-matching and prediction goals. And even though the models are not high-fidelity, some have produced significant discoveries, such as discovering unknown elements in a well-established pathway (Chandraskeharan and Nersessian 2015; MacLeod and Nersessian 2018). Such achievements are in fact a powerful endorsement of the power of this approach to modeling of complex biological systems, despite the ultimate goals of the field. In general, we think our research shows that bounded rationality can provide not only a good description of problem-solving practices, but also a justification for methodological decision-making that, on the surface, might seem sub-optimal in fields that are attempting to understand, predict, and control complex systems by relying heavily on computational modeling and simulation. As noted, although the strategies considered enable the production of models which meet limited representational goals, such models do not achieve the scale and fidelity at which systems biology aims. A major part of the rhetoric of systems biology is that modern computation opens up the possibility of building high-fidelity models of large-scale systems, which are required for robust predictions and to advance biological theory. However, as Voit et al. (2012, p. 23) describe it, the majority of models in systems biology are not ”large enough to approach the reality of cell or disease processes with high fidelity.” The strategies modelers use to build these limited, mid-sized models, seem, on the face of it, to push modelers away from these goals, especially since they are rarely capable of robust predictions (MacLeod and Nersessian, 2015, 2018). This has created a challenge for the field of justifying their current practices, and systems biologists are remarkably reflective on the nature of the cogntive challenges they face and on the need to rationalize the current production of relatively small-scale models.As pioneers in the field, Voit et al. (2012) and Voit (2014) argue that the value of such models to the field is that, once built, they create tractable cognitive platforms for richer model development and provide insight into the underlying systems, by virtue of their mesoscopic nature. That is, they create potential pathways for scaling models up in detail and complexity, such that higher fidelity models can be achieved sometime in the future. Interestingly, Voit et al. locate the cognitive basis of

> [their] strategy of locally increasing granularity ... in semantic networks of learning and the way humans acquire complex knowledge ... hierarchical learning is very effective, because we are able to start simple and add information as we are capable of grasping it.
>
> *2012, p. 23*

In other words, mesoscopic models, once constructed, serve as learning platforms, facilitating a modeler's ability to build richer, more detailed models that capture more of the biological system's complexity. Hence, although current systems biology might be failing in its aim of constructing high fidelity models, the kinds of models being built are anticipated to facilitate the ability of modelers to produce more realistic models without being overwhelmed by their complexity. Detail can be added incrementally, in controlled steps using heuristics modelers have been developing. For instance, once a mesoscopic model is in place, modelers can begin to de-black-box systems which were initially black-boxed, so as to produce a model with greater precision or a wider variety of accurate behaviors. A single term operating in the model can be expanded as a sub-network. Importantly, the fact that the more expanded representation must still reproduce the original behavior it produced as a single term is an important constraint that modelers can use to help build this sub-network and pin down parameters without having to reiterate the entire model. Similarly, interactions can be modeled using more complex relations, and so on.

The analogy Voit draws with semantic learning reinforces the essential cognitive function of mesoscopic models we gave above, which is to facilitate a process of cumulative controlled cognitive development in complex contexts. This attempt by scientific practitioners to rationalize their own activities, in effect, shares the main claims of the justification we have provided of the strategy of building mesoscopic models, i.e., containing structure supports more complex model development through building-out and other heuristic strategies. Modelers operate with learning constraints, among others, and their methodological choice behavior is boundedly rational. Although some kinds of purely computational systems might be capable of modeling whole complex systems upfront, distributed human-computer problem-solving systems are not.

Conclusion

In this chapter we have described the model-building processes of system biologists, arguing that these practices should be understood as boundedly rational responses to the complexity of the modeling problems they face. The concept of bounded rationality can thus play a useful role in the field of systems biology, providing cognitive means to justify its present activities. However, as we have seen, unpacking what precisely bounds problem-solving activity and why a set of given responses might be rational in the case of systems biology, or indeed in the case of any computational science, requires addressing the nature of the human-computational system interaction. Briefly, we have tried to argue that distributed cognition and simulative mental modeling can play this role. Indeed, our case above illustrates how Simon's notions of bounded rationality and problem solving as search can be combined with a distributed cognitive framework to help better articulate and study the complex nature of modern computational science.

Note

1 Given that one of the goals of systems biology is to be able manipulate biochemical systems, it thus important that the models they build do not just capture natural equilibrium behavior, but also what will happen in response to intervention on a system.

References

Alač, M., and Hutchins, E. (2004). I see what you are saying: Action as cognition in fMRI brain mapping practice. *Journal of Cognition and Culture*, 4(3), 629–661.

Barsalou, L. W. (1999). Perceptions of perceptual symbols. *Behavioral and Brain Sciences*, 22(4), 637–660.

Becvar, A., Hollan, J., and Hutchins, E. (2008). Representing gestures as cognitive artifacts. In M. S. Ackerman, C. Halverson, T. Erickson, and W. A. Kellog (Eds.), *Resources, co-evolution, and artifacts: Theory in CSCW* (pp. 117–143). New York: Springer.

Brass, M., Bekkering, H., and Prinz, W. (2001). Movement observation affects movement execution in a simple response task. *Acta Psychologica*, 106(1–2), 3–22.

Bryant, D. J., and Tversky, B. (1999). Mental representations of perspective and spatial relations from diagrams and models. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 25(1), 137.

Chandrasekharan, S., and Nersessian, N. J. (2015). Building cognition: The construction of computational representations for scientific discovery. *Cognitive Science*, 39(8), 1727–1763.

Christensen, B. T., and Schunn, C. D. (2009). The role and impact of mental simulation in design. *Applied Cognitive Psychology*, 23(3), 327–344.

Finke, R. A. (1989). *Principles of mental imagery*. Cambridge, MA: The MIT Press.

Glenberg, A. M. (1997). Mental models, space, and embodied cognition. In T. B. Ward, S. M. Smith, and J. Vaid (Eds.), *Creative thought: An investigation of conceptual structures and processes* (pp. 495–522). Washington, DC: American Psychological Association.

Gorman, M. E. (1997). Mind in the world: Cognition and practice in the invention of the telephone. *Social Studies of Science*, 27(4), 583–624.

Greeno, J. G. (1989). Situations, mental models, and generative knowledge. In D. Klahr and K. Kotovsky (Eds.), *Complex information processing: The impact of Herbert A. Simon* (pp. 285–318). Mahwah, NJ: Lawrence Erlbaum.

Hall, R., Wieckert, K., and Wright. K. (2010). How does cognition get distributed? Case studies of making concepts general in technical and scientific work. In M. Banich and D. Caccamise (Eds.), *Generalization of knowledge: Multidisciplinary perspectives*, New York: Psychology Press.

Hegarty, M. (1992). Mental animation: Inferring motion from static displays of mechanical systems. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 18(5), 1084.

Hegarty, M. (2004). Mechanical reasoning by mental simulation. *Trends in Cognitive Sciences*, 8(6), 280–285.

Hegarty, M. (2005). Multimedia learning about physical systems. In R. E. Mayer (Ed.), *Cambridge handbook of multimedia learning* (pp. 447–465). New York: Cambridge University Press.

Humphreys, P. (2009). The philosophical novelty of computer simulation methods. *Synthese*, 169, 615–626.

Hutchins, E. (1995a). How a cockpit remembers its speed. *Cognitive Science*, 19, 265–288.

Hutchins, E. (1995b). *Cognition in the wild*. Cambridge, MA: MIT Press.

Johnson-Laird, P. N. (1983). A computational analysis of consciousness. *Cognition & Brain Theory*, 6, 499–508.

Kirsh, D. (1996). Adapting the environment instead of oneself. *Adaptive Behavior*, 4, 415–452

Kosslyn, S. M. (1980). *Image and mind*. Cambridge, MA: Harvard University Press.

MacLeod, M. (2016). Heuristic approaches to models and modeling in systems biology. *Biology & Philosophy*, 31(3), 353–372.

MacLeod, M. (2018). Model-based inferences in modeling of complex systems. *Topoi*. https://doi.org/10.1007/s11245–018–9569-x

MacLeod, M., and Nersessian, N. J. (2015). Modeling systems-level dynamics: understanding without mechanistic explanation in integrative systems biology, *Studies in History and Philosophy of Science Part C – Biological and Biomedical Science*, 49(1), 1–11.

MacLeod. M., and Nersessian, N. J. (2018). Cognitive constraints, complexity and model-building. *History and Philosophy of the Life Sciences*, 40(17).

Nersessian, N. J. (1992). How do scientists think? Capturing the dynamics of conceptual change in science. In R. Giere (Ed.), *Minnesota studies in the philosophy of science* (pp. 3–45). Minneapolis, MN: University of Minnesota Press.

Nersessian, N. J. (2002). The cognitive basis of model-based reasoning in science. In P. Carruthers, S. Stich, and M. Siegal (Eds.), *The cognitive basis of science* (pp. 133–153). Cambridge: Cambridge University Press.

Nersessian, N. J. (2008). *Creating scientific concepts*. Cambridge, MA: MIT Press.

Nersessian, N. J. (2009). How do engineering scientists think? Model-based simulation in biomedical engineering research laboratories. *Topics in Cognitive Scien*ce, 1(4), 730–757.

Nersessian, N. J. (2012). Engineering concepts: The interplay between concept formation and modeling practices in bioengineering sciences. *Mind, Culture, and Activity*, 19(3), 222–239.

Nersessian, N. J., Kurz-Milcke, E., Newstetter, W. C., and Davies, J. (2003). Research laboratories as evolving distributed cognitive systems. In *Proceedings of the Annual Meeting of the Cognitive Science Society*. Mahwah, NJ: Lawrence Erlbaum Associates.

Osbeck, L. M., and Nersessian, N. J. (2006). The distribution of representation. *Journal for the Theory of Social Behaviour*, 36(2), 141–160.

Perrig, W., and Kintsch, W. (1985). Propositional and situational representations of text. *Journal of Memory and Language*, 24, 503–518.

Roschelle, J., and Greeno, J. G. (1987). Mental models in expert physics reasoning (Tech. Report No. GK-12). Berkeley, CA: University of California Press.

Schwartz, D. L. (1995). Reasoning about the referent of a picture versus reasoning about a picture as the referent. *Memory and Cognition*, 23, 709–722.

Schwartz, D. L., and Black, J. B. (1996). Analog imagery in mental model reasoning: Depictive models. *Cognitive Psychology*, 30(2), 154–219.

Shepard, R. N., and Cooper, L. A. (1982). *Mental images and their transformations*. Cambridge, MA: MIT Press.

Simon, H. A. (1976). From substantive to procedural rationality. In S. J. Latsis (Ed.), *Method and appraisal in economics*. Cambridge: Cambridge University Press.

Simon, H. A. (1996). *The sciences of the artificial*. Cambridge, MA; MIT Press.

Trafton, J. G., Trickett, S. B., and Mintz, F. E. (2005). Connecting internal and external representations: Spatial transformations of scientific visualizations. *Foundations of Science*, 10(1), 89–106.

Trickett, S. B., and Trafton, J. G. (2002). The instantiation and use of conceptual simulations in evaluating hypotheses: Movies-in-the-mind in scientific reasoning. In *Proceedings of the Annual Meeting of the Cognitive Science Society* 24(24).

Trickett, S. B., and Trafton, J. G. (2006). Toward a comprehensive model of graph comprehension: Making the case for spatial cognition. In *Proceedings of International Conference on Theory and Application of Diagrams* (pp. 286–300). Heidelberg: Springer.

Trickett, S. B., and Trafton, J. G. (2007). "What if…": The use of conceptual simulations in scientific reasoning. *Cognitive Science*, 31(5), 843–875.

Voit, E. O. (2014). Mesoscopic modeling as a starting point for computational analyses of cystic fibrosis as a systemic disease. *Biochimica et Biophysica Acta (BBA)-Proteins and Proteomics*, 1844(1), 258–270.

Voit, E. O., Newstetter, W. C., and Kemp, M. L. (2012). A feel for systems. *Molecular Systems Biology*, 8(1).

Voit, E. O., Qi, Z., and Kikuchi, S. (2012). Mesoscopic models of neurotransmission as intermediates between disease simulators and tools for discovering design principles. *Pharmacopsychiatry*, 45(1): 22.

Zhang, J., and Norman, D. A. (1994). Representations in distributed cognitive tasks. *Cognitive Science*, 18(1), 87–122.

Zwaan, R. A. (1999). Embodied cognition, perceptual symbols, and situation models. *Discourse Processes*, 28, 81–88.