Computational Neuroscience and Cognitive Modelling

a student's introduction to methods and procedures

Britt Anderson

SAGE

Los Angeles | London | New Delhi
Singapore | Washington DC

Chapter 1

An Introduction to the Ideas and Scope of Computational Methods in Psychology

Objectives

After reading of this chapter be able to:

- understand the goals of computational modeling for psychology and neuroscience;
- describe the limits of computational models;
- describe the role played by computers in the growth of computational modeling; and
- describe how the success of a model can be evaluated.

1.1 Overview

My overall objective in this chapter is to introduce some of the motivations for using computational models in psychology and to suggest some of the limitations. It is often easier to sharpen one's thinking on these points if there is some form of dialogue. If you do not have the chance to discuss these ideas with someone else, you might first read one of a number of recent position papers (e.g., McClelland, 2009; McClelland et al., 2010; Griffiths, Chater, Kemp, Perfors, & Tenenbaum, 2010) and try to come to your own personal opinion of their correctness. Try to form your own answers to the following questions:

- Why bother modeling at all?
- Can we make a brain? Should we try? Should we care? Is this important for computational psychology?
- What role should biological plausibility play in neural and psychological models?
- Do you need a computer to create a computational model?
- How do you evaluate a model?
- What is the importance of a model's assumptions?

1.2 Why Model?

One answer to this question might be because you believe that modeling is something that the brain itself engages in. By modeling you are trying to reproduce or discover the architecture of thought.

Mental models as a basis for cognition have their roots in the writings of Craik (1952). A more contemporary proponent for mental models as models of cognition is Johnson-Laird (e.g., Johnson-Laird, Byrne, & Schaeken, 1992). Just as a small mechanical model can be used to predict the oscillations of high and low tides, perhaps in our brains (or our minds) we have some small simulacra that reproduce the functional organization of our world. To model cognition then is to model our models of our world (Chapter 21 looks at one modeling platform for building mental models).

Modeling does not have to be justified so directly. Models can be viewed as simplified versions of complex phenomena. The phenomena may be too complex for us to directly grapple with, but by abstracting away some of the complications we can achieve a core set of features that we think are important to the phenomena of interest. This reduced, simpler set provides a more tractable basis for us to see how the pieces fit together and to understand better the causal relations.

This justification raises the question of how simple is too simple. How do we really know which features are unnecessary complications and which are core features? What are the criteria for making such a judgment? And doesn't this put us on the road to reductionism, where we reduce, ultimately, all the laws of thought to the models of particle physics? Would an explanation at that level, even if accurate, gain us a deeper understanding of human psychology?

1.3 Can We Make a Brain?

The trivial answer is that yes, we can make brains; and every parent has done so. A parent's actions lead to a new brain with all the attendant complexities and capacities. It is done with no insight into the process. We achieve no greater understanding. Therefore, we can argue that building a brain per se is not guaranteed to improve our understanding of the rules and procedures regulating human thought.

This example is not trivial. The tools of neuroscience are giving us an increasing understanding of neurons and their constituents. We know the make-up of many ion channels down to the level of genes and post-translational modifications. If we knew the position of every molecule in the brain, would we learn anything about how the brain as a whole functions? This can be restated as a question about whether we believe the brain and the mind to be merely the sum of their parts.

There are scientists pursuing this route to modeling the brain. The Blue Brain Project at the Ecole Polytechnique Federale Lausanne announces exactly these intentions on its website*:[1]

*Since website addresses change frequently, I use numbered page notes to refer to a URL, which is then listed in the back matter of the book. This will allow me to more easily update the links.

> Reconstructing the brain piece by piece and building a virtual brain in a supercomputer – these are some of the goals of the Blue Brain Project. The virtual brain will be an exceptional tool giving neuroscientists a new understanding of the brain and a better understanding of neurological diseases.

Do you share this conclusion?

A challenge to this approach claims that an understanding of the brain requires a consideration of scale. A sand dune is made up of individual grains of sand just as the brain is made up of individual neurons, but no amount of study of individual sand grains can explain the behavior of sand dunes. The behavior of a dune requires a large numbers of sand grains interacting. Is the same true of the brain?

Discussion: Are Psychological Concepts Experimentally Accessible?

We often use computational methods because something is experimentally inaccessible. This can be due to size or complexity or because, by its very nature, it is inaccessible to direct experimental manipulation. Do you agree that cognitive processes are inaccessible to direct experimentation? For example, can one *directly* manipulate working memory? And how does your answer to that question effect your position on the value of computational modeling?

I.4 Computational Models as Experiments

Another reason commonly given for using computational approaches in brain research is that they offer practical efficiencies. Generally, it is easier to run a computer program over and over again with minor variations than it is to repeat a behavioral task on numerous pools of subjects. In some cases it may not be possible to repeat behavioral investigations. Patient HM,[2] perhaps the most psychologically studied person of all time, had his problems as the result of a surgery that will never be done again. The only way to repeat HM's lesion is by modeling.

Under what circumstances is the quest for experimental efficiency justifiable as the basis for computational modeling? Some authors suggests that it is sufficient that computational models enable exploration. We have an idea on some topic, and before doing more focused, expensive, or time-consuming research we can "play" with our computer model to decide if further experiments are warranted, and which ones we should do first. But doesn't this presume we have more confidence in the model than we should? Isn't there a danger that a poorly chosen model could direct us in the wrong direction or lead us not to explore lines of research that we should? How do we monitor the use of exploratory modeling to make sure these outcomes do not occur?

Since computational models do not sacrifice animal participants or take the time of humans, they are touted as more ethical. Some studies involving conflict are not ethical in human research, but we can create a computer model of hundreds or thousands of little warriors and let them battle for territory as a way of exploring ideas about social conflict (see Chapter 22 for one modeling platform that can be used for this type of research).

Can we generalize from such data though? How do we validate hypotheses gained in this way?

Coherence and Concreteness

Words are subtle, and may mean different things to different people at different times. They may even mean something different to the same person at different times. Have you ever read something you wrote and wondered, “What was I thinking?” Models typically involve translating our ideas from words into formulas or computer programs. A by-product of this reformulation is a concreteness and clarity that may be lacking from our most careful phrasings. Even if a model does not say anything “new” or make novel predictions, this clarity can aid communication and improve reproducibility.

Modularity and Simplification

One attraction to modeling cognitive phenomena is that one can pull out a piece of a larger process and focus on it alone. This is justified, in part, by asserting that cognition is modular. The same rationale can be applied to modeling the “basal ganglia” as an isolated, functionally defined construct. Is this justification persuasive? Is cognition simply the result of isolated processes summed up? If we defined each of the important cognitive modules could we study them in isolation and by “gluing” them together understand the brain?

In complex environments it is often the embedding of elements in an environment that is necessary for observing a desired effect. A common example is the role of raindrops and rainbows. Water droplets are essential for the production of rainbows, but the study of rainbows cannot be reduced to the study of water droplets. It cannot even be reduced to collections of raindrops. One could have a complex modular model of water and its coalescence into drops, and still not have rainbows. If one wired a water droplet model to a sunlight model, would the models then have rainbows? Or would there need to be an observing visual apparatus (i.e., a person)? A name for phenomena that only appear in the large or from interactions is emergent; these are phenomena, like the sand dune, that cannot be understood from their individual pieces alone.

> Can you think of a school of psychology that concerned itself with emergent properties, a psychological school that emphasized that wholes were more than the sum of their parts? What do you think this school would have felt about computational modeling in psychology?

Models for Exploring the Implications of Ideas

One of the fathers of connectionism as an approach to modeling in psychology (Chapters 11 and 13 introduce two types of simple neural networks), James McClelland (2009), has emphasized two main purposes behind his own work. He states that he seeks to simplify complex phenomena, and that he wants to investigate the implications of ideas.

Implications

A justification for models as a route to implications rests on the assertion that models are more transparent than behavioral experiments. Computational models are suggested to be more transparent because they require greater specificity, and because they do exactly what you tell them to do (even if you did not realize what you were telling them at the time, a “logic bug”).

As modelers, we specify the rules and the players. There is no ambiguity. This give us an opportunity to demonstrate the sufficiency of our ideas. Models do not have to be biologically realistic to provide sufficiency proofs. If I assert that variation in light intensity across facial photographs contains information sufficient for identification, it is enough for me to build a program that can do it, even if I use methods that are inherently biologically implausible. My model represents an existence proof. While it does not show that people actually recognize faces from light intensity, it proves that they could. Note that this argument is one way. While a successful program shows that luminance information is sufficient for identification, we cannot draw a similarly strong conclusion from a failed program. We can only conclude that that program was inadequate, not that no program could ever be designed to do the same.

Typically, though, we have higher aspirations than simply stating our ideas clearly or providing existence proofs. We want to advance our understanding of the phenomena in question; we want to travel beyond where our ability to forecast can take us. When we specify a computational model we want to be able to ask, “What if?” In order for our models to be computationally tractable, and to permit us to still be able to follow what is going on, we can usually only implement simplified models. This brings us back to the questions considered above. Is our simplified model too simple? Can we make inferences about more complex systems from simpler ones? If the only circumstances in which we can take advantage of model concreteness for investigating model implications are those situations where we have had to simplify, reduce, or modularize the situation, do we ever really achieve this advantage? One criticism of studying cognition by computational simulation is that our simplifications are arbitrary. In the process of simplifying, paradoxically, we lose our transparency.

I.5 Do Models Need to Be Biologically Plausible?

> “There are more things in heaven and earth, Horatio, than are dreamt of in your philosophy.”

When we review neural networks (Chapters 11 and 13), we will see that modelers often draw inspiration from biology, but how close do models need to be to biology in order to be useful? Backpropagation is a tool used in neural networks. This error correction method relies on propagating the error signal from the end of a computation back to earlier nodes in the chain. It is not obvious that neurons in our brains have the requisite wiring or methods necessary to propagate errors backwards. Therefore, what can we conclude from a computational model that uses the backpropagation error correction algorithm?

One response is philosophical. This response claims that it is just too early in our research program to presume that our knowledge of biological mechanisms is sufficient to reject *any* possible account. Just because we do not know *now* of any way for real neurons to backpropagate an error signal does not mean that we will not know it next week, or the week after. This justification gives renewed vigor to the use of models as exploration and for examining the implications of ideas. If something turns out to be very useful in a computational model for solving a particular problem then maybe we should look harder for that mechanism in biology?

Another response is to simply answer no: models do not need to be biologically plausible. Models are abstractions. The scope of a model does not have to feature *neural* mechanisms to inform ideas about cognition. Further, goals from modeling can include coherence, concreteness, or sufficiency. None of those objectives require biological plausibility. Which answer prevails in a particular circumstance will depend on the questions being addressed with the model and the purposes behind the model's construction. Biological plausibility should not be used to abjure modeling, nor should it be required of all psychological and neuroscientific models.

1.6 Do Computational Models Require Computers?

Computation does not require a computer. It is true that presently we tend to associate computational models with computer programs, but this not a prescriptive relation, it is a descriptive one. Perhaps the best demonstration of this idea is that Alan Turing's ideas about computation came not from digital computers, but from human computers.

Before the word meant a machine, the term "computer" was used to describe a person. Turing thought about the stages in the process of doing a mathematical computation, the need for a temporary scratch pad to record interim results; it was from these ideas that he developed his thoughts on the limits of computation and ultimately the description of the Turing machine. One can assert a formal equivalence between people and machines, and Turing developed the idea of equivalence into a test to assess when machine intelligence rivaled a human's.

If human and machine computations are equivalent, then it is not *necessary* that computational models be implemented by machines. It may, however, be more convenient or efficient.

Many models that can be stated in simple formulas yield more insight when they are *not* converted to programs, because that makes it easier for humans to follow the cause and effect relationships. Simple probabilistic models (some of which appear in Chapter 15) often give a better glimpse into their overall behavior when they are stated as formulas instead of as computer programs. Models can be viewed as abstractions that emphasize some features and ignore others. The utility of a particular method of presentation depends on a model's purpose. Road maps and topological maps emphasize different features of a landscape and ignore others. Both are useful, but in different settings, and both are more helpful printed out than described as a computer program.

As further evidence that computational models do not have to be written as computer programs, there are the results of early computational investigations that preceded the invention of digital computers. For example, signal detection theory is a computational model commonly used in psychological experiments that was developed during World War II to characterize the challenges of radar operators attempting to distinguish the blips that were planes from the blips that were not. In the early 1900s Lapicque came up with a model of neural activity that was a mathematical model in the pre-computer age. Psychophysics was elaborated by Fechner drawing on the data of Weber. It was developed from first principles. Fechner conceived of how to make physical and psychological measurement scales commensurate and his justification was logic and calculus. While rather exceptional, these examples demonstrate that computational models do not have to be computer programs. The modern digital computer is a technical convenience. For some types of models, the implications are clearer when they are stated as formulas rather than computer programs.

However, for other types of models, the term "convenience" may not go far enough. It is true that all the calculations specified in a digital computer program might, in principle, be equivalent to those carried out by hand, but it is also often the case that the labor and time to do so would be prohibitive. This fact is the major reason that the growth of computational approaches in psychology and neuroscience has so closely paralleled the growth of desktop computing power.

Discussion: Can a Model Be too Simple?

Watch and read the following items and then discuss the following questions.

- Blue Brain TED talk by Henry Markham[3]
- Cat Fight Over Simulations[4]

What do we learn by simulating the brain as a large collection of single elements? How simple can those elements be?

1.7 How Do You Evaluate a Model?

The prior sections highlight issues about whether modeling can be justified, and for what purposes it can be used, but once you have decided to use a computational model how do you examine its quality? Specifically, are any of these the right way to evaluate a model:

- Goodness-of-fit
- Sufficiency
- Optimality
- Consistency with human performance?

It is probably fair to say that at this point in time it is premature to be prescriptive. Just as it may be premature for us to say that computational methods must comply with current biological knowledge, it may also be too early to require that our models fit data well. With our current knowledge limited, and with our current machines still not sufficiently powerful, requiring our model implementations to closely fit actual data might be a bar too high. Failure would lead us to throw away good ideas simply because we did not have the means to properly test them. Still, we cannot be too forgiving of ourselves or we would never discard one model in favor of another, and all models are always above average. In line with these ideas it seems we ought to expect that a model does what it claims to do. We cannot have a model of memory that never stores or recalls anything. If we are to be cautious in using models, we ought to expect new models to be in some way better than the models they seek to replace. Of course, better is a hard word to define in practice, and leaves us with the same sorts of options outlined above. One notion not included in the above list is the idea of simplicity or elegance. Are those adequate criteria for comparing models? Is it reasonable to think that a cognitive apparatus cobbled together over some tens of thousands of years, mostly in an effort to meet immediate environmental challenges, would be simple or elegant?

1.8 Do Models Need Assumptions?

Like almost any question of this type, the answer is, "It depends." Model assumptions can be conflated with model capabilities and model properties can be confused with claims about the process being modeled. Given that most current models will be much simpler than human cognition, it would be a mistake to look at what a model cannot do, and to infer that the model builders were making broad claims about underlying human cognition. Rather, models are built from certain points of view and with certain base claims. These claims provide properties for model operations and from these claims there emerge models that can do certain things. Such models can be used to argue whether certain properties are necessary for certain functions, but they cannot be used to claim that such demonstrations provide evidence for the mechanics of human cognition. Certain properties entail certain effects. It is the role of models, at this time, to point the way to properties that can be examined as possible assumptions about human cognitive processes.

1.9 Additional Questions

The preceding is meant to introduce some of the controversial and unresolved issues about the purpose and nature of computational modeling in psychology and neuroscience and to point out some of the remaining hard questions about what a model tells us and how we evaluate model success. Before rushing out to use a model, it is good to decide for yourself where you stand on these issues, especially how to evaluate model success. You might develop your own ideas by considering a few final questions:

1. Are neural network models better than other types of psychological models because they are built from biologically inspired elements?
2. What are the advantages and disadvantages of neural network models when compared to systems of if – then production rules?
3. What ultimately do you think it is reasonable to learn about human psychology from the use of computational models?
4. Is the current popularity of computational modeling deserved by new knowledge, or is it a fad for computers or taking the easy way out?
5. If behavior is a product of brain events, and our knowledge of the brain is very incomplete, is there any purpose at the present time to making computational models of neural or psychological phenomena?

1.10 Road Map

This chapter sets the broad agenda by asking each of us to examine what we think are the strengths and weaknesses of computational modeling and to decide what role we think it should have in psychology and neuroscience. The rest of the book is devoted to considering, in turn, a selection of mathematical topics and examples of how models are used. The organization for each section is similar. First, we consider a mathematical topic. To start simply, the mathematics is presented in isolation from much that has to do with psychology or neuroscience, although I do try to offer some motivation by giving glimpses of where we are heading. This may make things a bit dry, but it helps us to avoid distractions and it allows us to concentrate on terminology and notation. The early chapters on a section serve to highlight how a basic familiarity with mathematical topics, and a willingness to implement routines on a computer, can serve as the basis for some impressive simulations. The subsequent chapters develop the use of the mathematical ideas for a neuroscience or psychological model.