

Virtual Models and Simulations: A Different Kind of Science?

Peter R. Krebs
School of History & Philosophy
University of New South Wales

Abstract

The personal computer has become the primary research tool in many scientific and engineering disciplines. The role of the computer has been extended to be an experimental and modelling tool both for convenience and sometimes necessity. In this paper some of the relationships between *real* models and *virtual* models, i.e. models that exist only as programs and data structures, are explored. It is argued that the shift from experimenting with real objects to experimentation with computer models and simulations may also require a new approach for evaluating scientific theories derived from these models. Accepting the additional sets of assumptions that are associated with computer models and simulations requires ‘leaps of faith’, which we may not want to make in order to preserve scientific rigor. There are problems in providing acceptable arguments and explanations as to why a particular computer model or simulation should be judged scientifically sound, plausible, or even probable. These problems not only emerge from models that are particularly complex, but also in models that suffer from being too simplistic.

Introduction

In a recent volume by De Chadarevian and Hopwood (2004) a number of authors present some of their views on 3-dimensional models and the role such models play in science from a historiographical perspective. The various models discussed have in common that they are mostly *material* things, i.e. models made of clay, wood, plasticine and the like. In the last few decades computers have revolutionized scientific modelling, and the notion of *model* has changed. The use of ‘computer models’ does not just add another kind of model to the array of ‘traditional’ artifacts. In some disciplines computers have become *the* modelling tool, rather than merely playing a supplementary role. Indeed, in the field of Cognitive Science the computer model *is* the ‘traditional’ model, given the underlying computational theory of mind. Not only are many of the characteristics of computer models and simulations entirely different from material models, but the way we interact with models changes as a consequence.

Now, there are no longer real objects to probe, to measure or to collect, and all of our activities target mere *representations* of the world, i.e. mathematical abstractions, and computations with these representations (symbols). Moreover, a new layer of ‘*virtual reality*’ is often created with the aid of various visualization techniques. Experimentation with such models in an interactive and ‘interfering’ way that Hacking (1983) and Harré (1970) ask for is not possible. Instead, the experiments are conducted in the domain of the virtual and the computational paradigm. Yet computer models are sometimes deemed to be real world objects in the same way the objects that are modelled are real world objects.

Not long ago, the concept of simulation “invariably implied deceit” (Keller, 2003). I think that this sentiment also applied to the term model, albeit to a lesser degree. Simulations and models were thought of as merely mimicking, or faking, the real world. While modelling has become a widely used technique in almost any imaginable discipline, the term is still often associated with a certain amount of incredulity, or, skepticism. For every model that shows *A*, there seems to be

always an alternative model showing B , and it is significant that we quite often hear the expression 'it is only a model'. In contrast, the term *model* is also used to denote standards and even perfection: the *model* husband (Jordanova, 2004). Expressions like 'virtual models' and 'virtual experiments' might be preferable, because the term *virtual*, and in particular *virtual reality*, seem more positive and are usually associated with cutting edge computing and Artificial Intelligence (AI).

The arrival of modern computing machinery in the 1950s and the proliferation of inexpensive and very powerful PCs since have led to a revolution in terms of what kinds of models and simulations can be implemented. Computer models and simulations make use of many advanced techniques that introduce new, often exciting, ways to present aspects of models to scientists, science communicators and 'consumers' of science alike. Computer generated images (CGI) not only changed the way we think about pictures and movies, but also about how theories are formed. New and innovative methods have been devised to present data in both scientific and non-scientific contexts. There is a variety of powerful methods for visualization and presentation available, and many applications of these techniques have made their way into textbooks and journals in the form of illustrations and graphical representations. With advanced image processing techniques, it is not only possible to alter and to enhance pictures, but it is also possible to render images of phenomena that are not visible, or may not exist at all. Many of the computer aided experiments and visualizations may be helpful in understanding complex phenomena because "visualizations contribute to 'amplify cognition'" (Araya, 2003). However, due to some reservations about the validity of computer simulations as experiments and methods to gain 'scientific knowledge', it seems that virtual models introduce a different set of issues concerning scientific rigor. Accepting virtual models and virtual simulations as experimental or empirical tools in science, will force us to adopt some new form of '*virtual* scientific method'.

Building Models

During the process of building a computer model or simulation, several transformations, or translations, take place. In the first instance there is a transformation of the (sometimes) observable phenomena or theoretical entities and the relationships between them into their corresponding mathematical entities. The result is a mathematical model that has been described as an intellectual construct, or, a mathematical object (Jorion, 1999). Then there is a translation of the mathematical structures into computational entities that are designed to deal with the complexities of the calculations in an appropriate, effective and efficient manner. The third transformation takes place when the data, which has been generated or transformed by models, is translated into a format that is more easily interpretable by the experimenter. In models and simulations, where large amounts of data are involved, additional steps are usually taken to present the data in some sort of visual form. The final transformation occurs when the model is reinterpreted in the language of the initial problem, question or theory.

Mathematical Models

A mathematical model is an intellectual construct that is based on a mathematical object, which "does not tell anything about the world" (Jorion, 1999). Jorion believes that mathematical objects, without sensible interpretation, are all about syntax, and any of their meaning derives entirely from its structure. The inherent meaning held by a mathematical object is that

[...] some of the symbols which constitute [the mathematical object] impose constraints on others, some have no more meaning than the set of constraints they are submitted to (Jorion, 1999, 2).

This view of a *symbol* is comparable with that of a representational system (RS) of ‘Type 1’ suggested by Dretske, who says about these kinds of RSs that they are “*doubly* conventional: *we* give them a job to do, and then *we* do it for them” (Dretske, 1988). However, a mathematical model becomes more than just a collection of meaningless symbols, provided that a sensible interpretation is possible. Jorion goes as far as to say that the mathematical model and the part of the world that is modelled are *isomorphic*, provided that the interpreted model is meaningful, i.e. the model “makes sense”.

The benefits of a mathematical model for world comprehension are the following: if an interpreted mathematical model makes sense, then it is reasonable to assume that the type of relations which hold between the symbols in the model hold also between the bits of the real world which are represented in the interpretation of the model (Jorion, 1999, 3).

The analysis of many models, in terms of initial assumptions and the claims made later, reveals that there are many different opinions on what “makes sense”. Artificial neurons, for example, have very little in common with real neurons: they differ in their external functionality, their behavior, and their architecture. Other than a gross similarity in that they transform (integrate) several inputs to one output, they really share only the name. The isomorphism of biological neurons and mathematical neurons can barely be described as an ‘*approximorphism*’, let alone as an ‘*isomorphism*’, but the question of whether it “makes sense” to employ simplistic artificial neurons in cognitive models or not, is certainly not asked often enough. Psillos (1999) refers to ‘*modelling assumptions*’ that reflect the relationship between the model and the target physical system. He thinks that

[f]ar from being arbitrary, the choice of modelling assumptions for [the target system] *X* is guided by *substantive similarities* between the target system *X* and some other physical system *Y*. It is in the light of these similarities that *Y* is chosen to give rise to a model *M* of *X* (Psillos, 1999, 140).

I believe that these “*substantive similarities*” also capture the nature of the relationships between the mathematical description (model) and a theoretical entity, provided certain conditions are met. Some of these conditions are discussed later.

How do we derive a mathematical description of some relationship among physical (or mental) entities? There are several conceptual transformations and processes involved. In the following paragraphs I discuss some of the issues concerning *abstraction*, *formalization*, *generalization*, and *simplification*, because these operations should be considered fundamental steps in the process of building (or constructing) mathematical models.

Abstraction

Mathematical models refine the real world by introducing an element of abstraction. It is clear that a model should be simpler than that which is to be modelled. The process of abstraction involves several practices, all of which widen the gap between the sometimes observable

phenomena and an idealized description in mathematical terminology. Stufflebeam (1998) suggested that his cat Sophie's behavior, when dropped from two feet, "satisfies the distance function $D(t) = \frac{1}{2}gt^2$ ". The abstraction here includes the reduction of the cat Sophie to a point mass in Newtonian physics. The observable behavior of the cat Sophie in free fall differs from the idealized point mass. In fact, Stufflebeam's description of Sophie's behavior as $D(t) = \frac{1}{2}gt^2$ does not involve Sophie at all. Abstraction is the process of defining a general and idealized case for relationships between entities and processes. In the distance formula, g stands for the acceleration, and if we substitute the values for g of 9.8 m/s^2 or 32 ft/s^2 then we get a reasonable approximation of the conditions on Earth. However, we can also find the appropriate values for this model to work on the moon or on Mars. The distance formula holds *anywhere* for *any* object, provided we have the correct value for g . The most important aspect here is the introduction of placeholders like $D(t)$, which is the abstract notion of 'the distance of something at a particular moment in time'. This placeholder, or symbol, can now be manipulated within a formal system, like mathematics in this case. The introduction of symbols may put constraints on the type of operations and the methods for the model. For example, a sigmoid squashing function is selected in neuron models (perceptrons), because (1) the function's behavior is close to that of the step function at some level and (2) the function is differentiable at every point. While the qualities of the step function are desirable for the implementation of a neuron's functionality, some mathematical procedures (the back propagation of error algorithms for learning in this case) *require* that this function is differentiable everywhere. The point is that the mathematical methods that make up the model, or play an essential part in the model, are likely to dictate the kind of mathematical structures of the model at some level.

Formalization

The second and usually difficult process in building a mathematical model is that of *formalization*. A mathematical description of entities and the relationships that hold between them can only work as a useful model if there is a sufficient precision of terms. In areas of elementary Physics, like Newtonian dynamics, models work well because terms like *mass*, *velocity* and *force* and the interaction between these concepts are defined within a formal system, based on axioms. This is not the case in other scientific endeavors. The difficulty in Cognitive Science, for example, is that many terms describe mental things, like beliefs, behaviors and linguistic concepts, rather than physical things with properties that can be described and defined easily. Moreover, for mental concepts, we do not have clearly defined relations or processes to manipulate such concepts. Green (2001) suggests that some of the apparent success of connectionist models is due the lack of precision of terms (*vagueness*) and insufficient explanations of what it is that is actually modelled. The question is whether beliefs or behaviors can be modelled successfully, if it is not possible to provide a formal description of what we want to model. However, formal representations of a belief, for example, are needed in a computer program, because we need some way of encoding this concept. I suspect that formal descriptions of mental events, if it is at all possible to produce such descriptions, will not be in terms of simple placeholders. They will have to be either simple and relatively vague, or they will be very complex in order to provide some exactness and precision. But there is a catch: on the one hand there has to be sufficient precision to build a good model, on the other hand, precision in terminology and in detail makes it harder to build models that remain simple. Formalization ought to eliminate many of the 'soft' assumptions and descriptions about mental concepts. However, mental concepts are not easily defined or described in formal terms. For example, experience with the representation of knowledge in many applications in the field of AI have

shown, that it is very difficult to encode ‘facts’ and the associated rules. In these situations we have to face the additional problem of also having to encode the subjective *degree* of belief and quite likely fuzzy representations about what *is* believed. It is clear that we cannot choose suitable sets of symbols, sets of rules of inference and transformation rules for mental concepts in the same way we can choose $D(t)$.

Generalization

For some models it is desired that they work well for a theory about something “in principle”, rather than to target a particular instance of the theory in question. In other words, the model has to be able to produce data, if the model is designed to predict some cognitive behavior in humans, for *humans*, rather than the behavior of *Lucy* or *Bob*. There is, of course, the added problem of validating the model. In order to measure the success of the model, we need to compare data from the model with real data. The *real* data in this case has to be statistical data, because averaging data from many individuals can provide us only with generic *human* data. Generalization is usually achieved by omitting detail and allowing for a very broad interpretation of results. The danger here is to make models so general that they no longer capture the complexity of the theory or issue to be modelled (Krebs, 2005; Krebs, 2007). For example, the general formula for falling objects (e.g. cats) based on simple Newtonian physics is not a sufficiently precise model for what happens to a parachute jumper in free fall. For the latter much more specific case, it is important that drag and terminal velocity are considered in the model.

Simplification

One of the many criteria defining what makes a ‘good’ model is that the model is easier to work with. One way of making models easier is to simplify things, which can be achieved by disregarding details or external (environmental) issues that influence the model. For Sophie, the distance traveled by a free falling object on earth can be modelled using $D(t) = \frac{1}{2}gt^2 + v_0t + D_0$ where g is the acceleration of about 10 ms^{-2} , and v_0 is the vertical velocity at the beginning of the time interval t . $D(t)$ gives us the distance after the time t from the position D_0 , the position of the object at the beginning of the time interval. This is an ‘easy to work with’ model, because we do not take into account, among many other things, that (1) the acceleration is only *approximately* 10 ms^{-2} , and (2) the atmosphere causes drag. Even when taking drag into consideration, the mathematical model of Sophie’s behavior is still crude, because we have not considered the Reynolds numbers, the variation of the gravitational force over geographical regions, and many other perturbations. If we take drag into consideration we need to know that drag itself depends on, among other things, (1) the shape of the object and (2) air density. But, the density of the air is dependent on the temperature and the humidity, and the Reynolds numbers depend on the velocity of the object (cat), its shape and size, its surface, and so on. In the case of Sophie, the problem can not be fully described, because the cat could and would change its shape and therefore many parameters during the free fall.

At some point the model will become so complicated that it is no longer easy to work with, because the model is more difficult to understand than the original problem. $D(t) = \frac{1}{2}gt^2 + v_0t + D_0$ is likely to be sufficient as a model for most ‘dropping cat problems’. I consider the task of simplification to determine what must be included in the model, and what kind of detail can be omitted, the most difficult challenge. As computers become more powerful, the computational

complexity of models can be increased, which in turn, as one would expect, will increase the quality and power of the models. But this is not necessarily the case. A very complex model is no longer easy to use, and an increased complexity can also be an indication that the model needed many additions to explain or produce acceptable predictions. In the same way Tycho Brahe's model of the universe needed more and more epicycles to 'keep up' with the data that was gleaned from actual observations. It might turn out that the model is just not good enough to explain things adequately.

Experiments

It has been suggested by some Philosophers of Science, e.g. Harré and Hacking, that experimentation is not merely about the observation of phenomena and subsequent inferences to the explanatory theories. Instead, experimentation is also about observing and *interfering* with the objects in question (Hacking, 1983). The ability to manipulate objects is an essential and integral part of the process of experimentation, which is "to create, produce, refine, and stabilize phenomena" (Hacking, 1983). The close connection between the experiment, a material model and the real world is also a key requirement in a definition of the term *experiment* offered by Harré, who says that

[a]n experiment is the manipulation of [an] apparatus, which is an arrangement of material stuff integrated into the material world in a number of different ways (Harré, 2003, 19).

Harré suggests that the experimental setup (apparatus) is either an instrument that can tell us something about the world due to the causal relationships between the setup and the 'states of the world', or it is a "domesticated version of the systems in the world" (Harré, 2003, 26).

The kinds of experiments that fit the criteria suggested here are associated by some, naïvely perhaps, with what actually happens in a laboratory. These are the kinds of experiments that remind us of our high school days. However, it has become obvious that the vast majority of experiments are different from this stereotypic view (Morgan, 2003). When we conduct experiments with computational models and simulations, there are no materials that could possibly be manipulated. The material, the apparatus and the process of interference are all replaced by data structures and computational processes. The nature of the entities and the phenomena that are the points of interest in the field of Cognitive Science, for example, dictates that models and simulations are often the only way to do any experimentation at all. In Cognitive Science, the experiment is moved into the realm of the *virtual*, not just for convenience, but more often than not, out of necessity.

Virtual Experiments

Elements of computation can be part of a causal chain. Imagine an experimental setup where micro-electrodes are used to measure some voltage changes in a living cell in response to some stimulus introduced with another set of micro-electrodes. Instead of using a voltmeter that is built around a mechanism involving a coil, a magnet and a pointer with a dial, the voltage is displayed on a computer screen. The voltage differential at the electrodes is converted into a digital signal so that a particular voltage is represented (encoded) as a binary bit pattern. This data is fed through one of the computer's input/output channels, and a program performs the task to convert and display the data as a series of figures, i.e. numbers, on the screen. There are, in principle at least, no difficulties in explaining the causal chain between the number on the screen and

electrical potential at the micro electrodes. The numbers on the screen are elements of a Type II RS in the Dretskeian theory. RSs of Type II are grounded in the real world in that their power to *indicate* is linked to causal events in the real world. Linking meaning to causal events has also been suggested by Russell, who defines “causal lines” as

[...] a temporal series of events so related that, given some of them, something can be inferred about others whatever may be happening elsewhere. A causal line may always be regarded as the persistence of something - a person, a table, a photon, or what not. Throughout a given casual line, there may be consistency of quality, consistency of structure, or gradual change in either, but no sudden change of any considerable magnitude. I should consider the process from speaker to listener in broadcasting one causal line: here the beginning and the end are similar in quality as well as structure, but the intermediate links - sound waves, electromagnetic waves, and physiological processes - have only a resemblance of structure to each other and to the initial and final terms of the series (Russell, 1948, 477).

The meaning (i.e. our interpretation of the semantic content) is not bound to the real world in the same way. The power to indicate something about the real world has to be recognized by the observer of the sign. Scientific instruments, thermometers, or voltmeters, indicate temperature, electric potential and similar properties and phenomena. They function by exploiting (detecting) a known physical phenomenon. Instruments provide the observer with a representation of the state or relation of that phenomenon through a series of often complex transformations. For example, it is a property of the real world that the volume of a quantity of metal varies with temperature. A suitable arrangement of levers and a pointer on a dial can be used to exploit the relationship between temperature and volume to create an instrument that will indicate the temperature with some accuracy. There is a distinction between what the instrument indicates and what the observer believes that indication *means*. The pointer on the dial will only be meaningful to someone who *knows* that this instrument is indeed a thermometer. The instrument will *indicate* the temperature quite independently from the observer. To be an indicator of some particular property of the real world, the causal relationships must be maintained and the observer must attach the right kind of interpretation in terms of the indicator’s meaning. Dretske explains that

[i]f a fuel gauge is broken (stuck, say, at “half full”), it *never* indicates anything about the gasoline in the tank. Even if the tank *is* half full, and even if the driver, unaware of the broken gauge, comes to believe (correctly, as it turns out) that the tank is half full, the reading is not a sign - does not mean or indicate - that the tank is half full. Broken clocks are *never* right, not even twice a day, if being right requires them to *indicate* the correct time of day (Dretske, 1988, 308).

In the suggested example, i.e the computer indicating voltages and the like, the computer is an integral component of the experimental setup, but the computer is not implementing a *virtual* model. Consider the following changes to the experiment. The computer program is now modified to read the pattern and to display the corresponding number every second, and as an additional feature, the program records the time and the values from micro-electrodes in the machine’s memory. The information in the memory can also be *replayed* so that the sequence of numbers is displayed on the computer screen in one second intervals. Essentially the computer is now *simulating* the original experiment by re-playing what happened earlier. Is there now a problem in causally linking the patterns in memory to the micro electrodes? I suggest that there is not. There is only a time delay that has no bearing on the ‘causal chain’, because the data for

the simulation has been obtained by means that is (was) in principle ‘causally’ traceable. The experiment continues and a mathematically minded researcher recognizes that a pattern seems inherent in the data. She pushes the data through her favorite statistics program on her computer and finds a very good fit of the data for some function $f(x)$. Because of the difficulties when dealing with living neurons in these kinds of experiments, it is decided to build a model of the neuron’s function based on $f(x)$. The neuron and all micro-electrodes are dispensed with, and the stimulus is now generated within the computer model varying the values for x in the domain of f , and the corresponding values $f(x)$ is displayed on the screen. Obviously, we can now easily produce many more data points to fill any gaps. The data that comes out of the model is different in terms of its origin and therefore also different in terms of what conclusions we can draw about f .

While the causal connections of some symbol *could* be traced, at least in principle, the *actual* connection to the real world is merely *assumed*. This assumption, because it is at least in principle traceable, maintains or promotes the symbol to an element of a Type II RS. The relationship between the real world and representations of the world is of great importance in the context of models and simulations. The value of a model as a representation of the real world and any insights into the working of the world by investigating properties of the model depends on the kind of representations the model employs. A *meaningless* representation, in the sense that it can represent arbitrarily *anything*, can and will render the entire model meaningless, unless there is a syntactically correct procedure (probably a causal chain) to tie these representations down. We can accept that some mathematical constructs and computer programs produce useful data (predictions) or that they perform suitably in the context of a particular problem, without having any similarities to the entities and relations of the problem at hand. However not all such ‘models’ may be able to offer any explanations or insights in another domain. For example, some computer programs, which may be designed to follow principles from the field of AI, perform the task of reading aloud some arbitrary text surprisingly well. However, these programs do not offer anything in terms of how a human being performs the same task - these programs are *faking* it, even if Artificial Neural Nets are involved (Krebs, 2005).

A model, or representational system, that is to function as a representation of the real world ought not to contain any Type I elements. In addition, representations of Type II, by definition, must not have gaps or uncertainties in the causal chain linking them to the real world. A thermometer is only a thermometer if it has the power to indicate the temperature. Some apparatus may well indicate the temperature provided certain other conditions are given. An example will illustrate this point. Imagine a partially inflated balloon that is connected to a pressure gauge. The volume of air and the air pressure inside the balloon will change with the ambient temperature and the ambient air pressure. This setup will function as a thermometer, *if* the ambient air pressure is kept constant. However, if the temperature is kept constant and the ambient pressure is allowed to vary, then the instrument will indicate pressure. This simple instrument has the power to indicate either *temperature* or *pressure*, that is, the setup can function as a thermometer *or* a barometer. A scientific, or a merely *usable*, instrument would have to be engineered so that the relationships between pressure, temperature and volume are exploited. But the power to indicate one or the other must be constrained through appropriate means to guarantee an indication of either only pressure or only temperature.

Type II representational systems contain *natural signs* that are objectively connected to the real world and their power to indicate something about that world is exploited by *using* their natural meaning (Dretske, 1988), because

[i]n systems of Type II, natural signs take the place of symbols as the representational elements. A sign is given the job of doing what it (suitably deployed) can already do (Dretske, 1988, 310).

It is important to note that there is no intentionality associated with this type of representation. However, the potential intentionality (the meaningful interpretation) is constrained by the causal links to the real world. The variation in volume of metal, for example, may be due to the change of temperature, but this variation in volume cannot be reasonably attributed to the colour of the paper it is wrapped in.

Computer models and computer simulations have become tools for science in many ways. AI and Computational Neuroscience are special cases among the ‘hard’ sciences in that computation *is* the very nature of their activities. Other sciences might employ computational models and simulations as tools, however chemistry, for example, is essentially about elements, molecules, compounds, plastics or pharmaceuticals, even if computational models and simulations play a role in chemical research. AI, in contrast, takes computational models to simulate, even replicate, cognitive functions that are *computational* themselves. This would certainly be the case if the assumption that cognition *is* computation is true. If it turns out that cognition is merely computable, then AI would still be entirely about computation, but the contributions to Cognitive Science would need additional justification.

Three distinct types of models have been identified where, (1) computers are used to deal with theories and mathematical abstractions, which would otherwise be computationally intractable, (2) computers provide responses (data) in ‘what-if’ simulations, i.e. the behavior of a real world physical system is simulated according to some theory, and (3) computers simulate the behavior of non-existing entities, for example the simulation of artificial life (Keller, 2003).

The role of computational models and ‘virtual experiments’, i.e. simulations, as contributors in the framework of empirical science are of particular importance. This holds especially for Cognitive Science because many of the objects of inquiry in Cognitive Science cannot be observed directly or mediated by scientific instruments. Consequently, models and simulations are often the *only* method available to the scientist. It has been argued that computer simulations are essentially extensions of numerical methods, which have been part of scientific reasoning for a long time (Keller, 2003; Gooding, 2003). Human beings do not reliably maintain accuracy when they have to deal with a large quantity of numbers, and digital machines are much more efficient at doing logical and numerical calculations. The recognition of patterns and structures is much more the domain of human beings. The work of analysis and interpretation of patterns, whether these are observed directly or whether these are produced by a machine, remains largely the task of the scientist. Ziman (2000) suggests that what can be known to science is restricted to what is known to scientists, when he says that “[a]n empirical scientific fact originates in an observation - an act of human *perception*”(Ziman, 2000, 102).

Experiments that are conducted in a virtual and computational environment often do not allow access to the object of inquiry. The question, whether evidence from *virtual* experiments qualifies as *empirical*, is still debated. One of the issues within this debate concerns the relationship between behavioral models and simulated or virtual objects on one hand, and real world behavior and the real world objects on the other. Are these virtual entities *representations of* or are they *representative for* the real world object?

Computer Models as *Scientific Experiments*

Models representing theories (*conceptual* models) and models representing real entities (*representational* models) must be accommodated within the framework of scientific practice. The conceptual model is the kind of model that has been associated with the terms *metaphor* and *analogy* by Bailer-Jones (2002). The general claim is that *all* models are metaphors. In this view, models are

an interpretative description of a phenomenon that facilitates access to that phenomenon. [...] This access can be perceptual as well as intellectual. [...] Models can range from being objects, such as a toy aeroplane, to being theoretical, abstract entities, such as the Standard Model of the structure of matter and its particles (Bailer-Jones, 2002, 108).

Some models can be an adequate representation of real entities provided that there is sufficient accuracy with which a model represents the real world. *Sufficient accuracy* is not a clearly definable term. What is 'sufficient' is essentially a matter of one's subjective stance toward the question what *science* is and how it operates. Psillos, defending the position of scientific realism, says that

taking a realist attitude toward a particular model is a matter of having evidence warranting the belief that this model gives an accurate representation of an otherwise unknown physical system in all, or in most, causally relevant respects (Psillos, 1999, 144).

We can consider and may even be able to defend the view that models in science go beyond being 'interpretative descriptions' and that they are scientific truths instead. Psillos hints that the adequacy of a model as a representation can only be determined on a case by case analysis, when he refers to the realist attitude toward a *particular* model. We will have to accept that the judgment whether a model or a simulation, or any experiment with such a model, is grounded in some *scientific* method, will also have to be made on a case by case basis. I have already shown that there are no rules for building models, and that the process of building models is largely based on assumptions about what the relevant factors are, how things can be simplified, how we write a program, and so on. The question of whether *virtual* simulations and *virtual* models are valid tools for a *scientific* endeavor is even more problematic. Ziman, who argues for a normative view of science, comments that

[m]ost people who have thought about this all are aware that the notion of an all-conquering intellectual 'method' is just a legend. This legend has been shot full of holes, but they do not know how it can be repaired or replaced. They are full of doubt about past certainties, but full of uncertainty about what they ought now to believe (Ziman, 2000, 2).

I believe that thoughts by Popper (1959) on how science should operate are still normatively useful. Theories should be formulated such that they are testable, and neither magical ingredients nor magical methods should be allowed as part of the supporting evidence. This, of course, must also apply to any counter examples and counter arguments. The application of models is a part of the empirical process. Helping to flesh out the details of some theory or to formulate a new hypothesis using models and simulations is also part of a scientific framework (Popper, 1959,

106). The epistemological role of simulations and models in science in terms of the development of theories is closely linked to questions about scientific theories in general (Peschl and Scheutz, 2001). Nevertheless, I believe that models and simulations are scientific tools, provided 'good scientific practice', whatever that may entail, is applied. Claims for a particular model or simulation should be judged as an adequate representation, and as an adequate, i.e. suitable, model, need to be examined in each case. We need to check that each part and process of a model can be mapped onto the corresponding part of the real world object or process that is modelled. In the case of a computer model, the elements and links in the data structures links and their relations to the object that is modelled need also to be explained. A computer model has to be *testable* in two ways. Firstly, we can test that the model is adequate in terms of what it models, and secondly test how the model is implemented and whether the implementation itself is adequate.

Churchland and Sejnowski (1992) note that real worldness has two principal aspects, namely (1) that the world is more complex, so that scaling up models does not always succeed, and (2) that real world events do not occur in isolation. Consequently, virtual models and virtual experiments lack realism in several ways. Like many other more 'conventional' models, they do not scale well, both structurally and functionally, and the virtual implementation by means of computation, reduces the number of similarities to real world objects even further. In a way, computer simulations introduce a second layer of abstraction. The first layer is the abstraction or conceptualizing of real world phenomena into a model. The second is the simulation of the model and its dynamics into the realm of the virtual.

Levels of Explanation

Models and simulations are targeted at different levels of explanation. A model can be used to explain certain aspects of a neuron, a particular phenomenon within a neuron, or the behavior of a collection of neurons. Another way to specify levels of explanation of models concerns the model itself. Models and simulations have a high level task to explain something. This level is likely to be the most abstract, and much of the model's implementation and internal workings may be of little interest. If, for example, we are presented with a simulation of the behavior of a few neurons on a computer screen, the actual implementation is of no concern to the observer or experimenter. The neuron simulation works (hopefully) as it should - it should work according to a set of specifications, which the experimenter is aware of. However, there are many layers of programs, library functions, operating system, device drivers, integrated circuits, gates, resistors and wires. The laptop computer, which I am using now, has several quite different programs for neural simulations stored on it. Most of these models are trying to explain the same thing at the highest or abstract level. They are all about relatively simple artificial neural nets, Hebbian learning, learning algorithms e.g. back propagation, and so on. The fact that the 'neurons' in these programs are mathematical structures involving mostly linear algebra is not essential to know or understand in detail for many users of the computer programs. The implementation of the mathematical engine, the subsystem that evaluates and transforms the matrices, is accessible only to the mathematically oriented computer programmer. Then, of course, there are all the components and systems that are part of the implementation on an actual machine. Very few of us have a deep understanding of the technical details of these systems and components.

Conclusion

Models and simulations are part of what is considered *scientific method* in the empirical sciences, although it is not clear what the term *scientific method* actually denotes. In some scientific disciplines, like in the field of Cognitive Science, there are many phenomena which do not belong to the *observable* world; but as Peschl and Scheutz point out that “[i]t is exactly this ‘hidden character’ of many cognitive processes which makes this domain so interesting as an object of scientific research” (Peschl and Scheutz, 2001).

This holds true for other disciplines. The fact that many processes are not accessible for direct or indirect observation is also interesting in terms of what can be modeled and simulated. It is not so much the mode of experimentation. Whether real world objects or ‘virtual objects’ are the targets of the experiment does not seem to be that much a point of controversy. It is, I suspect, the human contribution during analysis and interpretation that makes the experiment and the results appear to be ‘reasonable’ in terms of their value as scientific evidence. We should not forget that with the ever increasing complexity of computer hardware and the operating system software, it is impossible for most application programmers to understand much of these system ‘operations’ in any detail. Some of the users of software that offers a friendly interface for experimentation with artificial neural nets, for example, may not understand how the neural nets work on a theoretical level, or how they are implemented mathematically or as programs. However this is a point of concern, in the same way it *should be* a concern when using any other kind of technical equipment in scientific experimentation. The challenge remains for the provision of suitable explanations of how the apparatus (computer) works, and more importantly how the model or simulation that is implemented on the computer relates to the real world. The explanations will need to be different, due to the inherent difficulties in demonstrating causal chains in a virtual world.

References

- Araya, A. A. 2003. “The hidden side of visualization”, *Technè*, 7(2):27–93.
- Bailer-Jones, D. M. 2002. “Models, Metaphors and Analogies”, in: P. Machamer and M. Silberstein, eds., *The Blackwell Guide to Philosophy of Science*, Malden: Blackwell, 108–127.
- Churchland, P. S. and Sejnowski, T. J. 1992. *The Computational Brain*, Cambridge: MIT Press.
- Dretske, F. 1988. “Representational Systems”, in: T.O’Connor and D. Robb, eds., *Philosophy of Mind: Contemporary Readings*, New York: Routledge, 304–331.
- Gooding, D. 2003. “Varying the Cognitive Span: Experimentation, Visualization, and Computation”, in: H. Radder, ed., *The Philosophy of Scientific Experimentation*, Pittsburgh: UoPP, 255–283.
- Green, C. D. 2001. “Scientific models, connectionist networks, and cognitive science”, *Theory & Psychology*, 11(1):97–117.
- Hacking, I. 1983. *Representing and Intervening*, Cambridge: Cambridge UP.
- Harré, R. 1970. *The Principles of Scientific Thinking*, Chicago: UoCP.
- Harré, R. 2003. “The Materiality of Instruments in a Metaphysics for Experiments”, in: H. Radder, ed., *The Philosophy of Scientific Experimentation*, Pittsburgh: UoPP, 19–59.
- Jordanova, L. 2004. “Material Models as Visual Culture,” in: S. De Chadarevian and N. Hopwood, eds., *Models: The Third Dimension Of Science*, Stanford: Stanford UP, 443–451.
- Jorion, P. 1999. “What do mathematicians teach us about the world? An anthropological perspective”, *Dialectical Anthropology*, 24(1):45–98 (reprint 1–25).
- Keller, E. F. 2003. “Models, Simulation, and ‘Computer Experiments’”, in: H. Radder, ed., *The Philosophy of Scientific Experimentation*, Pittsburgh: UoPP, 198–215.

- Krebs, P. R. 2005. "Models of cognition: Neural possibility does not indicate neural plausibility", in: B. Bara et al., eds., *Proceedings of the 27th annual meeting of the Cognitive Science Society*, Mahwah: Lawrence Erlbaum Associates, 1184–1189.
- Krebs, P. R. 2007. "Smoke Without Fire: What do virtual experiments in cognitive science really tell us?", in: N. Srinivasan et al., eds., *Advances in Cognitive Science*, Delhi: Sage. (in press)
- Lenat, D. B. and Guha, R. V. 1990. *Building Large Knowledge-based Systems*, Reading: Addison-Wesley.
- McClelland, J. and Rumelhart, D. 1986. *Parallel Distributed Processing: Explorations in the Microstructure of Cognition (Volume 2: Psychological and Biological Models)*, Cambridge: MIT Press.
- Morgan, M. S. 2003. "Experiments without Material Intervention", in: H. Radder, ed., *The Philosophy of Scientific Experimentation*, Pittsburgh: UoPP, 216–235.
- Peschl, M. E. and Scheutz, M. 2001. "Explicating the epistemological role of simulation in the development of theories of cognition", *Proceedings of the seventh colloquium on Cognitive Science ICCS-01*, 274–280.
- Popper, K. R. 1959. *The Logic of Scientific Discovery*, London: Routledge.
- Psillos, S. 1999. *Scientific Realism: How science tracks truth*, London: Routledge.
- Rumelhart, D. and McClelland, J. 1986. *Parallel Distributed Processing: Explorations in the Microstructure of Cognition (Volume 1: Foundations)*, Cambridge: MIT Press.
- Russell, B. 1948. *Human Knowledge: Its Scope and Limits*, London: Allen & Unwin.
- Stufflebeam, R. S. 1998. "Representation and Computation", in: W. Bechtel, W. and G. Graham, eds., *A Companion to Cognitive Science*, Malden: Blackwell, 636–648.
- Ziman, J. 2000. *Real Science: What it is, and what it means*, Cambridge: Cambridge UP.