The Philosophy of Computer Simulation

Kevin B. Korb and Steven Mascaro
Clayton School of Information Technology
Monash University
Clayton, Victoria
Australia

kbkorb@gmail.com
sm@voracity.org

Abstract

The growth in the application of computer simulation across the sciences, and especially the application of artificial life techniques (agent-based and individual-based modeling) to evolutionary biology and psychology, the social sciences, epidemiology and ecology, raises many philosophical questions. One basic question is: How can we learn about reality by simulating it? Attempts to answer that question revolve around the relation between experimental procedure and simulation. We consider the relation between simulation and real-world experiment and find it to be the identity relation.

Keywords: Computer simulation, epistemology of simulation, Monte Carlo, homomorphism, validation, verification, experimentation.

1 The Scope and Limits of Computer Simulation

Computers have rapidly become the primary intellectual tool deployed by humans. This is natural. One of the first things students of computer science learn is that computers are *universal:* within the range of computable functions, there is simply *nothing* that computers cannot do. Every normal programming language is, in fact, a universal Turing machine, as can be proved easily by programming a simple universal Turing machine in that language. This universality, coupled with the rapid expansion of computational power, means that computers support almost everything that occurs in developed economies. It also means that computers can be applied to nearly any intellectual task, if not as an independent source of innovation, a use waiting at least upon some considerable new developments in artificial intelligence, then as a helpmate and support. Computers can be, and have been, applied to further research in biology, chemistry, microphysics, macrophysics, economics, sociology, art, music and philosophy. Computer simulation has become a reliable and regular contributor to investigation in each of these fields of endeavor, and probably every science. This expanding reach of computation has led to extreme reactions, including those who see computation as essentially inferior to human inference, and these uses of computation as epistemologically suspect, and those who see no bounds to computer application and who call for a new epistemology to underwrite these activities.

Here we shall attempt to develop some ground between the more extreme reactions to scientific computer simulations. In particular, we find considering the relation between computer simulations and scientific experiments to be interesting and fruitful. Many within the new epistemology camp have been suggesting that simulating "lies between" theorizing and experimenting (Humphreys, 1993; Winsberg, 2003; Rohrlich, 1991); if simulation is part theory and part experiment, then the old stories of how we learn about theory from experiment can hardly apply. However, we do not agree with this "in-betweenness" theory; rather, we suggest that the old stories about the growth of scientific knowledge, whether right or wrong about science before the computer, are equally right or wrong about current science.

1.1 What Computers Can't Do

There is general agreement that (ordinary) computers cannot compute non-computable functions, e.g., solving Turing's Halting Problem or computing the Busy Beaver numbers. To generate solutions to such problems computers would need to have access to infinite precision real numbers, for example, which is something no finite digital machine can manage. There is not any consensus about what this restriction really means, however. Some, such as Penrose (1999), seem to think this implies that computers are significantly inferior in potential computational ability to analog computers, such as humans. But for such a potential to be manifested, one must find an analog means of taking advantage of infinite precision real numbers, which presupposes overcoming quantum limits and pervasive low-level thermal noise. Since such limits have been operative throughout the entire evolutionary history of humanity, and since human mental capacities have certainly evolved largely for their adaptive value, it follows that human mentation as it currently exists has no more ability to break the barrier of non-computability than does the humble desktop computer. If there is potential to break through that barrier, we can hardly expect unassisted evolution to find it. We find fully satisfactory Turing's original answer to a similar complaint put to the possibility of machine intelligence — that any formal system is constrained to a proper subset of the truth by having a Gödel sentence true of it: who is to say we are any different? Computers are limited. But only a fool can fail to see the many severe limitations of humans.

If we find something that humans *can* do, then we have a *prima facie* case that computers can do it too. If we are too stupid to figure out how to get them to do it, that is a problem not attributable to them.

2 What Is Simulation?

So, what are simulations? The PC game "The Sims" is a simulation: it simulates the life and times of various characters who worry about getting jobs and cleaning toilets. Aircraft and naval piloting simulators simulate conditions involved in normal and abnormal maneuvers of aircraft and ships. And Second Life simulates a large range of human and non-human activities. Despite many commentators on the philosophy of simulation taking these sorts of cases seriously (e.g., Frigg and Reiss, 2008; Humphreys, 2004; Kueppers et al., 2006), in all of these simulations a human user plays an essential and central role, which is not to the point in simulation science. It is also not to the point that in ordinary language these processes are called simulations; that usage is simply emphasizing that humans are being put into other-than-real-world situations. Such simulations are not in general being used to expand our scientific knowledge, and so they do not raise the epistemological questions we wish to engage here. The simulations of interest to us here are those in which the *entire* simulation occurs within a computer, as a computer process. Indeed, we shall argue that the simulations of interest here are computer processes which simulate *other processes* — whether chemical, ecological, astrophysical or from whatever other scientific study.

2.1 A Definition of Simulation

A commonly used definition is:

Definition 1 *A* **computer simulation** *is "the use of a computer to solve an equation that we cannot solve analytically."* (Frigg and Reiss, 2008)

See also, for example, Humphreys (1991); Pritsker (1979); Kueppers et al. (2006); Winsberg (2001).[1] A comment of Reddy's (1987, p.162) might be confused with this kind of definition: "Simulation is a tool that is used to study the behaviour of complex systems which are mathematically intractable." That, however, would be a confusion of an accidental with an essential property: we use tools where they are useful, and not where they can be used but are unhelpful.

Definition 1 itself, however, includes both too much and too little.[2] Whether "we" can or cannot solve an equation analytically is surely immaterial. For one thing, that would render the term absurdly relative to the individual; for example, many programs which for us would be simulations would not count as simulations for a John von Neumann. For another, as new analytic techniques become available, what once counted as a simulation may not any longer.

We do not want a concept of simulation which is relative to time, place or individual calculational ability; we want a concept which is secured by a

[1] It is worth noting that Humphreys has retracted this view, finding the arguments of Hartmann (1996) persuasive (Humphreys, 2004, p.108).

[2] This is a point originally made by Hartmann (1996).

methodological role within science. But focusing on the positive side of the definition, things only get worse.

It's true that computer simulation began with the work of von Neumann, Metropolis and others working out ways of computing solutions to equations required for the development of the hydrogen bomb. This lead to such procedures as "Metropolis sampling", Monte Carlo (MC) integration and Monte Carlo methods in general. Monte Carlo integration is a method of numerically solving an unanalysable (or difficult to analyse) integral;[3] it does so by averaging pseudo-randomly selected values of the function in question. One can think of it as throwing darts at a board where the curve is drawn and using the frequency of darts under the curve as an estimate of its area. MC integration contrasts with numerical quadrature, which sums the areas of rectangles bounding portions of the curve. Nobody talks of the latter as simulation, but it solves equations just as well as the MC approach (at least given moderately well-behaved curves and low dimensionality). However, under the definition above quadrature counts as simulation. And, if MC integration were to be counted as simulation, we can't see any reason to deny the application to numerical methods generally, since they are all about solving things with computers that we cannot solve in our heads.

However, we think it is far preferrable to deny that equation solving is simulation and reserve that term for (computer) processes which mimic relevant features of a dynamic physical process under study (which is Hartmann's definition (Hartmann, 1996, p.83); see also (Zeigler, 1976; Pritsker, 1984)). Racynski and Bargiela (2007) have recently put this nicely in their first sentence: "To put it simply, computer simulation is a process of making a computer [process] behave like a cow, an airplane, a battlefield, a social system, a terrorist, [an] HIV virus, a growing tree, ... or any other thing."

2.2 Dynamic versus Static

Frigg and Reiss (2008) have objected to the idea that simulations are inherently dynamic, being processes that model other processes. It does not matter, according to them, that the computer process *takes* time, so long as it *represents* time: "[A]ll that matters is that the computer provides states that come with a time index.... If ... we have a computer that can calculate all the states in no time at all, surely we don't feel we lose anything."[4]

It is, of course, true that if the computer process encodes a representation of all the time steps of the target process, whether simultaneously or not, then it contains all the *information* that a simulation would carry or convey. However, it hardly follows that it *is* a simulation. For example, we might have

[3]Note, however, that unanalysability is not a part of anyone's *definition* of Monte Carlo methods; it's just that analysable integrals are analysed instead!

[4]We should like to point out that, despite our differences on some particular issues, and especially the definition of simulation, Frigg and Reiss (2008) present a parallel argument to our own, in particular advancing our shared claim that the epistemology of simulation is the epistemology of experimentation.

function `state(cond,i)` which returns the simulated process's i-th state, given initial condtions `cond`. This function contains all the information contained in the simulation; indeed, by iteration it could be used to run the simulation. However, we can equally well use it to run all sorts of processes which are *not* the simulation, for example, the states indexed by the Fibonacci sequence.

For a more homely example, a similar point can be made about a feature film sliced into individual frames and put in an album: the album is not the feature film. The album contains all and only the information within the movie. But a movie moves, an album does not; and the album will remain not-movie until someone splices it back together. The methodologically relevant point is that one can poke a computer process and *then* see what happens. That is at least part of the point when people note that simulations embody aspects of experimentation. But if the process is already completed, one cannot poke it. There is no experimental side to things, even if the computer program incorporates all the information of the original simulation. Since the information is in there, presumably there is some way of extracting the same information as one would from experimenting with a simulation; but it would not be by some intervention which mimics experimentation. One might well say, along with Frigg and Reiss (2008), that since all the information is there, none of this matters. But keeping *some* connection with ordinary language and ordinary semantics is necessary, and calling a photo album a movie is just silly.[5]

2.3 Artificial Life Simulations

A final, and we hope decisive, objection to Definition 1 is that there have arisen very large regions of simulation research which are not plausibly described as equation solving at all, covering at least the vast bulk of agent-based modeling in artificial life and social simulation (which is our primary area of simulation research) and individual-based modeling in ecology (Grimm and Railsback, 2005). Although some equations will inevitably describe some characteristics of such simulations, it is at most *unusual* for the solution of equations to be the motivating factor in such investigations. The motivation is more typically the investigation of high-level properties of the system which emerge from an explicitly defined lower level of simulation. Some example motivations are:

- Demonstrating a feasible mechanism for the Baldwin effect in evolution (Hinton and Nowlan, 1987)
- Showing that flocking behavior can result from independent decision-making throughout a flock (Reynolds, 1987)
- Determining the minimal space requirements for beech forests to survive

[5]For a final analogy, you might consider Hans Moravec's proposal for life-extension: downloading the information content of your brain into a disk and "waiting" for technological development to support "your" reanimation. We suggest the incredulity this idea induces in most people is simply rational.

in isolated patches (Grimm and Railsback, 2005, section 1.2.2 & section 6.8.3)

- Finding conditions supporting or undermining the main postulated mechanisms for the evolution of dimorphic parental investments in offspring (Mascaro et al., 2005)

- Investigating the effectiveness of different possible public health interventions in response to a smallpox epidemic (Eidelson and Lustick, 2004)

If the philosophy of simulation is not to be left behind by the science of simulation, Definition 1 must be abandoned. Therefore, we shall adopt Hartmann's definition, but rendering it more explicit.

2.4 Another Definition of Simulation

One immediate benefit of Hartmann's definition is that it rules out the virtual reality scenarios directly: since the human user (trainee) is a necessary ingredient, these are not computer simulations. We will nevertheless now drop the word "computer" and talk about simulation most generally, as this will help us understand the relation between the epistemology of simulation and the epistemology of experiment. Our proposed semi-formal rendition of Hartmann's mimicking account is:

Definition 2 *S is a* **simulation** *of P if and only if*

1. *P is a physical process or process type*

2. *S is a physical process or process type*

3. *S and P are both correctly described by a dynamical theory T containing (for S; parenthetically described for P):*

    - *an ontology of objects O_S (O_P) and types of objects $\Psi_i(x)$ ($\Phi_i(x)$)*
    - *relations between objects $\Psi_i(x_1, \ldots, x_n)(\Phi_i(x_1, \ldots, x_n))$; hence, there are states of the system, s*
    - *dynamical laws of development (possibly stochastic): $f_S(s) = s'$ ($f_P(s) = s'$)*

In other words, both the simulation S and the target of the simulation P are physical processes with a common dynamical theory T. Computer simulations are then easily defined as:

Definition 3 *S is a* **computer simulation** *of P if and only if*

1. *it is a simulation of P*

2. *and it is a computer process or process type.*

An immediate objection to Definitions 2 and 3 might occur to you. It is symmetric: according to this, we could just as well use the sun to simulate our astrophysical programs as vice versa! This clearly won't do; our simulations, whether computer based or not, are surely *intrinsically* simpler than their targets of study. To use a metaphor of Giere's (1999) (borrowed from Borges (1954)), if we were to construct a map of the earth on a 1:1 scale, it's true that we could more accurately measure distances using this very fine resolution map than using cruder maps, but, obviously, all the other advantages of maps would be lost. (Borges' characters start shifting domicile from reality to map!) The problems with symmetry are both practical and theoretical. Theoretically, whatever one's view may be about the nature of scientific explanation and theories, it's entirely clear that they some how *summarize* features of the world. Computer scientists would say they *compress* information about the world. In short, they are shorter than any direct, exhaustive description of their objects.

All of this is well taken, but it doesn't follow that we need to acknowledge the point formally, within our definition. Simulations are typically constrained by both a lack of understanding of fine details of the objects of our simulations and by a lack of time to wait for the implications of fine details to filter through our simulations. Frequently, however, crude simulations are made less crude, as advances are made in both our understanding of the physical systems and in our computational capacities. If we were somehow to extend this advance in resolution power indefinitely, we might begin to approach the 1:1 scale contemplated by Giere. Admittedly, going all the way would be pointless. However, that doesn't mean that by going all the way we would no longer be dealing with a simulation; a pointless simulation remains a simulation.

3 Homomorphic Simulation

So, for practical and theoretical reasons, we require that our simulations *not* be as detailed as the processes we simulate. Instead we require that:[6]

Proposition 1 *There should exist a* **homomorphism** *h from P to S.*

Definition 4 *A* **homomorphism** *h from P to S is a mapping $h : P \to S$ such that*

1. *For every object $x \in O_P$, $h(x) \in O_S$.*

2. *For every relation Φ, $\Phi(x_1, \ldots, x_n)$ is true of P iff $h(\Phi) = \Psi$ and $\Psi(h(x_1), \ldots, h(x_n))$ is true of S*

3. *For every state transition function f in P, $f(s) = s'$ iff $f_h(h(s)) = h(s')$ (or, for stochastic laws, the distributions over states should be identical)*

[6] For a similar account, see Norton and Suppe (2001), although their account is somewhat cluttered with a variety of idealized, averaged and approximate models.

The application of homomorphisms to simulation, to be sure, should be taken with a grain of salt. That is, it is an ideal and one which we are unlikely actually to reach with non-trivial simulations. It is frequently noted in the literature that our simulations often diverge from reality in small ways and sometimes in large ways. Nearly every simulation diverges in at least this way: digital computational processes cannot exactly simulate continuous time, whereas real systems at least appear to develop in continuous time; thus, these systems support relations ("in-between times") that have no counterpart in their simulations. Nevertheless, at least for most problems, time can be discretized to a fineness where this difference does not matter. The epistemological problem is to sort out when the divergences do matter to inferences about the real systems.

Our central epistemological proposal is that simulations can be tested for adequacy by testing whether a homomorphism between the real and the virtual system holds. In the simulation literature this is called "validation".

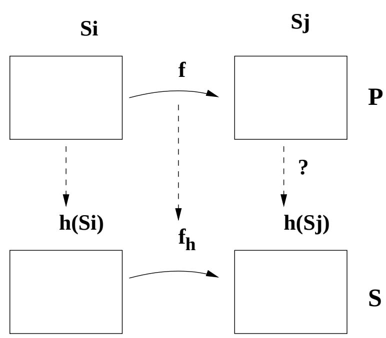

Figure 1: Testing for homomorphism.

3.1 Testing for Homomorphism (Validation)

By observing, or arranging for, the physical system P in state s_i and its subsequent transition to state s_j, we are enabled to test whether the simulation undertakes in homomorphic initial conditions the like transition (or vice versa). In Figure 1 this corresponds to checking that simulation S shows a transition from $h(s_i)$ to $h(s_j)$. Validation can be thought of as parallel to confirmation. Instead of confirming how well a theory represents reality, we are confirming how well a simulation maps reality. As such, validation comes in degrees, as there will be more or less severe tests possible for the adquacy of the mapping. Indeed, we would assert that the degrees come in the form of prior and posterior probabilities of the existence of a homomorphism, exactly as with ordinary confirmation theory, were we to allow ourselves the diversion into more traditional issues in the philosophy of scientific method.

(Grimm and Railsback, 2005, chap 9) present a likely account of how validation might proceed. They suggest first testing low-level submodels which describe non-emergent phenomena in the simulation and only subsequently looking at higher-level systems, including properties of the simulation that emerge from interactions between submodels. At the higher levels we are conducting simulated versions of controlled experiments (Grimm and Railsback, 2005, p.316): "We pose alternative theories for the individual behavior as the hypotheses to be tested, implement each hypothesis in the [simulation], identify some patterns as the 'currency' [standard] for evaluating the hypotheses, and then conduct simulations that determine which hypotheses fail to reproduce the patterns." For example, the beech forest simulation was designed to reproduce both the horizontal mosaic pattern of tree stands and the vertical pattern of tree cover. But subsequently unplanned for patterns in the simulation were discovered and put to good use (Grimm and Railsback, 2005, p.7; our emphasis):

> [The simulation] was so rich in structure and mechanism that it also produced independent predictions regarding aspects of the forest not considered at all during model development and testing. These predictions were about the age structure of the canopy, spatial aspects of this age structure, and the spatial distribution of very old and large trees. All these predictions were in good agreeement with observations, considerably increasing the model's credibility. *The use of multiple patterns to design the model obviously led to a model that was structurally realistic.*

Other than the fact that this procedure is dealing with a computer simulation rather than directly with an ecological theory, there is no interesting methodological difference between this and standard theory testing. A rich, multi-patterned simulation offers a variety of opportunities for testing its conformity to the target process. And just as in standard confirmation theory (Franklin, 1986, pp.123–129), the more varied the predictions of a simulation, or the submodels used to make them, that are tested and confirmed against reality, the greater our confidence that the simulation indeed maps that reality. With a sufficient variety of such tests, testing diverse transitions under diverse conditions, we may well be able to conclude that the simulation is, or is not, homomorphic, either approximately or exactly.

The existence of an approximate homomorphism is *crucial*: it underwrites the relevance of the simulation for the system being simulated and, in particular, its use both for explaining events in the real world and in predicting them.

As noted, homomorphisms may exist at a variety of levels of resolution. The level of resolution of the homomorphic simulation depends upon two major points:

1. How well do we (think we) understand P? How detailed a theory do we have to test?

2. Pragmatic constraints upon our simulation (e.g., how much time can we

spend waiting).

The levels of potential simulation may lie upon one-another Shrek-like, as in an onion (see Figure 2):
At the top-level is a simulation with such a small ontology that nothing useful

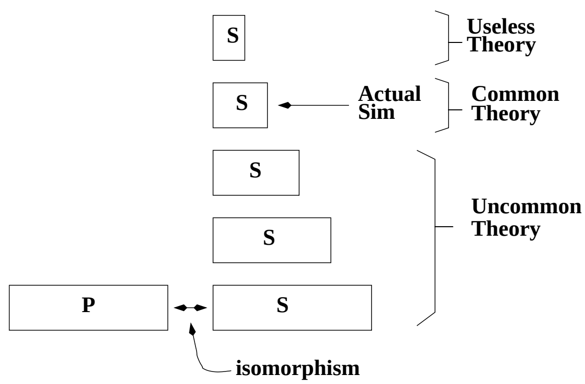

Figure 2: The simulation onion. (Smaller layers indicate fewer details, but greater generality.)

can be simulated. Below the actual simulation are more detailed potential simulations which are unused for such reasons as: the theory describing such detail has not been invented; the simulations at that level of detail are impractical; the level of detail describes events of no interest to us.[7]

4 Simulations as Experiments

Many people have been attracted to the idea that simulations have no empirical side to them and, in particular, that they are basically revved-up thought experiments. Oreskes and others claim that simulations cannot be used to acquire *any* empirical knowledge about the world, directly or indirectly (e.g., Oreskes et al., 1994; Axelrod, 1997; Di Paolo et al., 2000). Rather, simulations are limited to extending our understanding of the theories being simulated, by exploring their deductive consequences. Di Paolo, Noble and Bullock (2000), for example, examine the Hinton and Nowlan (1987) simulation study of the Baldwin effect — the acceleration of genetic evolution via learning by individuals in a population (Baldwin, 1896). Prior to that study, the Baldwin effect had been given little attention; it sounded too much like a Lamarckian process and the mechanism for fixing learned behavior genetically was not understood. The simulation of Hinton and Nowlan (1987) changed that by producing a plain, easily inspected mechanism, which demonstrably exhibited

[7]In other words, this kind of account is very far from requiring the "perfect mimesis" of isomorphism that Winsberg (2003, p.116) claims is implied.

the Baldwin effect. Di Paolo et al. (2000) argue that this is essentially a theoretical, deductive use of simulation, making plain what was implicit in the theory. While it is clear that simulations can be used to explore the deductive consequences of theories, it is not clear that that is the only role they may have in empirical science. Nor is it clear that Hinton and Nowlan's *mechanism* was in any sense implicit in Baldwin's theory. We now proceed to argue that they have potentially every role that experiments may have in empirical science.

4.1 A Comparison with Real Experiments

To further our claim that simulation studies share epistemology with traditional scientific experiments, we can consider Allan Franklin's experimental strategies (Franklin, 1990). Franklin emphasizes that his strategies are neither exclusive of other strategies nor exhaustive. Nevertheless, they provide a good indication of what happens in physical experiments; we annotate the list with reference to simulation studies (Franklin, 1990, p.104):[8]

1. Experimental checks and calibration, in which apparatus reproduces the known phenomena

2. Reproducing artifacts that are known in advance to be present
   Regarding 1 and 2, reproducing known phenomena is a standard check of adequacy in simulation studies.

3. Intervention, in which the experimenter manipulates the object under observation
   The relative ease of manipulating simulations is one of their key advantages in experimental studies.

4. Independent confirmation using different experiments
   In simulation research there is always an opportunity to test very different kinds of initial conditions, and sometimes an opportunity to test the operation of distinct subprocesses (Grimm and Railsback, 2005; Grimm et al., 2005). Replication of simulation results using distinct simulations is also a possibility (e.g., Axtell et al., 1996; Edmonds and Hales, 2005).

5. Elimination of plausible sources of error and alternative explanations of the result
   These are activities integral to both verifying and validating simulations.

6. Using the results themselves to argue for their validity
   By this Franklin meant that an experiment may create results which are highly unlikely to be artifacts of the measurement process or experimental procedure and so by themselves support the claim that they reflect an external reality. Similarly, simulation results may likewise be determined

[8] We have corrected Franklin's "corroboration" with "confirmation".

to be highly unlikely to be due to bugs, not just because of steps in the verification process, but also because of the results themselves.

7. Using an independently well confirmed theory of the phenomena to explain the results
   This is one leg of our triangle of Figure 3 below.

8. Using an apparatus based on well confirmed theory
   The apparatus here is the simulation, and associated software; verification is part of the process of justifying the claim that it is based on well confirmed theory.

9. Using statistical arguments
   It has been frequently remarked that the use of, or rather the need for, automated data analysis and data visualization techniques is a striking feature of simulation research. Epstein and Axtell (1996), for example, employ a variety of graphics to good effect.

Clearly, at a phenomenological level, simulation research is very akin to traditional experimental research. But this does not demonstrate that at an "epistemological level" they are again alike.

4.2 The Epistemology of Simulation

There are two acknowledged steps to justifying claims that a simulation is informative of the real world:

Verification: Determine whether the simulation correctly implements the theory being investigated,

- e.g., by performing design verification, debugging the simulation, and consistency checks.

Validation: Determine whether the simulation as implemented conforms to the target process.[9]

- This is testing for the existence of a homomorphism, by comparing simulation results with the target process and vice versa.

These steps are portrayed graphically in Figure 3. A theory T has been developed for the type of physical process P.[10] A range of token processes, $P_1, \ldots, P_n$ instantiate that type. And a simulation has been developed for those processes and/or the process type, leading to simulation processes $S_1, \ldots, S_m$. We can think of the computer program used to launch the simulation processes as a process type S, not depicted in the figure. This situation presents

[9]From a history of the philosophy of science perspective, these uses of "verify" and "validate" are backwards. However, this usage comes from the software engineering tradition.

[10]Incidentally, we are favorably inclined towards the semantic interpretation of scientific theories (Suppe, 1977), but nothing in our account hangs directly upon that.

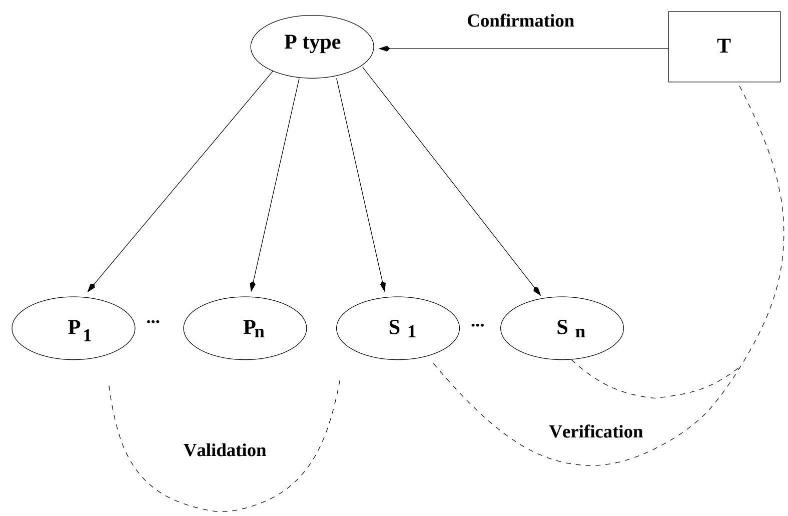

Figure 3: The verification, validation and confirmation triangle.

us with a triangle with three legs of possible justificatory test for the relevance of simulation S to its target P: whether the theory T represents the reality P (confirmation); whether S properly implements T (verification); whether S corresponds to P (validation). Each test can be carried out independently of the others. Once any two tests have been conducted, and assuming their outcomes are not in dispute, then the third test becomes irrelevant, since we then know everything there is to know about the relations between T, P and S. This explains a range of observations previously made about simulation. Given verification, S can be manipulated to investigate the implications of T (this is the "S as thought experiment" above). Again, given verification, determining that there is a homomorphism between P and S is tantamount to confirmation; given validation, S can be used in exploratory theorizing as well as predicting the consequences of intervening in P systems; given confirmation, failures of correspondence between P and S indicate verification failures. When any two justificatory steps have been successful, the S_k are just as much instantiations of P as are the P_i: they all provide supervenience bases for that type of physical process.

4.3 Experiments as Simulations

An obvious rebuttal to the claim of epistemological sameness between experiments and simulations is: "Unlike simulations, when you're testing the real-world at least you know what you're testing is *real*! You can't be testing the wrong thing!" Morgan (2002), for example, claims there is an inferential gap between simulation and reality that doesn't exist between real-world experiments and real-world target systems. This is a seductive thought, but it is wrong. The inferential gap is always there in any scientific study. It is nearly an everyday occurrence to hear about some medical study in which the experimental groups turn out to be unrepresentative of some target population. And it is an old joke that experimental psychology has accumulated a large body of evidence about the psychology of college students.

We previously observed that, following our Definition 2, the targets of our computer simulations — the physical processes in the world — might be construed as simulations (but not *computer* simulations) of our computer processes, even if that is not pragmatic. The general point is that some (non-computer) physical processes may be used to simulate others. For example, scale models built with clay and water are used to assess water flow and tidal action; or, wind tunnels are used to assess aeordynamic flows. Again, of course, the simpler and smaller processes are generally said to be simulations of the more complex and larger processes which are the targets of investigation. Which is the simulation is, in the first instance, driven by which process is wanting to be understood. But other factors play a role, including accessibility, ease of intervening in the process, and the ethics of intervening in the process. If there is a single process, and not a process *type*, of interest, and if it is accessible, etc., then there may be no recourse to a simulation: a simple intervention in the target process may be attempted. There is then no inferential gap, because the studied system is also the target system. Engineering applications and most treatments by medical doctors are of this type. If our interest spans an entire type of process, and if we are not looking for a specific outcome other than learning about that type, then a direct examination of all instances of the type is unlikely to be possible. In that case, we shall have to simulate the process type of interest with another that is more accessible. For example, we might use a sample of adult humans in Australia today to simulate the category of all adult humans across all of time and space. This is the common practice in medical research. Similarly, we might be interested in physical conditions immediately after the Big Bang, but have little prospect of directly measuring them; instead, we might simulate such conditions using high-intensity collisions of subatomic particles. Again, Galileo famously tested his telescopes on terrestrial objects to gain support for his inferences concerning celestial objects (Chalmers, 1982, p.72). While there are many complications and nuances to each of these stories, in all these cases and all such experiments we are using one physical process, or type of process, to simulate another. And, in all such experiments, there is the same potential for things to go wrong, for the experiment to be uninformative, because the experimental subject fails to simulate the target subject, because the homomorphism between reality and experimental process fails.

In short, "studying proximate systems as stand-ins for target systems of interest . . . pervades all science" (Frigg and Reiss, 2008). The epistemology of computer simulation is the same as the epistemology of experimentation for the simple reason that all experiments are simulations and computer simulation experiments are just a special type.

4.4 Special Epistemology

Regardless of our arguments and proposed interpretation of the epistemology of simulation, various philosophers of science have claimed that computer simulation is a new methodology demanding a new epistemology (e.g., Winsberg, 1999, 2003, Humphreys, 2004, p.54). Computer simulation is certainly a new methodology, involving new tools and techniques. Expertise in experimental physics is hardly interchangeable with expertise in physics simulation. But what additional reasons have been advanced for demanding a new epistemology? Some of the features of simulation said to require new epistemological thinking are:

- *Visualization.* Coupling computer simulations with visualization, including animations, is very common; indeed, computer simulation is being used specifically to create artistic animations (e.g., McCormack, 2005). Humphreys (2004, pp.111–114) seems to think that the use of visualization is a defining characteristic of simulation, which is clearly going too far. For example, in our discussion below of our simulation of parental investments, we have not found it necessary to include any graphics or imagery, nor would our argument be weakened had we never produced any. Regardless, Frigg and Reiss (2008) are surely right that the importance of visualization in coping with massive amounts of data is a property that simulation shares with experimentation.

- *Approximation.* All computer simulations, short of simulations isomorphic to the target process, of course, are approximate. But this again is no unique property of computer simulation. There are many examples of the use of physical experiments which are known to distort the properties of target systems. Wind tunnels are used to investigate aerodynamics, however their walls introduce "unnatural" turbulence which can affect the process under study (Norton and Suppe, 2001). Some such distortion is true of any scale model and, more generally, of any physical system not strictly identical to the target system.

- *Discretization.* Something which may well be distinctive about computer simulations is that they are housed within discrete von Neumann machines. We've already mentioned that time must be represented as a sequence of time steps, whereas real processes do not step through time. More generally, any state variable must be represented with some finite degree of precision, implying misrepresentation of real numbers. However, once again, this point of potential epistemological concern generalizes to experimental apparatus quite generally: although in a real-world experiment both the experimental and target processes may well both be continuous processes, the experimenter will have no way of taking advantage of that (in either manipulation or observation) beyond some finite degree of error.

- *Calibration.* Typically simulations have various parameters that need to be calibrated so as to reproduce known phenomena of the real system. For example, in a simulation of the evolution of group selection the strength of altruistic behavior needed to be adjusted to produce a stable population of altruists (Appalanaidu, 2007). This might suggest that you can get whatever result you want by recalibrating your simulation. It's implausible that science constructs reality according to its wishes, despite the more extreme views of social constructivists, but it's far more plausible that computer scientists can construct virtual reality according to their wishes. We accept that there is some danger here; the flexibility of universal computation can cover many faults of theory, if allowed to do so. However, we again assert the parallelism between simulation study and experimental study: given the validation of structural properties of a simulation, the calibration of parameters of the simulation can only push the results so far and not infinitely far. Such calibration serves the identical purpose with calibration in physical experiment, that of finding the settings which support previously observed measurements of a target system under given initial conditions, and so supporting the claim that measurements under new conditions will be informative.

5 Example Simulation: The Evolution of Parental Investment

We would like to illustrate some of the issues that have arisen with an artificial life simulation. Such simulations have received considerably less attention in the philosophical literature than Monte Carlo estimation, yet they are becoming much more prominent in the work of biologists and social scientists, usually under the rubric "agent-based" or "individual-based" modeling. Here we discuss an example of evolutionary ALife simulation, one of a number we have used to investigate issues in the theory of evolution which have been contentious and which resist any easy recourse to ordinary experimental test (Mascaro et al., 2005). Because of the contentiousness of these issues and because of the complexity of the real systems being hypothesized about, it is not likely that these simulations can soon resolve the problems addressed. However, it should be clear that there is potentially a rich field of patterns to validate such homomorphisms as we can create.

Our simulation is an agent-based evolutionary ALife simulation. This means that there are agents with a phenotype (behavior) arising from an interaction between a heritable genotype and their environment. The agents' behavioral repertoire includes reproduction, when suitable mates are present and when their level of health is sufficient. Offspring are created by chromosomal crossover and mutation and receive an initial donation of health from each parent. The health of agents is a function of their foraging abilities, their levels of activity and how much health they donate to offspring. When an agent's health drops to zero, the agent dies. The agent will die regardless when it reaches its maximum life span.

We used this simulation to investigate various hypotheses proposed to explain the evolution of dimorphic sex-linked traits and, in particular, why parental investments differ between males and females, with females characteristically (but not universally) investing more in their offspring than males (Trivers, 1972). In order to do this, the amount of health investment in offspring was a heritable, sex-linked trait in our simulation; we also simulated variable gestation terms. (For further details of our simulation see Mascaro et al., 2005.)

5.1 Hypotheses

We designed simulation experiments to test three widely discussed explanatory hypotheses for the evolution of differing parental investments:

1. *Concorde hypothesis:* The sex that makes the greater initial investment has the more to lose and is thus the sex more likely to evolve further investment (Trivers, 1972). Thus, differential investments may arise by chance and then be fixed by subsequent evolution.

2. *Desertion hypothesis:* The sex which has the first chance to desert an offspring (leaving it with the other parent) will do so (Dawkins and Carlisle, 1976). For example, in many fish it is the male who looks after the offspring; since males must wait for females to spawn their eggs before fertilizing them, females have the first opportunity to desert. Since fertilization occurs inside the female in mammals, males have the first opportunity to desert and, indeed, generally invest less.

3. *Paternal uncertainty hypothesis:* The sex which is less certain of being the parent of an offspring will invest less in its (apparent) offspring, particularly in species where females go through a gestation period (Trivers, 1972).

5.2 Experiments

5.2.1 Concorde Hypothesis

Dawkins and Carlisle (1976) attacked Trivers' first hypothesis above, suggesting it involves fallacious reasoning of the sort used to defend continued spending on a project based on sunk costs rather than on future potential. They used the then topical example of continuing government spending on the Concorde supersonic airliner, "justified" by prior extravagent waste on the project. It seems unlikely that the forces of evolution, unlike the forces of government, would succumb to such fallacious reasoning, rather than responding to the better supported principle of evolving traits to maximize expected fitness. As we agreed with Dawkins and Carlisle's reasoning, we fully expected our simulation of the Concorde hypothesis to fail to establish or sustain dimorphic parental investments.

It's clear that our basic simulation must favor some combined level of investment from both parents, since investment is necessary for offspring to achieve a level of health required for reproduction and since over-investment will be punished by an inability for the parents to further reproduce. However, the basic simulation offered no advantage to one sex or the other for differential investment, as is clear from the random walk in the relative size of sex-linked investments that results from running it.

In order to test the Concorde hypothesis, the simulation needed to be set up with the initial condition that a randomly selected sex starts with a higher investment than the other. Otherwise, the basic simulation was unchanged. The result was that all parental investments rapidly converged on 1/2 the level of combined investment optimal for the basic simulation. As the simulation incorporates all and only the basic ingredients of evolution, two sexes and parental investments, this seems to be a clear experimental refutation of Trivers' conjecture; alternatively, assuming the falsity of that conjecture was not in doubt, this is a clear experimental support for the adequacy of our basic simulation.

5.2.2 Desertion Hypothesis

To test the desertion hypothesis we allowed parents to invest health for an evolvable period after birth, contingent upon their maintaining contact with their offspring. The child needed a minimum total investment period; if one parent quit investing before that minimum, the other parent was forced to make up the difference as required by the offspring. Furthermore, females had a fixed minimum investment period in addition to an evolvable variable period, making up their total investment period. Males only had an evolvable variable period, allowing them an opportunity to desert first. Thus, if the hypothesis is correct, then average investment periods should lengthen for females and shrink for males, when they are initialized to be roughly equal. In most of our simulation runs there was a clear diminution of investment periods for males, although they did not drop to zero; female investment periods were reliably sustained by evolution well above the fixed minimum. Under some circumstances male and female investments would not diverge, such as when the females could not make up the difference or when agent mobility was reduced, so the opportunity to desert failed to arise. In summary, our simulation appears to have supported Dawkins and Carlisle's hypothesis that an earlier opportunity to desert combined with a fixed minimum amount of parental investment will result in sexual dimorphism.

5.2.3 Paternal Uncertainty Hypothesis

We tested the paternal uncertainty hypothesis by fixing the probability of paternity as a parameter of the simulation. Mothers always invested in their own offspring. However, males were chosen by the mother to invest in her offspring, according to the probability of paternity. In other words, if the paternity probability was set at 1, then the female always chose the true father for investment; if 0, then the female never chose the true father; and otherwise she chose males randomly, but with a probability of paternity fixed by the simulation parameter.

We're sure the reader can anticipate our results at this point. In simulations where the paternity parameter was 0, health investments by "fathers" were largely altruistic, and the investment level evolved downwards.[11] On the other hand, the parameter 1 resulted in maintaining high levels of investment and intermediate parameters resulted in intermediate levels of investment evolving.

5.3 Parental Investments

None of our experimental work with parental investment is particularly profound. We implemented a basic, straightforward simulation framework suitable for testing theories about parental investment, and then we implemented straightforward, clear mechanisms for each hypothesis, enitrely in accord with Grimm and Railsback's account of simulation as an experimental inquiry. The result is we have two candidate explanatory mechanisms and one which is unviable. There are, of course, other conceivable explanations for the evolution of dimorphic investments, which can be tested likewise. And moving on from demonstrating viable mechanisms for evolution to *asserting* the correctness of one or another hypothesis as an explanation for any actual case of dimorphic investments will require, minimally, investigations establishing the existence of a corresponding (homomorphic) mechanism in the real system and the non-existence of other processes which would negate or overwhelm it. So, what we have done with our simulations here is modest: we have made clear which of these three mechanisms is *capable* of being incorporated in legitimate explanations, in circumstances reasonably close to those we have simulated.

6 Conclusion

The experimentation with computer simulations that has become a prominent feature in the sciences is more than experimentation in name only. It is full-blooded experimentation. It carries problems, techniques and methods which are clearly new, such as debugging methods. It carries with it problems,

[11] Since mobility is a factor in our simulations, there was in fact kin selection pressure in favor of sustaining some level of investment. That is, even though the designated fathers were never the true fathers, they were more likely to be near relatives than unselected (because distant) agents.

techniques and methods which are old, as well, such as figuring out which statistics to capture to obtain an informative view of what is happening. None of the issues raised in this paper to this point actually identify any interesting limitation on computer simulation or the need for any new epistemology. The difficulties with sorting out the epistemology of experimental science are not yet adequately resolved; but there is no reason to believe that that epistemology won't have rich enough resources to accommodate what scientists are today doing with their computers.

The limits of computer simulation are, thus far, the limits of Turing computation. We know some of what lies beyond those limits, what has been called hypercomputations (Copeland, 2002), computations which, for example, infinitary machines can deal with. Despite the pessimism of Dreyfus, Penrose and Humphreys (Dreyfus, 1992; Penrose, 1999; Freedman and Humphreys, 1999), and many others, however, we have been given no reason to believe that human mental capacities are beyond the capabilities of Turing computability. In consequence, the prospects for the "ultimate" simulation, that of the human brain — completing the practical goal of producing a genuine artificial intelligence, are very real, if also rather distant.

References

Appalanaidu, C. (2007). Group selection: An investigation into the potential for the evolution of virtual ecosystems. Clayton School of IT, Monash University. Honours Thesis.

Axelrod, R. (1997). *The Complexity of Cooperation: Agent-Based Model of Competition and Collaboration*. Princeton, NJ: Princeton University Press. Suggests that agent-based modelling (not specifically simulation) is used to substitute for deduction, when formal mathematics can not generate appropriate conclusions and is a kind of intuitive technique.

Axtell, R., R. Axelrod, J. M. Epstein, and M. D. Cohen (1996). Aligning simulation models: A case study and results. *Computational and Mathematical Organization Theory 1*, 123–141.

Baldwin, M. J. (1896). A new factor in evolution. *The American Naturalist 30*, 441–451.

Borges, J. (1954). *Historia Universal de la Infamia*. Buenos Aires: Emece.

Chalmers, A. F. (1982). *What is this thing called science?* St. Lucia, Qld: University of Queensland Press.

Copeland, B. J. (2002). Hypercomputation. *Minds and Machines 12*, 461–502.

Dawkins, R. and T. Carlisle (1976). Parental investment, mate desertion and a fallacy. *Nature 262*, 131–133.

Di Paolo, E. A., J. Noble, and S. Bullock (2000). Simulation models as opaque thought experiments. In M. A. Bedau, J. S. McCaskill, N. H. Packard, and S. Rasmussen (Eds.), *Artificial Life VII: Proceedings of the Seventh International Conference on Artificial Life*, Cambridge, MA, pp. 497–506. MIT Press.

Dreyfus, H. L. (1992). *What Computers Still Can't Do: A Critique of Artificial Reason* (3rd ed.). Cambridge, Mass: MIT Press.

Edmonds, B. and D. Hales (2005). Computational simulation as theoretical experiment. *Journal of Mathematical Sociology 29*, 209–232.

Eidelson, B. M. and I. Lustick (2004). Vir-pox: An agent-based analysis of smallpox preparedness and response policy. *Journal of Artificial Societies and Social Simulations 7*(3).

Epstein, J. M. and R. Axtell (1996). *Growing Artificial Societies: Social Science from the Bottom Up*. Cambridge: MIT Press.

Franklin, A. (1986). *The Neglect of Experiment*. Cambridge University Press.

Franklin, A. (1990). *Experiment, right or wrong*. New York: Cambridge University Press.

Freedman, D. and P. Humphreys (1999). Are there algorithms that discover causal structure? *Synthese 121*, 29–54.

Frigg, R. and J. Reiss (2008). A critical look at the philosophy of simulation. *Synthese*. forthcoming.

Giere, R. (1999). Using models to represent reality. In L. Magnani, N. Nersessian, and P. Thagard (Eds.), *Model-Based Reasoning in Scientific Discovery*, pp. 41–57. New York: Kluwer.

Grimm, V. and S. Railsback (2005). *Individual-based Modelling and Ecology*. Princeton: Princeton University Press.

Grimm, V., E. Revilla, U. Berger, F. Jeltsch, W. M. Mooij, S. F. Railsback, H. Thulke, J. Weiner, T. Wiegand, and D. L. DeAngelis (2005). Pattern-oriented modeling of agent-based complex systems: Lessons from ecology. *Science 310*, 987–991.

Hartmann, S. (1996). The world as a process: Simulation in the natural and social sciences. In R. Hegselmann, U. Müller, and K. Troitzsch (Eds.), *Modelling and Simulation in the Social Sciences from the Philosophy of Science Point of View*, pp. 77–100. Kluwer.

Hinton, G. E. and S. J. Nowlan (1987). How learning can guide evolution. *Complex Systems 1*, 495–502.

Humphreys, P. (1991). Computer simulations. In *Philosophy of Science Association 1990*, Volume 2, pp. 497–506.

Humphreys, P. (1993). Numerical experimentation. In P. Humphreys (Ed.), *Patrick Suppes: Scientific Philosopher*, Volume 2. Dordrecht: Kluwer.

Humphreys, P. (2004). *Extending Ourselves: Computational Science, Empiricism, and Scientific Method*. Oxford: Oxford University Press.

Kueppers, G., J. Lenhard, and T. Shinn (2006). Computer simulation: Practice, epistemtology, and social dynamics. In J. Lenhard, G. Kueppers, and T. Shinn (Eds.), *Simulation: Pragmatic Construction of Reality*. Dordrecht: Springer.

Mascaro, S., K. B. Korb, and A. E. Nicholson (2005). An alife investigation on the origins of dimorphic parental investments. In H. A. Abbass, T. Bossamaier, and J. Wiles (Eds.), *Proceedings of the Australian Conference on Artificial Life*, pp. 171–185.

McCormack, J. (2005). A developmental model for generative media. In M. S. Capcarrere, A. A. Freitas, P. J. Bentley, C. G. Johnson, and J. Timmis (Eds.), *Proceedings of the 8th European Conf. on Advances in Artificial Life*, pp. 88–97.

Morgan, M. (2002). Model experiments and models in experiments. In L. Magnani and N. J. Nersessian (Eds.), *Model Based Reasoning: Science, Technology, Values*, pp. 41–58. Springer.

Norton, S. D. and F. Suppe (2001). Why atmospheric modelling is good science. In C. Miller and P. Edwards (Eds.), *Changing the Atmosphere: Expert Knowledge and Environmental Governance*, pp. 67–105. Cambridge, MA: MIT Press.

Oreskes, N., K. Shrader-Frechette, and K. Belitz (1994, February). Verification, validation and confirmation of numerical models in the earth sciences. *Science 263*(5147), 641–646. 2003-02-17 A fairly strong position against simulation contributing to knowledge.

Penrose, R. (1999). *The Emperor's New Mind : Concerning Computers, Minds, and the Laws of Physics* (2nd ed.). Oxford: Oxford University.

Pritsker, A. A. B. (1979). Compilation of definitions of simulation. *Simulation 33*, 61–63.

Pritsker, A. A. B. (1984). *Introduction to Simulation and SLAMII*. John Wiley & Sons.

Racynski, S. and A. Bargiela (2007). *Modeling And Simulation: Computer Science of Illusion*. Research Studies Pr.

Reddy, R. (1987). Epistemology of knowledge-based systems. *Simulation 48*, 161–170.

Reynolds, C. W. (1987). Flocks, herds, and schools: A distributed behavioral model. *Computer Graphics 21*, 15–34.

Rohrlich, F. (1991). Causal discovery via MML. In *PSA 1990*, Volume II, pp. 507–518.

Suppe, F. (1977). *The Structure of scientific theories*. Urbana: University of Illinois Press.

Trivers, R. L. (1972). Parental investment and sexual selection. In B. Campbell (Ed.), *Sexual Selection and the Descent of Man*, pp. 136–179. London: Heinemann.

Winsberg, E. (1999). Sanctioning models: The epistemology of simulation. *Science in Context 12*, 275–292.

Winsberg, E. (2001). Simulations, models, and theories: Complex physical systems and their representations. In *Proceedings of the 2000 Biennial Meetings of the Philosophy of Science Association (Supplement to Philosophy of Science)*, Volume 68, pp. S442–S454.

Winsberg, E. (2003). Simulated experiments: methodology for a virtual world. *Philosophy of Science 70*, 105–125.

Zeigler, B. (1976). *Theory of Modeling and Simulation*. New York: Wiley-Interscience.