Computational modelling, explicit mathematical treatments, and scientific explanation

John Bryden and Jason Noble
School of Computing, University of Leeds, LS2 9JT, UK

Abstract

A *computer simulation model*, can produce some interesting and surprising results which one would not expect from initial analysis of the algorithm and data. We question however, whether the description of such a computer simulation modelling procedure (data + algorithm + results) can constitute an explanation as to why the algorithm produces such an effect. Specifically, in the field of theoretical biology, can such a procedure constitute real scientific explanation of biological phenomena? We compare computer simulation modelling to explicit mathematical treatment concluding that there are fundamental differences between the two. Since computer simulations can model systems that mathematical models can not, we look at ways of improving explanatory power of computer simulations through empirical style study and mechanistic decomposition.

Introduction

It seems possible that computer simulation modelling could become the new modelling paradigm in biology. As modellers build transparent, tractable, computer simulation models their relaxed assumptions will, in comparison with traditional explicit mathematical treatments, make for considerably more realistic models that are close to the data. The 'Virtual Biology Laboratory' is proposed (Kitano et al., 1997) where a cycle is applied through comparing computer models with empirical evidence: the results from each procedure inspiring the direction of the other. Animals, such as *C. elegans*, have been well studied using computational models, e.g., (Bryden and Cohen, 2004). Indeed the formation of a complete model of the organism has been identified as a potential grand challenge for computing research (Harel, 2002). However, a full exploration of the relationship between mathematical and computational models in biology has not yet been achieved. Questions remain: for instance, whether both forms of modelling can peacefully coexist, whether mathematical models should aspire to the complexity of computational models, and conversely whether computational models can ever be as precise as a mathematical treatment.

In this paper we are mainly concerned with the scientific modelling of biological systems, however we hope that the findings can be applied more generally. Biological systems are made up of many different subsystems at different levels. Alife models often reside in the interface from one level to the next and can become extremely complex, especially as entities from any level can interact with entities from other levels.

The discipline of computer simulation modelling allows modellers previously unheard-of freedom to build and understand systems of many interacting parts. This new expressive freedom appears to have the potential to become the new modelling paradigm in science, perhaps overriding traditional techniques which use explicit mathematical treatments. However, this freedom does not come without a cost: as more and more detail is added computer simulation models can quickly become unwieldy and too complicated to understand.

How then can computer models contribute to the task of producing scientifically acceptable explanations? The use of a complex yet poorly understood model may be acceptable as some sort of loose analogy. However, Di Paolo et al. (2000) have argued that without a proper understanding of the internal workings of a computer simulation model, it can be impossible to say whether such a model makes a valuable contribution to the scientific problem it is addressing. They describe such problematic models as 'opaque thought experiments', arguing the need for explanations of the phenomena modelled. They suggest that modellers should use an 'experimental phase' in which manipulations are made to the computer model, the results of these manipulations hopefully generating insights into the workings of the system. Once the internal mechanisms are understood, the transparent model can then not only give new insights into the system being modelled but can also become a powerful predictive tool.

We question whether a computer simulation model can, in and of itself, constitute a scientific explanation. For example, one might produce a model in which individual organisms are explicitly represented and a particular population-level phenomenon appears to emerge. But this does not constitute an explanation of how entities from one level of a biological hierarchy produce interesting phenomena at another level. Di Paolo et al. (2000) argue that some explanation is required above a basic description of the model and the system it represents. In this paper we look further into what an adequate explanation of a model's mechanisms should entail. We will compare the account that we construct with the more basic position, sometimes seen in the artificial life literature, that a bare-bones description of a biological system with a computer model that qualitatively produces similar behaviour—with little or no extra analysis or explanation—can constitute a scientific explanation of some phenomena.

Given the above picture we must also consider the traditional methodology of explicit mathematical treatment. By explicit mathematical treatment we mean a model which is complete and contains no implicit steps, the steps can be logical statements and do not need to be formally written using mathematical symbols. While computer simulation models are fundamentally mathematical constructions, they, in the way they are reported, contain *implicit* mathematical steps rather than the *explicit* steps used by formal mathematical models. An explicit mathematical treatment takes logical axioms and specifies a number of clear explicit steps that deductively generate some result. In this paper we compare this traditional treatment with the new computational approach.

Firstly we set the context, we look at a framework for scientific modelling. Then, by looking at two examples of a similar system, we identify some properties that characterise an explicit mathematical treatment and which a computer simulation is unlikely to share. Having established that explicit mathematical treatment is the ultimate goal of any modelling enterprise, we look at how computer simulation models do indeed still have value. We look at how complex and unwieldy computer simulations may be simplified to more easily generate explicit mathematical treatments—proposing that this can be done by decomposition into simpler systems. Finally we set out, in an order of merit, the various different modelling approaches discussed.

A framework for scientific modelling

To understand how modelling is important and relevant within scientific investigation, we present a framework for scientific investigation with the scientific modelling cycle highlighted. Figure 1 presents a diagram of the framework.

The primary focus of scientific investigation is the building of a good *conceptual model* of the *real world*. Explanations of the real world reside in the conceptual modelling area of the framework, these are recorded in the *scientific corpus*. The basic scientific process involves the submission of concepts to the twin tests of *empirical science* and scientific modelling. The main focus of the framework, however, is on scientific modelling and the interface between a conceptual model and a *working model*.

Both computer simulation models and explicit mathematical treatments reside in the working model area of the framework. We take a working model to be a deterministic and completely specified model of a system. (Whereas a conceptual model may remain vague in places, a working model must be completely fleshed out.) Logical processes are applied to the axioms and the results of this process are recorded. Logical processes can include mathematical equations, logical deductions and computations. Working models produce *results* which are used to refine and update the conceptual model.

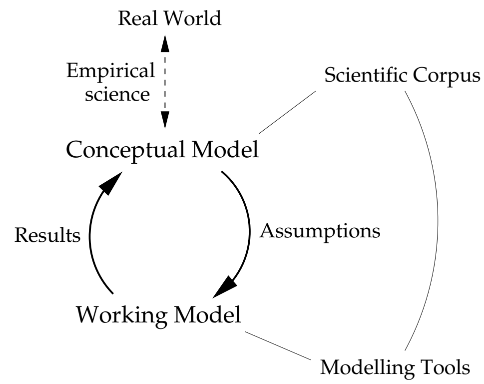

Figure 1: The cycle of enquiry in scientific modelling within the context of scientific investigation.

Before we specifically look at the sorts of results that can be generated by explicit mathematical treatments or computer simulation models, we discuss the types of *assumptions* that can be used to generate a working model. An assumption is essentially an abstraction from a more complex system. There will be many abstractions from the real world in the conceptual model (tested by empirical science) and it will normally be necessary to make further abstractions for ease of modelling. One of the main benefits of computer simulation modelling (Di Paolo et al., 2000) is that assumptions can be very easily added to or removed from models to see if they are significant or important. Explicit mathematical treatments tend to be more fixed in their assumptions. The types of abstractions used by either explicit conceptual models or computer simulation models can be distinguished into two groups, *reductionist* and *analogous* abstractions. We take inspiration for this distinction from Bedau's discussion of 'unrealistic' models (Bedau, 1999).

In order to highlight the important differences between the use of computational and mathematical techniques in building a working model, we must first consider the outcomes of a successful working model for the broader scientific project. The more valuable results generated by a working model will form some kind of explanation of why some phenomenon is present in the conceptual model. Other, less valuable, results include those that generate predictions. With an explanation generated by a model to hand, an empirical scientist can easily and quickly generate good empirical experiments to test whether an explanation is valid or not. A working model may indicate that some factors are more important than others for a particular phenomena. This may direct empirical science toward a more fruitful direction. The value of a result can depend not only on the type of working model used to generate the result, but also the assumptions used to generate the working model in the first place.

Competence and performance in scientific modelling

The previous section has set out the tasks necessary before embarking on a modelling enterprise: Once a conceptual model has been chosen that builds a picture of what is known about some real-world phenomenon, assumptions are then chosen to simplify this conceptual picture into logical units and axioms that can be built into a model. Up to this point everything is quite similar between the two logical modelling styles. Perhaps it is natural to assume that since both modelling techniques are analytical, the style of the results will also be quite similar?

To answer this question we must consider a thought experiment based on a specific example which can easily be understood and modelled by either a computer simulation model or an explicit mathematical treatment. The Lotka-Volterra system is a mathematical treatment of a predator-prey system. Two equations model the dynamics of the system:

$$\frac{dx}{dt} = Ax - Bxy \quad (1)$$

$$\frac{dy}{dt} = -Cy + Dxy \quad (2)$$

where x is the prey, y is the predator and A, B, C, D are constants. This system famously generates oscillations between the predator and prey populations. This mathematical treatment can be considered alongside an individual based computer simulation model of the same phenomenon.

A typical example system might be as follows. In a computer simulation model, individuals may have a location on a spatial grid moving at random each turn. If a prey individual encounters some food in its square it will receive an energy bonus, if it encounters a predator it will be eaten with the predator receiving an energy bonus. If either a predator or prey individual's energy level goes above a threshold then it will reproduce, and if any individual's energy level goes below a threshold, it will die.

Without wanting to go into too much detail, we assume, for the purposes of argument, that the computer simulation has very similar dynamics to the mathematical system. That is, both systems will make the same predictions about any particular predator-prey system to which they might be applied. The two systems can now be compared against each other and we can review our initial question concerning the nature of the scientific explanation that may be derived from each modelling enterprise.

To answer that question we draw on a distinction introduced by Chomsky between *competence* and *performance* (Chomsky, 1986). Chomsky's approach considers whether the linguistic corpus can be used as a source of empirical evidence for linguistic enquiry. He distinguishes between competence (our internal unconscious capacity for language) and performance (actual instances of language production). Regarding linguistic inquiry, he argues that we should take this distinction into account considering models of linguistic competence above models of linguistic performance.

We use Chomsky's distinction to shed light on the differing styles of scientific explanation that are likely to follow from the use of computational versus mathematical treatments of a particular problem. From this point of view, the computer simulation model must merely be considered as a performance of a scientific explanation, whereas the explicit mathematical treatment can be considered as having competence (an innate capacity) as a scientific explanation.

Simulation runs have the same sorts of problems as those Chomsky identifies for linguistic performances. They are subject to faults (in code as well as in run-time conditions) and each simulation model is merely a single data point and may not reveal the complete potential of a system. In a similar way, it is possible to hide flaws in the performance from the audience. Simulations can be set up so that the data points presented make the best possible case for whatever it is the modeller is trying to argue.

Alternatively, explicit mathematical treatments, assuming they are done correctly, are analytically complete: flaws in the system are immediately obvious. In addition, mathematical treatments are not limited to some narrow range of parameters but provide universal coverage of all variables included in the model. These two properties were identified by Chomsky as arguments in favour of looking at linguistic competence over linguistic performance.

Furthermore, explicit mathematical treatments have more powerful identity conditions than do computational models. By this we mean that one mathematical treatment can automatically be established as the same as, or different to, another treatment, just by comparing the logic. Computer simulation runs, on the other hand, may produce similar results for the same problem, but have very different underlying explanations. The opposite can also occur, in that two computer simulations may be driven by the same underlying process without this being obvious to an observer.

Mathematical treatments are more reusable than computer simulation models. Some give good clean results which can instantly be applied to systems, others benefit from the ease with which they can be written down in full and passed on. Such models can then be used as logical axioms for other models with their competence passed on. In contrast, although computer models can certainly be transferred from one author to another their results are rarely used, in practice, as axioms for other models.

One might argue at this point that we can distinguish the code for a computer simulation model from an individual execution of the code. The argument continues that a simulation run is merely a performance of the code, the code itself has competence. To answer this point we look at the style of computer model chosen in the Lotka-Volterra example above. It was chosen specifically so that the code would demonstrate an emergent phenomenon (Bedau, 1997). There are only two cases possible here. Either, without an execution of the code its macroscopic function is opaque, or, if the macroscopic function is deducible from the code, then this deductive process would necessarily form an explicit mathematical treatment. If this deductive process is impossible, any explanation generated must be teased out by analysis of simulation runs.

At this point, we are left with a conundrum. If computer simulation models are viewed as mere instances (performances) rather than as systematic explanations (having competence), how can they be of use to science? The answer is that there are many areas, identified especially in the AL-ife field, which do not yet yield to mathematical modelling but in which simulation models can already be produced. Such simulation models not only have scientific power as proofs of concept and for generation of insights for performing empirical science, but they can also have some explanatory power (Di Paolo et al., 2000).

When considering a complex simulation in which there is no explanation of the effects produced, some explanation can be deduced by performing experiments on the simulated system in the same way that one would do for an empirical investigation. In this mode of enquiry a *control* simulation is generated in which some important phenomenon does not happen. This is normally done through some manipulation of the system. The control simulation is compared with the untampered system and the results are used as evidence that the changes made by the manipulations are part of the explanation of the phenomenon.

The above procedure is very similar to the normal mode of empirical science. A conceptual model can be built of the working model system and this conceptual model acts as an explanation. We will now look further into how this form of explanation relates to an explicit mathematical treatment.

Analytic explanation versus synthetic explanation

To attempt to understand the difference between an explanation generated through the use of a working model in explicit mathematical form and an explanation generated by experimental manipulations of a computer simulation model, we consider a distinction used by the logical positivists—that of *analytic* and *synthetic* truths.

According to Frege's reworking (Frege, 1980) of Kant's original distinction, an analytic truth is one that can be deduced through logical laws alone. A synthetic truth is one which needs some other means, generally empirical investigation, to establish its truth or falsity.

We use this distinction to identify modes of truth for explanations generated by a working model. As pointed out previously, we assume all working models are using the same assumptions, i.e., they start from the same set of logical axioms. We distinguish between an analytic explanation—one which follows logically from the initial assumptions—and a synthetic explanation—one which must be determined by some other means.

Naturally an explicit mathematical treatment is in itself an analytic explanation. However, empirical experiments done on a computer simulation can only form synthetic explanations. These synthetic explanations require validation in the same way empirical science must be validated. The evidence backing up these validations relies on measurements taken from performances and is thus open to disconfirmation, reproduction and revalidation.

There is an ongoing debate about the analytic/synthetic distinction, some arguing that it is not a black and white distinction but more a question of degree (Quine, 1953). While Quine's arguments are concerned with statements about the real world rather than statements about a closed set of logical axioms, we agree that our distinctions of explanations should not be black and white. A working model can, like a biological system, be large and complex. Some parts of such a system will yield to explicit mathematical treatment, whereas with other parts we may have to rely on empirical-style experiments of the kind discussed by Di Paolo et al. The final explanation generated through such a process will consist of a mixture of analytic and synthetic statements.

In the next section we present an account of how systems can be decomposed into smaller parts to identify explicit mathematical treatments. Successful mathematical treatments will render the resulting explanations more analytic in the way we have just described.

Decomposition of systems

A system can become hard to analyse when it is made up from many inter-dependent subsystems. In fact, the identification of subsystems is a good first step when tackling such a complicated system. However, this is rarely simple. When subsystems are inter-dependent it is not possible to manipulate one subsystem independently without affecting another: both subsystems, at the same time, affect the overall system. The situation becomes increasingly difficult when the subsystem's components are not mutually exclusive from each other.

Simon (Simon, 1996) describes a 'nearly decomposable system' as being one in which components are independent in the short term, but dependent in the long term. This is a useful way to divide a system up and this has been expanded further (Watson, 2005; Polani et al., 2005) considering modular dynamical systems. Watson introduces a concept called *modular interdependence* to describe a system with modules that are decomposable but not separable. A hierarchy can be formed from subsystems and it is easy to see how complex behaviour can be generated. This hierarchical perspective is a valuable decomposition of a complex system. If it is possible to divide up a set of microscopic entities into subsets this will allow us first to tackle the mechanisms of the subsets, before understanding how they interact with each other.

In the next section we consider a more general perspective for decomposing systems. Rather than breaking up the set of microscopic entities into subsets, we consider a more arbitrary way of decomposing a system into subsystems that contain a simplified version of the dynamics of the supersystem.

Mechanistic subsystem

We propose information theoretic definitions of a mechanistic subsystem and interdependence in mechanistic subsystems. This style of definition has been used in (McGregor and Fernando, 2005) to formalise *hyperdescriptions*. We then go on to discuss how these definitions relate to our intuitive notions of these concepts before looking at examples in the next section.

Define a system S as being a set of mathematical entities, their interactions and their parameters. Take a descriptor function $d(S) = M$ that will map the system S to a set of descriptors M. Define the *entropy* of a random variable X as $H(X) = -\sum_{x\in X} p(x)\log p(x)$, the *conditional entropy* between two random variables X and Y as $H(Y|X) = -\sum_{x\in X} p(x)\sum_{y\in Y} p(y|x)\log p(y|x)$ and the *mutual information* as $I(X;Y) = H(Y) - H(Y|X)$.

Take a system S_1, such that $d(S_1) = M_1$. Then, S_1 is a mechanistic subsystem of S if

$$S_1 \subset S \quad (3)$$

$$H(M_1|M) = 0 \quad (4)$$

$$I(M_1;M) > 0 \quad (5)$$

$$I(M_1;M) < H(M)\,. \quad (6)$$

The mechanistic subsystem S_1 is a constrained version of its supersystem S. The constraints can take place in the parameter space, the number of entities, the nature of the entities, or their interactions. We list the Equations [(3) to (6)] and describe their meaning: (3) S_1 is a subset of S; (4) all information in M_1 is predicted by M; (5) M_1 and M share some information; (6) there is infomation in M that is not predicted by the information shared by M_1 and M.

The information theoretic definition presented includes many of the important concepts of a mechanistic subsystem. However a useful mechanistic subsystem should have two further properties. Firstly, it should be transparent, i.e., it is possible to understand why and how it produces its macroscopic effects. Secondly, its macroscopic effects should be of interest when compared to the macroscopic effects of the main system. We need to avoid specifying macroscopic subsystems that are either equally complex to the main system with only some neglible reduction, or are so simplistic that they are of no analytic value.

Following on from this definition of a mechanistic subsystem, we draw on Polani et al.'s definition (Polani et al., 2005) of a system that is decomposable but not separable to identify how two mechanistic subsystems can be interdependent. Take a system S and two mechanistic subsystems S_1 and S_2, the subsystems are interdependent if $0 < I(M_1;M_2) < \min[H(M_1), H(M_2)]$. The two subsystems are neither independent nor completely dependent.

With this approach identified, we can see how it is possible to break up a complex system of many interacting parts into simpler mechanistic subsystems.

Examples of Mechanistic Subsystems

We consider, as an example, the spatial embedding of reproducing agents. Space has been shown to be an important factor in the maintenance of cooperation in a population (Boerlijst and Hogeweg, 1991; Di Paolo, 2000). The common feature of these models is that two regimes are considered. The models are considered and analysed in a non-spatial environment before being placed in a spatial environment. The non-spatial treatment is a mechanistic subsystem of the spatial treatment. In this treatment agents are thought to be in a perfectly mixed spatial environment, a special case of the spatial component. A comparison of the interactions of agents in the spatial and non-spatial environments demonstrates how cooperation is increased. The mechanistic subsystem (the non-spatial model) functions primarily as a control in these experiments.

A different model (Bryden, 2005b) considers collective reproduction in amoebae. This non-spatial model demonstrates that cells that reproduce individually must reproduce more slowly to maintain high energy reserves for periods of low resources. By reproducing collectively during periods of low resources, individuals can avoid the need to reproduce slowly and can dominate periods of high resources by reproducing more quickly. The model is complex and it is not easily apparent why this is occurring. A mathematical treatment (Bryden, 2005a) analyses a mechanistic subsystem of the main model only considering individuals that reproduce individually. This treatment shows that, when there is a greater cost to individual reproduction, the rate of decimation, at times of low resources, will be proportionately greater than the rate of growth at times of high resources. Reproducing more slowly will decrease the cost of reproduction, and so the mathematical analysis explains why this occurs in the full model.

A further model (Bryden, 2005c) considers the effects of space on the individual reproduction mechanistic subsystem: agents in the model live and reproduce on a spatial grid. This model can be broken up into two mechanistic subsystems, firstly a non-spatial subsystem with individuals reproducing [as treated by (Bryden, 2005a)] and secondly a spatial subsystem without individuals reproducing. Results indicate that the spatial effects increased the frequency of both periods of high resources and periods of low resources. The mathematical model has shown that this would increase the tendency for individuals to conserve resources and reproduce more slowly. This is an example of a system that combines two interdependent mechanistic subsystems (a spatial and a reproductive system) that interact with each other to produce a macroscopic phenomenon.

Away from the field of agent based modelling, we consider models based on neural biological systems. Neural systems have extremely complex dynamics, which are resistant to mathematical analysis. However, the use of linear stability analysis has proved useful in identifying mechanistic subsystems which can be used as building blocks within larger systems. For example, a system of coupled oscillators, based on the FitzHugh-Nagumo model, has been analysed as a mechanistic subsystem (Buckley et al., 2004). This analysis demonstates how, when the oscillators are linked to a simple gas net, the system can produce temporally distinct oscillations. Much other work continues into the identification of simple oscillatory models, such as that done in Central Pattern Generators (CPGs) (Marder and Bucher, 2001). CPGs can work as mechanistic subsystems within models of animal locomotion systems.

In this section we have demonstrated how a complex system that does not yield to explicit mathematical treatment may be simplified into mechanistic subsystems which are more likely to yield to explicit mathematical treatment. We can observe from the examples chosen that the working models arrived at through such a process consist of both synthetic and analytic explanations.

The process of simplification identified above is not the only way of making simpler models. By choosing different assumptions and approaching a conceptual model from a different perspective it is also possible to open up a system to explicit mathematical treatment. With computer simulation it is increasing easy to change the assumptions of a model and get a feel for how the system changes. This sort of approach is invaluable as a tool for the sort of lateral thinking needed when generating an explicit mathematical model.

Discussion

While this paper has argued that an explicit mathematical treatment will provide a superior explanation of a scientific phenomenon to an equivalent computer simulation, it must be made clear that the overrarching goal of the scientific modeller is to build better models which explain important phenomena which are not as yet understood. To this extent computer simulation is still a crucial part of the modellers toolbox. The ease with which models can be produced with computers is extremely valuable. Furthermore, not only can these early efforts lead to some important scientific results, but they can also point towards new directions for mathematical models. We list below, in increasing order of merit, different styles of working models and explain how valuable each one is in generating scientific explanation. By starting with models at the beginning of the list and progressing up the list, models can become better explanations of scientific phenomenon.

- A description of an opaque computer simulation and some vague rhetorical statements that it consists of an explanation of what it is trying to model. We have argued that this approach is merely setting down a procedure for producing a performance of explanation of some phenomenon. However, this approach can still yield a proof of concept for some topic under debate, or generate insights for empirical experiments.
- The same computer simulation as before, but this time complete with well documented source code, parameters and other data that can easily be tested by other users and reused in new simulations. While this approach does not yet produce a competent explanation, it allows for more simple reproduction of the model which will help others develop it further.
- An opaque computer simulation (with well documented source code) with some manipulations and simulation runs that demonstrate how various attributes of the model explain various phenomena. We have argued that this approach can yield a competent explanation of sorts, but this is merely a synthetic explanation and is not logically grounded.
- An opaque computer simulation (with well documented source code) that has been decomposed into mechanistic subsystems. Some subsystems have been treated mathematically. Such a working model can also yield a competent explanation of sorts, this explanation is more analytic than in the previous case.
- An explicit mathematical treatment. Such a working model yields a competent fully analytic scientific explanation.

As set out above, clearly the best option is to produce an explicit mathematical treatment. However this is rarely simple, and in many cases mathematics is not yet mature enough to approach this goal. Since we must live in the real world, science must answer questions about systems that cannot be yet modelled by mathematical approaches. Computer simulation modelling provides us a working methodology for approaching these complex or complicated systems and making important steps toward understanding them.

Further to this, it is important to note that computer simulation models can extend already established mathematical treatments. By extending or relaxing the assumptions made in the purely mathematical treatment, the new model will rely on the mathematical treatment as a mechanistic subsystem but may produce new results or important insights on the mathematical model (Harris and Bullock, 2002). Since computer simulation models lend themselves to more accurate, relaxed assumptions, when explanations become available they are more likely to be of value to the conceptual model under question.

What is important is that scientific models progress up the order of merit listed. A novel modelling approach that identifies a new style of working model may have value even if it merely provides a performance of some scientific explanation. Such a system can be experimented with and decomposed into mechanistic subsystems and the standard of explanation will improve. This is one of the benefits of computer simulation modelling in that it gives us tools to break down a problem so that we can get closer to an explicit mathematical treatment through an iterative process. Computer simulation models can be thought of as providing tools for developing imagination and lateral thinking in modelling approaches.

Acknowledgements Many thanks to Chris Buckley, Eric Silverman, Netta Cohen and the rest of the Biosystems group at Leeds University.

References

Bedau, M. A. (1997). Weak emergence. In Tomberlin, J., editor, *Philosophical Perspectives: Mind, Causation, and World*, volume 11, pages 375–399. Blackwell Publishers.

Bedau, M. A. (1999). Can unrealistic computer models illuminate theoretical biology. In Wu, A. S., editor, *Proceedings of the 1999 Genetic and Evolutionary Computation Conference Workshop Program*, pages 20–23. Morgan Kaufmann, San Francisco.

Boerlijst, M. C. and Hogeweg, P. (1991). Spiral wave structure in pre-biotic evolution. *Physica D*, 48:17–28.

Bryden, J. A. (2005a). Population growth and decimation rates for reproducing individuals. *Submitted*.

Bryden, J. A. (2005b). Slime mould and the transition to multicellularity: the role of the macrocyst stage. In Capcarrere, M. S., Freitas, A. A., Bentley, P. J., Johnson, C. G., and Timmis, J., editors, *Advances in Artificial Life: 8th European Conference, ECAL 2005, Canterbury, UK, September 5-9, 2005. Proceedings.*, pages 551–561. Springer.

Bryden, J. A. (2005c). Space: What is it good for? In Kim, J. T., editor, *Systems Biology Workshop at ECAL 2005*. Online: `http://www.ecal2005.org/workshopsCD/systemsbiol/index.html`.

Bryden, J. A. and Cohen, N. (2004). A simulation model of the locomotion controllers for the nematode *Caenorhabditis elegans*. In Schaal, S., Ijspeert, A. J., Billard, A., Vijayakumar, S., Hallam, J., and Meyer, J.-A., editors, *Proceedings of the eighth international conference on the simulation of adaptive behavior*, pages 183–192. MIT Press / Bradford Books.

Buckley, C., Bullock, S., and Cohen, N. (2004). Toward a dynamical systems analysis of neuromodulation. In Schaal, S., Ijspeert, A. J., Billard, A., Vijayakumar, S., Hallam, J., and Meyer, J.-A., editors, *Proceedings of the eighth international conference on the simulation of adaptive behavior*, pages 334–343. MIT Press / Bradford Books.

Chomsky, N. (1986). *Knowledge of Language: Its Nature, Origin and Use*. Praeger Publishers: New York.

Di Paolo, E. A. (2000). Ecological symmetry breaking can favour the evolution of altruism in an action-response game. *Journal of Theoretical Biology*, 203:135–152.

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

Frege, G. (1980). *The Foundations of Arithmetic*. Basil Blackwell. English translation by Austin, J.L.

Harel, D. (2002). Towards full reactive modelling of a multicellular animal. In *Proceedings of the EPSRC/NESC UK Workshop on Grand Challenges in Computing Research*. UK Computing Research Committee.

Harris, D. and Bullock, S. (2002). Enhancing game theory with coevolutionary simulation models of honest signalling. In Fogel, D., editor, *Proceedings of the Congress on Evolutionary Computation*, pages 1594–1599.

Kitano, H., Hamahashi, S., Kitawaza, J., Takao, K., and Imai, S.-i. (1997). The virtual biology laboratories: A new approach of computational biology. In Husbands, P. and Harvey, I., editors, *Proceedings of the Fourth European Conference on Artificial Life (ECAL'97)*, pages 274–283. MIT Press / Bradford Books, Cambridge, MA.

Marder, E. and Bucher, D. (2001). Central pattern generators and the control of rhythmic movements. *Current Biology*, 11:986–996.

McGregor, S. and Fernando, C. (2005). Levels of description: A novel approach to dynamical hierarchies. *Artificial Life*, 11:459–472.

Polani, D., Dauscher, P., and Uthmann, T. (2005). On a quantitative measure for modularity based on information theory. In Capcarrere, M. S., Freitas, A. A., Bentley, P. J., Johnson, C. G., and Timmis, J., editors, *Advances in Artificial Life: 8th European Conference, ECAL 2005, Canterbury, UK, September 5-9, 2005. Proceedings.*, pages 393–402. Springer.

Quine, W. V. O. (1953). *From a Logical Point of View*. Harvard University Press.

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

Watson, R. A.and Pollack, J. B. (2005). Modular interdependency in complex dynamical systems. *Artificial Life*, 11:445–458.