Complexity beyond agent based models

D. Gross[1] and David Jefferies[2]

[1]*Senter for Vitskapsteori, University of Bergen*
Bergen, Norway
Email: `Dominique.Gross@svt.uib.no`

[2]*Department of Electrical Engineering, University of Surrey*
Guildford GU2 7XH, UK
Email: `d.jefferies@surrey.ac.uk`

Abstract. Even if there does not yet exist a generally accepted definition of complexity, there are a number of definitions which serve well for specific purposes. Most of these can be implemented in so-called agent based models. However, authors such as Funtowicz and Ravetz (Funtowicz & Ravetz 1994; 1997) or Rosen (Rosen 1991; 1999) argue that a full understanding of complexity has to go beyond computer models. This article tries to elucidate this by means of two examples.

Keywords: measures of complexity, agent based modelling, genuine novelty, reductionism

1. Introduction

Today the term "complexity theory" or the "science of complexity" is often associated with agent based models. Indeed these models, as for example described by Casti (Casti 1997) or Holland (Holland 1995; 1998), have proven to contribute substantially to the understanding of a variety of natural systems that are generally considered to be complex; examples include evolution, economic systems, the immune system and many more.

In spite of this success, one might argue that a too strong association of the "science of complexity" with agent based (computer) models might lead to an incomplete picture of the nature of real complex systems. Funtowicz and Ravetz (Funtowicz & Ravetz 1994; 1997), for example, consider "real complexity" or as they call it "emergent complexity", exclusively realised in social systems, whereas agent based models are only "ordinary complex".

> Emergent complex systems, by contrast, cannot be fully explained mechanistically and functionally; in them, at least some of the elements of the system possess individuality, along with some degree of intentionality, consciousness, foresight, purpose, symbolic representation and morality. Attempts to reduce human society completely to ordinary complexity can result either in unrealistic theories (as those of B.F. Skinner) or catastrophic policies (as those of Pol Pot). (Funtowicz & Ravetz 1994, p. 570)

Following Funtowicz and Ravetz a complete theory of complexity has to go beyond "ordinary complexity", i.e. agent based models. For another account of complexity in social systems see Dimitrov (Dimitrov 1997). Rosen (Rosen 1991) comes to a similar conclusion. He considers systems which include...

> ...closed loops of efficient causation. Systems of this type cannot be simulated by finite-state machines (e.g. Turing machines); hence they themselves are not machines or mechanisms. [...] I call these systems complex. (Rosen 1999, p. 24)

Thus complex systems in Rosen's sense cannot be modelled by means of agent based models, without leaving out "essential" parts. What this actually could mean will be illustrated in section 4.1.

Inspired by Rosen and Funtowicz and Ravetz, I will, by means of two examples, illustrate that there are relevant effects in nature that cannot be implemented in computer models. Sometimes it will be both intuitively correct and relevant to include these into a notion of complexity. This notion will then have to surpass what can be implemented in agent based computer models.

The paper is organised as follows. In section 2 I will try to show that complexity is currently used pragmatically. In section 3 inherent limitations of computer models are reviewed and in section 4 these limitations are contrasted with two examples of real complex systems. Section 5 discusses briefly the relevance of the extended notion of complexity.

2. Different Ideas about Complexity

There is no general consensus about how to formally define complexity (for a review of different notions and measures see for example Edmonds (Edmonds 1999) or Rescher (Rescher 1998)). It seems that the choice of an appropriate complexity measure in practical work (either modelling or others) is to a large extent solved pragmatically, i.e. the specific circumstances of the problem or the task constrain to a large extent the appropriate measures of complexity.

Perhaps one of the oldest quantitative definitions of complexity, which is still in use today, measures the complexity of a system by the length of its shortest description or the shortest algorithm that generates it (see Kolmogorov (Kolmogorov 1965) and for a more refined version see Bennett (Bennett 1986); for a good introduction to such informational complexity measures see also Adami (Adami 1998) or Bar-Yam (Bar-Yam 1997)). Informational measures of complexity are also often connected to entropy. An advantage of such measures is clearly that they allow us to pin down the complexity of a system to a number. However, although in principle applicable to formal systems (see for example Fletcher *et al.* (Fletcher *et al.* 1997) or Langton (Langton 1990)) and to real systems (see for example Bar-Yam (Bar-Yam 1997)), it is often unclear what the shortest description of a real system is. Thus algorithmic measures are often rather uninformative.

A rather different approach to complexity is chosen by McShea (McShea 1996). His initial aim is to test for a trend of complexity during a certain period of biological evolution. In the absence of a general definition, he uses a list of a few structural criteria in order to "decide" the complexity of organisms. As McShea himself notes, the approach is only valid within some limitations, which in the first place stems from the *ad hoc* choice of the criteria. He stresses that his particular choice leads to counterintuitive results in some cases, which does not prevent its usefulness within its scope, as long as one is conscious about it. From a more general point of view one can consider this as an example of a very pragmatic use of complexity. In the absence of a general definition one might create a "local definition" that serves well within a limited scope. It has to be noted that agent based models can be complex in the sense of McShea's notion.

Very often complexity is associated or even identified with so-called CAS (complex adaptive systems) (see for example Holland (Holland 1998; 1995) or Casti (Casti 1997)). CAS are real or artificial systems (models) with a medium sized number of autonomous "agents", which act and interact with each other and their environment according to simple local rules. Typically these agents can, to a certain extent, adapt to changes in their surroundings. Examples of real systems that can be viewed as CAS include many economic and evolutionary, ecological and social systems. The CAS-stereotypes are extremely easy to implement in agent based models, which makes them the method of choice when modelling this kind of real system.

Often the complexity of a system or a model is connected to a specific state between order and chaos, or at the edge of chaos (EOC). The importance of the EOC for the understanding of complex systems has been proposed by Bak (Bak 1997), Kauffman (Kauffman 1993; 1995) and Langton (Langton 1990). The basic idea is that a system can only be "creative" if there is enough disorder present to guarantee the creation of variation on the one hand, but enough order to conserve variations that may come up. If the system is too ordered it usually freezes into a stable state, whereas deep in the disordered region, there is a chaotic regime, which behaves quasi-stochastical and does not allow for the preservation of creative novelty. It is no problem to test for EOC in artificial systems, whereas it is much more difficult, though feasible, to apply the concept to real systems.

The various notions of complexity that have been reviewed in this section are in no way mutually exclusive. In general one might select a notion of complexity with respect to the appropriateness of the specific problem that has to be attacked. In one context a notion of complexity that is closely connected to the EOC might be illuminating, in another context one might prefer a structural notion of complexity *à la* McShea. At present there is no argument, I am aware of, that would make one specific notion of complexity in general preferable to any other. A choice between them will always depend on pragmatic considerations.

As I have stressed, all the definitions of complexity reviewed in this section have one thing in common: assuming any of them, both real system and artificial systems can be (at least to a certain degree) complex. In the next section I will illustrate that this is not the case with all possible notions of complexity.

3. Limitations to Complexity in Agent Based Models

3.1. Models

In order to prevent any confusion due to unclear meanings of various concepts I will first of all clarify my notation. The distinction between *micro-level* and *macro-level* is the most fundamental one in connection with agent based models. The former consists of the interaction of the agents with one another and with their environment and the parameter settings of the model. The *micro-state* at time t is a complete description of the agents' and the environment's states in terms of *micro-variables*. The *micro-dynamics* specifies how a certain micro-state m_{t_0} at t_0 is to be transformed into another micro-state m_{t_0+1} at $t_0 + 1$. The micro-state and the micro-dynamics are part of the micro-level. In agent based models the micro-dynamics is often formulated by simple rules that specify how the agents react to a given environmental situation; often this has the form of a look-up table that specifies a response to a given input. Agent based models are best implemented by means of modern object-oriented programming languages such as objective C/C++ or Java; today there are even pre-programmed packages such as *Swarm* (Swarm Development Group 1999) available, which facilitates the process of modelling significantly.

The *macro-variables* describe the aggregate behaviour of the model as a whole or a well defined part of it. Usually they are obtained from the micro-variables by means of arithmetic operations, averaging or statistical analysis. Thus at any time t macro-variables can be defined. The behaviour of the macro-variables over time define the *macro-behaviour* of the model, which can be submitted to various numerical analyses. In general, the macro-states and the macro-behaviour can be referred to as the *macro-level* of the model. In this view the macro-level is redundant relative to the micro-level, i.e. it does not contain any information that is not already present on the micro-level. The transformation of micro-variables to macro-variables involves purely syntactic transformations.

In computer models in general and specifically in agent based models one has to strictly distinguish between the syntactic structure of the model and its semantic contents. In principle agent based models are, to put it in John Holland's words, nothing but "streams of numbers in a computer" (Holland 1995) devoid of any intrinsic meaning. In order to serve as models of real systems, these naked data structures have to be interpreted. This fills the syntactic structures of the computer with semantics and creates the connection to a real system. The semantics of the model can be thought to be independent of the syntactic structure of the model and the system (at least in first order approximation). This is often described by the so-called modelling-relation (see Casti (Casti 1992) or Rosen (Rosen 1991)).

3.2. Agent Based Models Do Not Create Genuine Novelty

Accepting this it becomes immediately clear that in agent based models the macro-level cannot be semantically independent from the micro-level. This has two consequences. Firstly, new meaning cannot emerge on the macro-level; the semantic content of the macro-level can be completely deduced from the semantic content of the micro-level. One might imagine an agent based model of an ideal gas by explicitly representing each molecule. Then it certainly is possible to calculate at each time-step an average kinetic energy of the molecules, thus deducing a kind of temperature of the virtual gas, but there will not emerge "hotness" (compare Baas and Emmeche (Baas & Emmeche 1997) and Baas (Baas 1992)). It must, however, be allowed for the possibility that the macro-behaviour of the model induces a number of fruitful associations and insights that go far beyond what actually can be deduced from the model. Holland sees a certain danger in this and warns of unjustified interpretations.

> There are cases in the literature, where a relatively simple algorithm, such as linear regression, is identified with some sophisticated real world process, such as "perception." [...] Facile labelling of what are, after all, streams of numbers in a computer, leaves too much to the eye of the beholder. (Holland 1995, p. 317)

The emergence of novelty in agent based models is not only limited at the transition between hierarchical levels, but also in the temporal development of the model. In principle the behaviour of the models is fixed for all time by their micro-dynamics. Deviations from a deterministic picture can be caused by the introduction of pseudo-random numbers. This is not quite true, though, since these numbers are themselves generated by a deterministic algorithm; thus they do not change the fact that the model is fully determined by its initial state and its micro-dynamics. Even if genuine random numbers were included into the model, it would not change the principle situation very much, since it would only introduce a variation, which is *not* the same as genuine novelty. Still the micro-dynamics of the model sets narrow limits to what kind of behaviour (both on the micro- and on the macro-level) can be expected to happen (compare the notion of the "microstructural frame" (Gross & Strand 2000)).

3.2.1. Tierra

A number of well-known agent based models seem to be counterexamples in that they seem to produce genuinly novel behaviour on the macro-level. One might think of Thomas Ray's artificial life system *Tierra* (Ray 1996), which consists of self-reproducing computer programmes, that compete for limited CPU-time and memory and are written in a machine code-like instruction set of a virtual processor. The *Tierra* organisms' only task is to self-reproduce; this process is, with a low probability, subjected to mistakes, i.e. an instruction might be exchanged with another one in the daughter programme; the "offspring" then slightly differs from its mother, and might enable it to reproduce more effectively, but might also result in an inability to reproduce. Ray interprets this as a mutation. Planting an arbitrary self-reproducing ancestor into the *Tierra*-world, the mutations together with mortality and finite resources will lead to the development of ever more sophisticated organisms throughout an evolutionary process. In the course of time the strategies these digital organisms employ in order to reproduce as effectively as possible, take fascinating forms. Amongst other things one can see the emergence of parasitism, hyper-parasitism, immunity to parasites and so on. These features certainly have not been explicitly programmed into the micro-level of *Tierra*; in this sense one might be inclined to see parasitism as "genuinely novel" on the macro-level.

Looked upon from a different angle, this point of view might be relativised. Firstly *Tierra* is after all nothing but the execution of an algorithm or "streams of numbers in a computer". At the end of the day it is equivalent to a Turing machine, which transforms input into output according to a programme. Secondly it has to be noted that there is a lot of structure implicit to the *Tierra* system. Ray himself stresses the specific conditions the machine code instruction set has to fulfil in order to be a workable basis for an evolutionary system like *Tierra*. Most machine-languages are by far too "brittle" to allow progress by mutations. Furthermore each simulation run has to be initiated by an ancestor-programme that already is capable of self-reproduction. A great deal of information about how (digital) life processes should work is thus already implicit within the micro-level of the model. Thus, what *Tierra* does, is nothing but to introduce pseudo-random variations into a pre-specified self-reproducing algorithm and then to compare the effectiveness of the varied algorithms according to given criteria.

Thus, what seems to be genuine novelty is nothing but a very narrowly confined exploration of variants. *Tierra* organisms will never develop consciousness or start to change their environment or start to populate new areas of RAM (comparable to the transition of life (as for example described by McMenamin (McMenamin & McMenamin

1994))); they will always stay within the limits that are given by the programmer. In this sense there is no room for genuine novelty in computer-based systems.

In this context it is also worth mentioning that Bedau (Bedau 1997) acknowledges the fact that present artificial life systems, including *Tierra*, are relatively limited in their adaptive potential, as opposed to the real biosphere. He nevertheless seems to be confident that this can be overcome in principle.

3.3. *Agent Based Modelling is Reductionist*

Very closely connected to the problem of genuine novelty, is the age-old reductionist / anti-reductionist debate (see for example Dupré (Dupré 1993) or for social sciences Bhargava (Bhargava 1992)). There exists a lot of confusion about this topic, partly because authors are often unclear about what they mean by reductionism. Thus before proceeding, I will explain my use of the notion.

In this article reductionism is the view that every phenomenon that we observe in natural or artificial systems is strictly determined and generated by some (in a sense) more fundamental entities and their interactions; these interactions must be computable, but are allowed to contain stochastic elements. Often a distinction between practical and ontological reductionism is made. The former is the view that the reduction of all phenomena to some elementary particles (usually those of physics) can and will actually be done one day, whereas ontological reductionism is weaker and only maintains the in principle possibility. Most people would agree that it is unlikely that practical reductionism is true. Bedau's "weak emergence" (Bedau 1997) is clearly compatible with reductionism in this sense. Let me stress that the negation of reductionism does not necessarily imply vitalism (for possibilities of non-reductionist modelling see for example Rosen (Rosen 1999) or Holt (Holt 2000)).

In agent based models the macro-behaviour is exclusively generated by lower level entities (the agents) and their interactions (with each other and their environment). It can thus be said that agent based modelling is reductionist. It becomes now obvious that the aforementioned failure of agent based models to generate genuine novelty is just a consequence of their reductionist structure. However, Epstein (Epstein 1999) for example sees the reductionism of agent based models as a possibility to demystify a gap between macro and micro "by identifying micro-specifications that are sufficient to generate—robustly and replicable—the macro(whole)" (Epstein 1999, p. 55).

Putting this together we are facing three possibilities. Firstly one might hold the view that every complex system can be appropriately modelled by agent based models and, provided the availability of sufficient computing resources, they actually can be realised. Secondly, one might think that agent based models are in principle appropriate, but because of practical limitations, there will always be some complex systems which cannot be modelled sufficiently. Finally one might insist that certain phenomena cannot, not even in principle, be completely represented by reductionist models, but crucially depend on constraints on the macro-level.

The first view is, I think, implausible, since it would ultimately imply practical reductionism. Thus there remain the second and the third possibilities, which at least imply that there are some systems that cannot actually be completely reduced to a lower level, i.e. no agent based (or any other reductionist model) can capture them. I think this is not a very controversial claim and is in fact supported by a number of authors (see for example Dupré (Dupré 1993), Rosen (Rosen 1999) or Rose (Rose 1997)). The third view would additionally imply that there is some novelty creation at the transition to higher levels in natural systems. By means of the following two examples of biological systems, I will try to show that the failure of agent based models to capture some systems might in some situations have some bearing on the notion of complexity.

4. Non-Computable Complexity

4.1. Utricularia Floridana

The first example is due to Ulanowicz (Ulanowicz 1997) and is centered upon the water-plant *Utricularia floridana*. This plant secretes different polysaccharrides on the surface of its leaves, which attract various bacteria and algae that feed upon the sugars. Various microscopic animals (zooplankton) use these algae and bacteria as food. From time to time the microscopic animals touch tiny hairs of the plants, which then open a hole on the surface into which the zooplankton is sucked. The *Utricularia* is able to extract nutrients from the decomposing animals. Ulanowicz sees this as an example of an autocatalytic process in nature. He writes:

> It is important [...] to note that in any biological system the components maintain some plasticity or indeterminacy. Such is obviously the case with the periphyton and zooplankton communities, for their compositions change with various habitats. Plasticity applies as well over the longer time scale to *Utricularia* itself, which has evolved into numerous species, and even exhibits a degree of polymorphism over rather short intervals[...]. Such plasticity or adaptability contrasts with the usual situation in chemistry, where the reactants in any autocatalytic process are fixed, thereby contributing to the stereotypical image of autocatalysis as a "mechanism." (Ulanowicz 1997, pp. 44-46)

Ulanowicz stresses that this symbiotic relation depends heavily on the specific circumstances of the environment, i.e. the symbionts are always constituted of different species.

As modellers we can now be interested in two aspects, namely how the system came into existence and how it develops further. Let us first concentrate on the first problem. Which ingredients do we need in an agent based model in order to generate this system? We certainly need at least three types of agents, A, B and C, where A feeds on B, B on C and C on A. Furthermore, in order for the system to function, we need to open it to the environment, that is we have to allow some kind of nutrient inflow to the system. This would certainly not be enough, since the model does not yet possess any biological significance and is thus indistinguishable from any chemical autocatalytic "mechanism" of the kind Ulanowicz mentions. Furthermore such a system would not answer any ontogenetic question. What we actually want is that the components, to use an anthropocentric expression, "discover" that they are better off together, and so "initiate" a symbiosis; in other words, we want to initiate an adaptive process that leads from a rather unspecified state of the three "ancestor-agents" with "ancestor-rules" to the autocatalytic configuration. In agent based models this is, in principle, not very hard to achieve (see for example Kauffman (Kauffman 1993) or Bagley and Farmer (Bagley & Farmer 1991)). However in the specific case, a minimum requirement to the model is the specification of meta-rules that describe how the adaptive process works. These meta-rules and the ancestors again have to be of a very particular kind.

Let me just mention two conditions. Firstly the rules and meta-rules have to be such that an evolutionary process really can yield fit variants. The underlying coding of the genotype must not be too "brittle". The instruction set of the *Tierra* model, for example, is of this kind but, as mentioned above, a typical CPU-machine language such as those that we find in our PCs is not. Secondly, the ancestor-agents must be able to devop the ability to interact with each other. In *Tierra*, for example, interaction between agents is very limited and is not rich enough to produce a direct interaction of the kind we want in order to implement the relations of the *Utricularia* example. These two conditions are already very specific.

There are now three possibilities of how we can have ancestor agents and rules and meta-rules of this kind. They could be explicitely specified by an external programmer; this would not solve the ontogenetic problem only reduce it to the problem of the ontogeny of the meta-rules and the ancestors. The second possibility is to implement meta-meta rules and ancestors to ancestors. This would lead into an infinite regress. Finally one could close the chain somewhere, thus producing a self-consistent loop; unfortunately, according to Rosen, this is not possible in finite state machines (see section 1). I conclude that the ontogeny of this system cannot be modelled in a computer model without leaving out relevant information. The reason for this is clearly the lack of genuine novelty and the reductionist approach. In models the agents will forever stay within their limits given by the programmer; no matter how much time passes, they will always remain within their boundaries. Real world agents cannot, at least not in a scientific world-view, be considered as specified by somebody in this way. In this sense they are examples of the creation of genuine novelty.

Let us now turn to the second modelling problem, to predict or at least understand the behaviour of the system, once it exists. Agent based models seem to be a good option to choose in order to come to a better understanding of the system. However, even for this purpose Ulanowicz would consider reductionist approaches poorly suited. Consider this quote as evidence:

> Although the system requires material and mechanical elements, it is evident that some behaviours, especially those on a longer time scale, are, to a degree, *autonomous* of lower level events[...]. Attempts to predict the course of an autocatalytic configuration by ontological reduction to material constituents and mechanical operation are accordingly, doomed over the long run to fail. (Ulanowicz 1997, p. 49)

Following this, one sees immediately that the symbiotic relations of the *Utricularia* cannot be represented by agent based models without losing relevant biological content, because the macro-level is to a certain degree independent of ("supervenes") the micro-level. There is genuine novelty at the higher level, which cannot be implemented in agent based models. Thus, at least if one believes Ulanowicz, one has to accept that there are real systems whose behaviour surpasses what can be observed and reproduced in computer models. Intuitively it also seems that these properties are in fact relevant for the notion of complexity. The next example, I think, makes this statement much stronger and shows very clearly that it is useful to consider notions of complexity that go beyond what we can see in computer models.

4.2. Cichlid Fishes in Lake Victoria

The description of the next example will largely follow Goldschmidt (Goldschmidt 1997). The ecology of Lake Victoria as it developed throughout the last several hundred thousand years was dominated by a large number of cichlid fish species (furu). These fishes were abundant but mostly small to medium sized; there were no large predators in the lake. None of the indigenous fishes were ideal for commercial use, which was one of the reasons why it was suggested to introduce the Nile perch—a large predator fish, perfectly suitable for commercial fishing—into the waters of Lake Victoria.

As one might easily imagine, the cichlid fishes—not used to predation—were an easy prey for the Nile perch who drove them to the edge of extinction. This led to a number of second-order effects; amongst others the decline of detritus eating fishes (marginalised by the Nile perch) led to a spread of algae in the lake; furthermore the now (almost) extinct fishes partly fed upon mosquito larvae and helped to control the number of adults; in the absence of this control mosquito abundance greatly increased. Apart from this second-order effect one might expect that the Nile perch itself, after a rather short time, would have driven to extinction all indigenous fish and thus its own basis of food.

> Nile perch in open waters continued eating fish until there were virtually none left. But where were the thousands of starving Nile perch floating moribundly on the water's surface? Why hadn't their number declined rapidly? Why had the predicted collapse of the Nile perch population not taken place? (Goldschmidt 1997, p. 226-227)

The solution to this is connected to the prawn that up to then played a minor role in the ecosystem, but managed to take advantage of the changes in the lake by changing its diet to organic waste, which was an easy source of food since there was no longer competition for it by other species. For this and other reasons the prawn population managed to increase significantly in size. The Nile perch, in turn, changed its diet to prawns, thus managing to avoid its extinction and the expected economic disaster. A similar story can be told with the sardine that also was one of the winners of the introduction of the Nile perch.

I think in this example one sees another instance of what Ulanowicz calls "plasticity". The question is now how agent based models could have been used to predict (qualitatively but to a pre-specified degree of accuracy) the consequences of the introduction of the Nile perch. Clearly, such a model needs several types of agents representing the different species in the lake plus their interaction rules. Retrospectively one knows that additionally at least the organic waste on the bottom of the lake was of major importance, but this could not have been anticipated beforehand; thus presuming that some details play a major role, but not knowing which, one has to include as many as possible into the model in order to be able find out which of them are of importance; this actually blows up the model significantly, because each additional feature demands at least one separate interaction rule for the agents. Additionally one needs to include meta-rules that specify how the fishes' behavior changes in reaction to environmental changes. It is easy to see that the resulting model would have been huge. I presume that nobody could actually succeed in constructing such a model, which would have predicted, say, which species would survive the introduction of the Nile perch. One would have had to produce a giant model. This is not only against good modelling habit, but also often considered to be counterproductive (as discussed by Brooks and Tobias (Brooks & Tobias 1996)).

This situation can now be interpreted in two different ways. Either it is in principle impossible to model the system reductionistically (a view Ulanowicz would have probably held), i.e. neither of the reductionist positions is true, or it is at least too hard to implement all necessary specifications into the system, i.e. ontological reductionism is true. If the latter is the case, then it is easy to see that more computing resources could not solve the problem because at least parts of the problem is the specification and not the computation of the model. This specification is necessary because the model fails to generate genuine novelty by itself.

5. Discussion and Conclusion

Intuitively, I think, most would agree, that the two preceding examples were about complex systems. Furthermore it is rather clear that these systems displayed general novelty, in the case of the *Utricularia* at the transition between the levels and in the case of the cichlid fishes in the temporal dimension. Let me now set the last example in a broader context to demonstrate that novelty can be a relevant factor of complex systems.

A complete understanding of the consequences of the introduction of the Nile perch into Lake Victoria has to go beyond a mere investigation of its impact on the ecosystem. Crucial human interests are closely connected to the state of the lake. Firstly, the communities around the lake depend at least partially on the fishing industry; a drastic decrease of the biodiversity might lead to an impoverishment of the region. Moreover there are health issues to be taken into account. A drastic increase of mosquitos could have resulted in a malaria epidemic. Thus, what was at first a mere uncertainty of the outcome of an intervention into a system, now becomes risk in the broader perspective.

An intervention into a complex system might, as the cichlid example shows, produce unanticipated consequences that even transgress system boundaries and produce genuinely novel behaviour. Blindly trusting agent based models, which are limited in their ability to display genuinly novel behaviour, could be fatal, when this novel behaviour is potentially dangerous. I think this is what the quote by Funtowicz and Ravetz in section 1 attempts to express.

Thus it might make sense to include genuinely novel behaviour in the notion of complexity especially under certain circumstances, for example when modelling interventions into complex natural systems. Then reductionist methods such as agent based models can only be expected to give a limited representation of the system and are strongly constrained by the programmer's choices. Natural systems are not, thus limiting oneself to mere variational models might lead to an underestimation of risk and inappropriate policy recommendations.

References

C. Adami. (1998), *Introduction to Artificial Life.* Springer-Verlag, New York.

W. Arthur, B. LeBaron, B. Palmer, and R. Taylor. (1997), "Asset Pricing under Endogenous Expectations in an Artificial Stock Market." In *The Economy as an Evolving Complex System II*, *Studies in the Sciences of Complexity*, Santa Fe Institute.

N. Baas. (1992), "Emergence, Hierarchies, and Hyperstructures." In *Artificial Life III* (1992), vol. XVII of *Studies in the Sciences of Complexity*, Santa Fe Institute, 515–537.

N. Baas and C. Emmeche. (1997), "On Emergence and Explanation." *Intellectia*, **2(25)**, 67–83.

R. Bagley, D. Farmer. (1991), "Spontaneous Emergence of a Metabolism." In *Artificial Life II*, 93–140.

P. Bak. (1997), *How Nature Works.* Oxford University Press.

Y. Bar-Yam. (1997), *Dynamics of Complex Systems.* Addison-Wesley, Reading.

M. Bedau. (1997), "Weak Emergence." *Philosophical Perspectives 11*.

M. Bedau. (1997), "A Comparison of Evolutionary Activity in Artificial Evolving Systems and the Biosphere." In *Proceedings of the Fourth European Conference on Artificial Life, ECAL97* , 125-134.

C. Bennett. (1986), "On the Nature and Origin of Complexity in Discrete, Homogeneous, Locally-Interacting Systems." *Foundations of Physics*, **16(6)**, 585–592.

R. Bhargava. (1997), *Individualism in Social Sciences.* Clarendon Press, Oxford.

R. Brooks, A. Tobias. (1996), "Choosing the Best Model: Level of Detail, Complexity, and Model Performance." *Mathematical and Computer Modelling 24*, 4 , 1–14.

J. Casti. (1992), *Reality Rules.* John Wiley & Sons, New York.

J. Casti. (1997), *Would-Be Worlds.* John Wiley & Sons, New York.

J. Casti. (1998), "The Science and Surprise of 'Would-be' Artificial Worlds: The Collaborative Process of Developing an Industry-wide Model." Business Processes Resource Centre.

V. Dimitrov. (1997), "Use of Fuzzy Logic when Dealing with Social Complexity." *Complexity International*, vol. 4. `http://life.csu.edu.au/complex/ci/vol4`.

J. Dupré. (1993), *The Disorder of Things.* Harvard University Press, London, 1993.

B. Edmonds. (1999), *Syntactic Measures of Complexity.* PhD thesis, University of Manchester. Available at `http://www.cpm.mmu.ac.uk/ ~bruce/thesis`.

J. Epstein and R. Axtell. (1998), *Growing Artificial Societies.* The MIT Press, Cambridge.

J. Epstein. (1999), "Agent-Based Computational Models and Generative Social Science." *Complexity*, **4(5)**, 41–60.

J. Fletcher, M. Bedau, and M. Zwick. (1997), "Dependence of Adaptability on Environmental Structure in a Simple Evolutionary Model." *Adaptive Behavior*, **4(3)**, 283–315.

S. Funtowicz and J. Ravetz. (1990), *Uncertainty and Quality in Science for Policy.* Kluwer, Dodrecht.

S. Funtowicz and J. Ravetz. (1993), "Science for the Post-Normal Age." *Futures*, **25(7)**, 739–755.

S. Funtowicz and J. Ravetz. (1994), "Uncertainty, Complexity and Post-Normal Science." *Environmental Toxicology and Chemistry*, **13(12)**, 1881–1885.

S. Funtowicz and J. Ravetz. (1994), "Emergent Complex Systems." *Futures*, **26(6)**, 568–582.

S. Funtowicz and J. Ravetz. (1997), "The Poetry of Thermodynamics." *Futures*, **29(9)**, 791–810.

T. Goldschmidt. (1997), *Darwin's Dreampond — Drama in Lake Victoria.* The MIT Press, Cambridge.

D. Gross and R. Strand. (2000), "Can Agent-Based Models Assist Decisions on Large-Scale Practical Problems? A Philosophical Analysis." *Complexity* **5(6)**.

J. Holland. (1994), "Echoing Emergence: Objectives, Rough Definitions, and Speculations for ECHO-Class Models." In *Complexity—Metaphors, Models, and Reality*, volume XIX of *Studies in the Sciences of Complexity*, Santa Fe, 310–334. Institute.

J. Holland. (1995), *Hidden Order.* Addison-Weseley, Reading.

J. Holland. (1998), *Emergence.* Oxford University Press, Oxford.

A. Holt. (2000), "Understanding Environmental and Geographical Complexities through Similarity Matching." *Complexity International*, vol. 7. http://life.csu.edu.au/complex/ci/vol7.

S. Kauffman. (1993), *The Origins of Order.* Oxford University Press, Oxford.

S. Kauffman. (1995), *At Home in the Universe.* Oxford University Press, Oxford.

S. Kauffman and W. Macready. (1995), "Technological Evolution and Adaptive Organizations." Santa Fe Institute Working Paper 95-02-008.

A. Kolmogorov. (1965), "Three Approaches to the Quantitative Definition of Information." *Problemy Peredachi Informatsii*, **1(1)**, 3–11.

C. Langton. (1990), "Computation at the Edge of Chaos." *Physica D*, **42**,12–37.

D. McMenamin and M. McMenamin. (1994), *Hypersea: Life on Land.* Columbia University Press, New York.

D. McShea. (1996), "Metazoan Complexity and Evolution: Is there a Trend?" *Evolution*, **50**, 477–492.

T. Ray. (1996), "An Approach to the Syntheses of Life". In *The Philosophy of Artificial Life*, 111–145. Oxford University Press.

N. Rescher. (1998), *Complexity – A Philosophical Overview.* Transaction.

S. Rose. (1997), *Lifelines.* Penguin Books, London.

R. Rosen. (1991), *Life Itself.* Columbia University Press, New York.

R. Rosen. (1999), *Essays on Life Itself.* Columbia University Press, New York.

Swarm Development Group. (1999), *The Swarm Simulation Systems.* Web-based Documentation, URL: http://www.swarm.org.

R. Ulanowicz. (1997), *Ecology, the Ascending Perspective.* Columbia University Press, New York.