

Kent Academic Repository

Full text document (pdf)

Citation for published version

Chu, Dominique (2011) Complexity: Against Systems. Theory in Biosciences . pp. 182-196. ISSN 1431-7613.

DOI

<https://doi.org/10.1007/s12064-011-0121-4>

Link to record in KAR

<https://kar.kent.ac.uk/30776/>

Document Version

UNSPECIFIED

Copyright & reuse

Content in the Kent Academic Repository is made available for research purposes. Unless otherwise stated all content is protected by copyright and in the absence of an open licence (eg Creative Commons), permissions for further reuse of content should be sought from the publisher, author or other copyright holder.

Versions of research

The version in the Kent Academic Repository may differ from the final published version.

Users are advised to check <http://kar.kent.ac.uk> for the status of the paper. **Users should always cite the published version of record.**

Enquiries

For any further enquiries regarding the licence status of this document, please contact:

researchsupport@kent.ac.uk

If you believe this document infringes copyright then please contact the KAR admin team with the take-down information provided at <http://kar.kent.ac.uk/contact.html>

Complexity: against systems

Dominique Chu

Received: 29 April 2010 / Accepted: 17 January 2011
© Springer-Verlag 2011

Abstract This article assumes a specific intuitive notion of complexity as a difficulty to generate and/or assess the plausibility of models. Based on this intuitive understanding of complexity, it identifies two main causes of complexity, namely, *radical openness* and *contextuality*. The former is the idea that there are no natural systems. The modeler always needs to draw artificial boundaries around phenomena to generate feasible models. Contextuality is intimately connected to the requirement to simplify models and to leave out most aspects. Complexity occurs when contextuality and radical openness cannot be contained that is when it is not clear where the boundaries of the system are and which abstractions are the correct ones. This concept of complexity is illustrated using a number of example from evolution.

Keywords Complexity · Evolution · Artificial life · Contextuality

Introduction

It can be extraordinarily difficult to develop models to predict or indeed explain phenomena. Some phenomena resist being modeled. This resistance is often called “complexity.” Of course, modeling always requires considerable skill and experience and is therefore hard to do. Theoretical physicists, for example, must be acquainted with quite deep mathematics. In some areas of theoretical physics, modeling techniques can get so difficult that there

may be only a handful of people who master these methods. This is not what we mean by “complex.” Modeling in physics still leads to models that are general in scope and often relatively simple (even aesthetically pleasing) and they also have a clear path to experimental corroboration. This is not so in the above mentioned domains that are commonly associated with complexity. They effectively resist being captured by general and universal models that allows the unification of a large number of phenomena within a single elegant theory. Climate change, economics, and evolution are all complex in the sense that they require large, messy models that are difficult to formulate and uncertain in their status.

To the untrained eye, this kind of complexity should not come as a great surprise. In our daily experiences, we do expect to observe some regularity. The sun sets and rises every day (in most parts of the world anyway), organisms are born and die, objects fall to the ground, bedrooms get messy, and so on. There are many such regularities, but we do not, intuitively expect any deep connection between them, any universal law from which everything can be deduced.

The view from physics, where simple and universal laws, “grand unified theories” or “theories of everything” are expected and searched for, seems somewhat spoiled. Real life is messy. Yet, often there is a demand to model phenomena that are more like real life than the idealized world of theoretical physics. Most fields of enquiry are not captured by simple and general theories. This is what we mean by complexity.

This article is about complexity. Specifically, we will investigate the causes of complexity. In order to do this, we have to start with a basic intuition about what we mean by complexity. Within the context of this article, we will interpret complexity as the difficulty to find suitable (formal) models (either predictive or explanatory) for the

D. Chu (✉)
School of Computing, University of Kent, CT2 7NF Canterbury,
United Kingdom
e-mail: d.f.chu@kent.ac.uk

phenomena we are interested in. We acknowledge that there are other ideas about complexity. For example, in the literature complexity is often associated with certain types of computer simulations [“complex adaptive systems” (Holland 1995; 1998; Bedau and Smith 1999)] that show interesting (“emergent”) behavior arising from simple interactions only (such as, for example, in the famous Game of Life). These models convinced some that complexity can arise out of simplicity. Motivated by models of insects that demonstrate how elaborate structures can be built by very simple automata-like agents, there has been a lot of excitement for some time: One of the implications of this notion of complexity is that complexity is only apparent and can, if one is clever enough, be reduced to fundamentally simple rules that elegantly unify many of the apparently messy phenomena we observe [see, for example, the notion of universality classes (Bak 1997; Stanley et al. 1996a, b,)].

Another idea of complexity is to see it as some sort of intricate order that is resilient to random shocks. For example, gene regulatory networks, while large and highly connected, are also well controlled and resilient to noise and most minor structural changes. This is sometimes called complexity [see, for example, (Noble 2009)]. Complexity is also used to denote the opposite, namely, chaos and disorder, possibly arising from very simple interactions. Sometimes complexity is associated with certain anti-reductionist approaches to science, sometimes complexity is equated with non-linearity, sometimes complexity is used to justify reductionism (Wynne 2005). And then there are countless more or less well defined other meanings of the word complexity.

Altogether, it seems that there is a lack of consensus not only on the precise definition of complexity, but also on the basic intuition behind the concept. Having arrived at this point one now has to make a choice. Either one accepts that complexity is a word that can mean many things without doing harm. Or, one decides that there is something of deeper interest behind one or more of the possible meanings and decide to investigate this in more detail. In this article, we take the latter approach. We think that a notion of complexity-as-a-difficulty-to-model is relevant and deserves a deeper analysis. It crops up in all the big questions of our time, from climate change to credit crunch. Understanding complexity in this sense means understanding what precisely it is that makes it so hard to model certain things, while others are easy to model. This is, in essence the question we will address in this contribution. We will therefore, in what follows, understand complexity exclusively as this difficulty to find suitable models, and exclude other meanings.

In order to understand some of the factors that make modeling hard, we will introduce the concepts of *radical*

openness and *contextuality*. These concepts will be defined in Sects. 5 and 6, respectively, and refer to very specific difficulties of the modeling process. Most practicing modelers will, we suspect, be quite familiar with the basic ideas behind these two concepts although many will not recognize them as particularly problematic. Below we will argue that in many cases both radical openness and contextuality can be easily overcome, but then there are cases when this is no longer the case. In this latter case we have complexity.

Common misconceptions about complexity

There exists a vast body of literature on complexity. Much of it is concerned with modeling, some of it probes the notion of complexity itself. Either way, there are two themes that dominate much of this literature and have perhaps become accepted as an integral part of complexity: the notion of “complex system” and the idea of measures of complexity. At least as far as the intuition about complexity underlying this article is concerned, these ideas are not useful.

Complexity and complex systems

The notion of complexity is often associated with the notion of system (as, for example, in the term “complex systems”). The idea of systems is widely used in the sciences and normally refers to a well-defined part of the world which is in some sense independent from the rest of the world (the “environment”). Complex systems are then just the same only that the systems are complex rather than “not complex.” Often cited examples of such complex systems are organisms, economies, or the climate, to name but a few.

The concept of “system” seems rather innocuous at first, and within the context of everyday scientific communication it is. There is no harm done in referring to, say, a particular operon in *E. coli* as the *system*, or to the economy of Germany as a *system*. Within the normal discourse in science, everybody knows what is meant. Unfortunately, in the context of complexity (as we understand it here) it will become clear that the notion of a *system* being *complex* is misleading. The reason is that the difficulty of modeling (i.e., the complexity) does not primarily depend on some inherent properties of real world entities, as the notion of a “complex system” might suggest. This notion implicitly assumes that there is a natural partition of the world into system and non-system (or systems and environment) and that this partition can somehow be discovered by the modeler. For every modeling exercise the modeler picks out one of these systems and represents them mathematically or

otherwise. If she is unlucky, then the one she picked is complex and therefore quite difficult to represent.

We think it does not take a long argument to convince the reader that this picture is not particularly realistic. At the heart of the problem is the idea that there are natural systems out there or, as we will call them henceforth, *systems an sich*. Instead, we will argue below that any particular system definition is a choice that mostly depends on the specific purposes that motivated the modeling exercise. Once we accept this, then complexity is first and foremost dependent on these purposes rather than the innate properties of any system *an sich*.

This idea that complexity is a property of the modeling purpose rather than of the “system” to be modeled will be argued over the remainder of this contribution. For the moment we will illustrate this point using an example. Let us assume that it makes sense to assume that complexity is a property of a system *an sich* and for the sake of the argument let us also assume that organisms are such complex systems. We now note that the alleged complexity of organisms does not prevent anybody from modeling some of their aspects or parts, sometimes to a high degree of accuracy using very simple models.

For example, gene expression can be modeled using a single differential equation. Such models, simple as they are, can give very good insight and are regularly used in the theoretical biology research literature. So apparently no complexity there.

One can of course argue that the complexity of an organism comes from the interaction of the genes. Genes themselves are simple, but many genes together are not—organism are not (there are some echos of complexity out of simplicity here). This may well be true, but then again, it is not inconceivable that in some modeling contexts entire organisms could be represented extremely simple. How much detail one takes into account in models depends entirely on the perspective one takes that is on what needs to be taken into account. To see this let us consider another way to model the very same gene whose expression we captured by a single differential equation, but now we want a detailed model of the polymerases reading the DNA and all its chemical interactions. Suddenly, we will find ourselves bogged down by dependencies and interactions that stretch our technical ability to the limit.

We conclude that It is possible to model gene expression in a simple way by leaving out most of the detail. It is also possible to make models of gene expression nearly arbitrarily difficult by including more and more detail or indeed by including more and more of the processes on which gene expression depends. There is nothing inherently simple about a gene and the process of gene expression, but there are ways to represent genes in a simplified matter that makes it possible to study at least some of their aspects

using simple models. Similar ideas apply to any kind of modeling exercise.

The question we have to ask now is the following: How do we choose what is included into a model and what is left out, how do we pick those phenomena and effects that are important enough to be included and those that are not? The key message that we will argue over the next pages is that complexity is connected to the difficulty of making this choice that is the difficulty of deciding what needs to be included into a model and how, and what can be neglected and left out.

This difficulty of creating and justifying the design of a model (i.e., complexity), does not depend on the properties of the world but on the purposes of the model and the kinds of questions the modeler wishes to address. We can say this more concisely that complexity depends on the modeler’s purpose not on some properties of a “system *an sich*.” This idea will be central to this article.

Complexity as size of a model

Before we can go on to develop a notion of complexity based on the difficulty of modeling, a few comments on previous attempts to define complexity are in place. Within complexity research there exists a tradition to quantify complexity. Measures of complexity are normally arrived at by quantifying some aspects of a model and to take this as a proxy for the complexity of an assumed system of which the model is a representation. Measures of this sort have been extensively discussed in the complex systems literature [see, for example (Edmonds 1999; Rescher 1998; Gross and Strand 2000)] and this contribution will therefore refrain from reviewing quantitative notions of complexity. Suffice to say that after more than a decade of research into complexity, there is still no agreement on any preferred measure of complexity. Common to all measures is that they in one way or another attempt to put a number to the size of a model, where “size” could mean the actual amount of information to describe a system (Bennett 1986), the number of components or types of components (McShea 1996), measures of the interactions of individual components, or a mixture of those. The difference between the individual measures is how these various indicators of size are combined to form a particular measure of complexity. There are many ways to do this, and perhaps there are nearly as many measures as there are authors. Unfortunately, various notions of complexity tend to contradict one another in the sense that rankings of models according to their complexity will be very sensitive to the particular measure used. While a model *X* may be more complex than model *Y* by one measure, another measure may rank *Y* as more complex.

In the practice of modeling considerations of the size of a model are of course important. Large models tend to be

computationally more expensive to analyze and they usually also require more parameters to be determined, and take more time to code/formulate. A large part of the work of a modeler is precisely to find simplifications that reduce the size of models while preserving the particular key-features and behaviors of the model that she is interested in. Indeed, the success of a modeling venture will depend on the modeler ignoring most of the features of a system she is modeling, primarily concentrating her attention on a few aspects only. Conversely, one could make any model nearly arbitrarily large (in terms of size) by choosing to include irrelevant features that do not impact on the final result.

Intuitively, one would be inclined to believe that there is a strong correlation between the minimum size of a model and some fundamental properties of the natural system the model is supposed to represent. Apart from some natural variations due to different skills of modelers—one may argue—the size of the model must certainly reflect *some* property of some system *an sich*. At the very least, one would expect that there are models of minimal size, for a specific system. This intuition is, presumably, also what motivates quantitative measures of complexity based on properties of a model. However, we think this intuition is misleading.

Before a modeler can formulate a model and interpret it in the way one would normally find it published in scientific articles, she has to go through a long and difficult process of struggle. Scientific reports tend to give the impression that the choices that went into a particular model are self-evident and entirely determined by the nature of the phenomenon of interest. This could not be further from reality. The finished product is the result of a process in the course of which the modeler modified and re-modified the model many times; she will have simplified, approximated, assumed, and omitted. The nature of these simplifications, approximations, assumptions, and omissions will be partially dictated by the purpose of the modeling exercise, but partially they will be choices that are not necessarily rational or logically justifiable. We conclude that any particular model of a system is not an unbiased representation of some objective reality.

On the other hand, the idea of taking the size of a model as a shortcut for the complexity of some natural system tacitly assumes that this process of formulating the model leads to an unbiased representation of the system, at least as far as complexity is concerned. In particular, the idea of “measures of complexity” implies that there is an objectively defined “thing” in the world (the system *an sich*) to which the modeler has access. This *system an sich* also has a degree of complexity, and the size of this model of the system *an sich* reflects this complexity property. It has been indicated above, and will be argued below that this idea of

a system *an sich* is incoherent. A more coherent view is to accept that models are not truthful representations of the real world, but they are abstractions that reflect a particular purpose that motivated the modeling exercise. The system is then defined by the model, rather than the model being defined by the system.

In this article, we will therefore dismiss quantitative notions of complexity, based on properties of models, as irrelevant. Instead, we will develop and defend an alternative qualitative notion of complexity with the aim of connecting complexity to the difficulty of finding suitable abstractions of systems. This shifts the notion of complexity away from being a property of models or elusive systems to being a property of the process of finding suitable models. We will argue that this interpretation gives a more coherent view of complexity and is also of practical use in that it helps understand why sometimes it is harder to come up with a model than other times. Finally, we will illustrate this view of complexity using an example from computational modeling.

Formal and semantic models

Scientifically (or otherwise) we cannot reason directly about systems *an sich*. There is no sense in which we can mentally manipulate objects of the real world. All we can do is to manipulate *representations* of these objects—models. A corollary of this is that we cannot even compare our models with systems *an sich* directly, but only with other models. This is perhaps a commonplace, but at the same time also a fine point in many contexts. We suspect that much of the confusion about complexity is due to an unclear notion of models and how they relate to the “real world.”

Working scientists distinguish between system and model as a matter of routine and spend a great deal of time worrying about how well the two match. Fundamentally, this is of course sloppy thinking. Even scientists do not have direct access to any system *an sich*, but they can only compare the outputs of their measuring devices with the predicted values of their models. In many contexts, it is acceptable to be sloppy about the precise notions of and relations between systems and models. In this contribution, we cannot afford to be sloppy, and we must make clear distinctions. We are interested in complexity, and we want to understand how models are created and what makes this process of model creation difficult. We will therefore need to take a closer look at models and their relation to the *system an sich*. To do this, we will motivate the distinction between *formal* models and *semantic* models. We will argue that the former are, in a sense, a minor element in the fabric of scientific knowledge/understanding which is

woven from semantic models. Formal models are primarily crutches to help reasoning, but do not contain any understanding in themselves. They need to be interpreted in order to be useful. Once we have made this point, complexity measures that focus on properties of formal models will seem unnatural.

As a computational or mathematical scientist, one tends to primarily think of models in terms of *formal models*, i.e., systems of equations, computational simulation programs and so forth. Within certain communities in the natural sciences, the word “model” is nearly universally identified with such formal models. The main purpose of formal models is to state relationships between variables. These could be formulated in a variety of ways, including systems of equations, sets of chemical reactions, computer code to name but a few possibilities. Common to all these formal models is that they are strictly objective and once specified there can be very little disagreement about their properties. A computer program, once written, will give more or less the same results wherever it is run; equations have an objectively verifiable set of solutions.¹ Not everybody may be able to find these solutions, but once a solution is found a consensus concerning its correctness can quickly be established.

There are many different types of formal models, including differential equation models, computer simulations, network representations, lists of chemical reactions, and so on. In terms of the practice of research, there is much to be said about the advantages and disadvantages and the various types of formal models. Mathematical models usually allow more general insights than simulation models, but also suffer from tractability problems. The choice of the correct type of formal model can be crucial for the success of a project. However, as far as our present purpose is concerned, various types of formal models only differ in how they are solved, but are equivalent with respect to their role in the modeling process overall. Moreover, all formal models can be represented as computer simulations, be it analytic mathematical models, equation-based models in physics, or models of chemical reaction systems. Conceptually, we may therefore (for convenience) identify formal models with computer models, even though in practice there is a big difference between the various types of formal models.

¹ The issue is not always quite as clear-cut. Particularly when it comes to large-scale computational models there can be significant problems to create a consensus on the properties of the model when it is run on different platforms, or even differently configured versions of the same platform. These problems certainly do exist, are interesting and indeed important to consider and are also relevant for complexity. However, as far as the following discussion is concerned, they are an unnecessary complication in that they do not interfere with the conclusion we will reach. We will therefore assume the simplest case that formal models are unproblematic.

While formal models are conceptually unproblematic, they are also, *by themselves*, scientifically useless. The process by which formal models come about may require a significant degree of ingenuity and creativity. Once the formal model exists, then it is just a (normally quite complicated) set of re-writing rules that can be applied to some input (for instance an array of number) to produce an output (which is often also an array of numbers). In and by themselves, these numbers are, of course, just numbers. To be informative for the modeler they need to be processed. This processing step may involve some type of formal processing, statistical, graphical, or otherwise. Yet, formal processing of a set of meaningless numbers only gives another set of meaningless numbers. To be of any use the output of the formal model must be interpreted in the context of existing models and woven into the fabric of scientific understanding.

This scientific understanding/knowledge is contained within the *semantic models*. The precise nature of these semantic models is a difficult question. For example, it is not entirely trivial to understand how they are created, communicated, and processed internally. For the purpose of this contribution it is not necessary to clarify these potentially difficult philosophical questions in much detail, although we do acknowledge that this needs to be addressed eventually. Suffice to say that semantic models form the main part of our “understanding” of the world and of scientific theories. Formal models are just crutches to help our reasoning. Semantic models are the real keepers of scientific knowledge.

Semantic models are inexact and fallible. There can be no precise agreement on semantic models because they cannot be shared in an unambiguous manner between scientists; this is simply a limitation of natural language. Loosely, in scientific articles reporting modeling one can identify the semantic models with the text that is in-between the formulas, algorithms, and definitions. One can assume that semantic models are (up to errors of reasoning and imprecisions) largely consistent with the formal models, but they are of course not equivalent.² Sometimes there are mismatches between semantic models and formal models leading to incorrect interpretations of the results. An example is reported by McMullin and Varela (1997). We are content with the following characterization of semantic models:

- Semantic models are representations of scientific knowledge. They are internal to the agent.
- Internal reasoning on semantic models is possibly imprecise and always fallible.

² The various interpretations of Quantum Mechanics are an example of disagreements on the semantic models while there is an agreement on the syntactic models.

- There is no clear distinction between semantic models relating to scientific activities and trivial semantic models about everyday life.

Crucial to the notion of complexity that we wish to develop here is to understand how semantic and formal models relate to one another. A formal model is an abstract representation of a semantic model. In order to arrive at a formal model, the modeler will go through an intense process of simplification during which she will ignore most of the aspects/contents of semantic models concentrating on a few select aspects only. For one, the modeler needs to limit herself to a specific domain of interest (say, gene expression in mammals, chromatin re-modeling, etc). Once this domain of interest is chosen, she then needs to decide on a level of abstraction. This entails deciding on the desired detail of the model, the dependencies and interactions that are taken into account and those that are left out. Both of these steps are essential to ensure feasibility of the formal model.

Normally, formal models encode only a small subset of the semantic models pertaining to the particular domain. So, for example, a great deal is known about gene expression, the function of polymerases, the chemical interactions necessary for initiation of the transcription process, the interaction of the polymerase with the DNA, and so on. The modeler will have semantic models of all of these processes and interactions. Yet, in a formal model of gene expression many of these known interactions may be ignored; a formal model may represent gene expression simply as a single differential equation.

The system *an sich* does not feature anywhere in this. However, to facilitate the following discussion, it is helpful to re-introduce the notion of a system, although with a slightly changed meaning (compared to the idea of system *an sich*). Every formal model will have an interpretation which will be a semantic model that exactly corresponds to the formal model. One can identify the system with this particular semantic model. In this sense, one can talk about systems without having to refer to the unknown world of systems *an sich*. This definition is also compatible with the colloquial use of the word “system;” yet we stress that the system in this sense is not an immutable thing out there, it is not the system *an sich*, but it is choice that emerges from the struggle of the scientist to develop a model. In this sense, the system is determined by the modeling process rather than the model being determined by the system.

How does the modeler arrive at the system definition, or equivalently, how does the modeler choose the formal model? This process of extracting the key-features of a system is not a process that follows rules. Instead, it is an iterative process of trial, resistance and success that eventually leads to the system at hand. A very enlightening

account of this process is given by Pickering (1995). Neither do we require such a description to understand complexity. Yet, there is one key-factor that above all shapes the abstraction process: The purpose of the modeler, i.e., the aims and objectives that motivate the modelers choice of what to include and what not to include into the formal model/system.

If one wishes, for example, to model the melting point of DNA molecules, it is not necessary to model quantum mechanical effects, even though the interactions between the nucleotides, and hence their melting point, are ultimately determined by these interactions. Instead, one will seek to remove nearly everything known about the DNA and only represent those features that are relevant for the melting process, which in this case is the base-pair composition of the molecule together with information about the stability of the individual bonds. On the other hand, a formal model with the purpose of understanding how transcription-factors find their binding sites will be different. It will normally not consider the binding strength between base-pairs at all and instead only focus on the sequence of nucleotides on the DNA string and calculate from this the protein–nucleotide interactions. The purposes behind these two models are very different from one another, motivating very different formal models and thus systems.

Once one has chosen a formal model, the question arises how one can know that this formal model is useful, in the sense that it either reflects the quantitative behavior of the system accurately, or at least provides a qualitative explanation for it. Again, this question is fraught with philosophical difficulties; treating these in detail would not only go well beyond the scope of this article, but would also be distracting from the main question we are trying to address here. Without committing to any particular view on this we can state that laboratory system often plays a role in instilling confidence in a model. Laboratory systems are real world entities that are artificially screened from environmental disturbances, usually by elaborate experimental techniques. Key to laboratory systems is that they can be manipulated, if not at will, then at least to considerable extents. The primary aim of experimental science is to be able to create and manipulate ever more sophisticated laboratory systems. For the modeler experimental systems can be the reference against which she assesses the quality of her models. They may enable her to assess the usefulness of the abstraction. Ideally, the laboratory system is a real world implementation of the formal model (i.e., it realizes “the system”). Laboratory systems are an important tool in science to check the usefulness of models, but they are not the only way to do so, nor are they unequivocally useful.

In summary, the abstraction of the formal model from semantic models, i.e., the choice of system, will crucially depend on the purpose of the model, the aims, and objectives of the modeler. The degree to which details are represented and most of all which details are represented is heavily influenced by the specific modeling purpose motivating the model. This entails that the “size” of the formal model, however defined, is not indicative of any properties of an independently existing system *an sich*. Attaching a measure of complexity to a formal model may be useful to find out how hard it is to analyze the model, how long it takes to simulate the model or how difficult it will be to find solutions. Yet, such a notion of complexity is more akin computational complexity measures (i.e., the time complexity of an algorithm), which is conceptually a very different thing to the kind of complexity we are after.

Complexity

In this section, we aim to develop a notion of complexity that conceptually captures and describes the difficulty of creating formal models from semantic models. Before doing so, however, we will briefly introduce a concept of complexity originally developed by the late theoretical biologist Robert Rosen. During his life, Rosen has published a number of forward looking books on theoretical biology. One of the *leitmotifs* of his writings is complexity. It would go well beyond the scope of this contribution to give a comprehensive summary of the notions of complexity due to Rosen. Suffice to say that his work has gathered a small but dedicated followership among scholars from all corners of the academy. By far the best known of Rosen’s works is his book “Life Itself” (Rosen 1991). One of the central themes of this book is the development of a particular notion of complexity. The book is widely read, and its notion of complexity makes a very good starting point for the wider exploration of complexity. We will therefore discuss it in some detail here, however, without claiming that this is a comprehensive review of Rosen’s work on complexity.

Rosen based his notion of complexity on organizational properties of systems. So Rosen still assumes that the notion of a system *an sich* is meaningful. However, rather than analyzing any particular model of his system, Rosen considers the space of all possible models. He then asks whether or not within this space of formal models there is at least one that represents the system as having the particular organizational property of being closed with respect to efficient causation (CWEC); if it does, then the system is complex, if not, then it is not complex. Rosen’s approach avoids many of the pitfalls of quantitative notions of complexity, in that it judges complexity based on a

spectrum of possible formal models, rather than on a single formal model. The problem with this approach is that it is still based on the idea of an independently existing systems *an sich* constraining the choices of the modeler. As discussed above, this is a problematic assumption.

According to Rosen, a system is complex if it is CWEC.³ The idea of CWEC is often illustrated in terms of (abstracts) models of metabolic systems. At one level, one can think of metabolic systems as converting some input into output. Rosen represents this as mappings from some input set A to an output set B . This seems to be very similar to familiar mathematical representations of biochemical systems, except for one thing. Rosen recognizes that the mapping between A and B , let us call it f , must be somehow realized by a (material) component in the system. This is what Rosen calls the *efficient cause* of the conversion from A to B .

Once we accept that f is not just some mysterious action in the system but must be implemented by an actual component, then we must also accept that f could fail, needs to be created and will decay over time. This begs the question of the origin of f , how it is maintained or repaired. To use Rosen’s notation, we can ask about the efficient cause of f itself. To solve this, Rosen introduces, abstractly, a new component, ϕ that maps B into f ; Rosen calls ϕ the efficient cause of f , and f the efficient cause of B . In this new, expanded system every component has now an efficient cause, except for ϕ . There are two ways to solve this. First, one may introduce yet another component to repair/replace ϕ ; this is not satisfactory, because, as can be easily seen, it would just create another component that is not efficiently caused and thus lead to an infinite regress. The second possibility is to close the circle by postulating that one of the components already within the system acts as the efficient cause of ϕ . This then leads to a system where, in Rosen’s parlance, every component is efficiently caused by something within the system—the system is CWEC.

Rosen and his commentators have been particularly interested in a particular type of CWEC system, the so-called (M,R) system, which is allegedly the smallest possible CWEC system. We will not discuss this in any more detail here. Detailed discussions of (M,R) -systems and CWEC can be found in the primary literature (Rosen 1991; Rosen 1999), but also in a number of reviews of Rosen’s ideas (Chu and Ho 2006; Chu and Ho 2007b; Letelier et al. 2004; Wolkenhauer 2002). This will not be repeated here. However, a short reminder of the main idea of CWEC is appropriate.

³ It has to be pointed out that complexity is not Rosen’s primary interest, which is directed towards organizational principles of living systems. Complexity is more an accidental fallout rather than the main focus of his investigations. Hence he formulates complexity, rather specifically, in terms of abstract models of metabolic systems.

1. (M,R) systems are the simplest organizational realization of an organism.
2. The organizational features of (M,R) systems cannot be implemented in computers.
3. A system is complex if it is CWEC.

The first statement reflects Rosen's interests as a theoretical biologist. More important for the current purpose is the second statement. Rosen (1991) sketches a formal mathematical proof that Turing machines (and thus by extension all computers) by their very nature cannot have this CWEC property. This proof has been controversial in the past and is most likely wrong (Chu and Ho 2006, 2007a). We will argue below, however, that the conclusion is nonetheless meaningful. The third statement is an interpretation of the second statement in that it defines a system as complex when it has this CWEC property. Given that, as Rosen claims to show, computers are not CWEC, this concept of closure possibly tells us something about the limitations of models, and hence about complexity.

At this point we have to be careful because what we are really interested in are properties of models. It is easy to see that formal models are not CWEC. In the case of computer models, the model itself could be identified with the program code that defines it. Instead of entering the program into a computer to compile and run it, one could write it on a piece of paper. The information on this piece of paper would be sufficient for anybody familiar with the relevant programming language to understand what the model is about. However, by itself, the code is not doing anything. It just is some text on a piece of paper. In order to be of any use, the instructions specified by the code on the piece of paper, need to be executed, or in Rosen's word, it needs an efficient cause. This requires some external agency or simply put a computer. So, nearly trivially, formal models are not CWEC.

While formal models themselves are not CWEC, the question we are actually interested in is whether or not formal models can represent systems that are CWEC. We think they can, or at least there is no reason to assume that they cannot. Indeed, there are a number of previous contributions reporting specific realizations of (M,R) systems (Letelier et al. 2004; Letelier et al. 2006; Wolkenhauer 2002; Wolkenhauer and Hofmeyr 2007) as formal models. A potential problem with these *in silico* implementations of (M,R) systems, is that they rely on an external efficient cause for execution. Typically, the components in these models are matrices (to represent mappings). Clearly, matrices do not map anything, but they are rather instructions for how to map and the mapping is then done either by hand or by a computer. As such, these models or realizations of (M,R) systems probably fail to demonstrate anything because they do not contain an efficient cause

within themselves (only descriptions of the action of an efficient cause). We can think of more convincing ways to formally model (M,R) systems. One way to do this is to separate the model of the efficient cause from the efficient cause of the simulation by one step.

Let us clarify this. One concrete possibility is to model the components of Rosen's (M,R) systems as dynamical hierarchies (DH) (Gross and Lenaerts 2003). This means that the simulation is based on some simulated "elementary particles" that interact with one another and through their interaction can aggregate to higher level objects—the DH. The DHs have their own properties and behaviors and can themselves aggregate to even higher level DHs. The relevant point is that the properties of these higher level objects are not explicitly programmed by the modeler (although implicitly they are of course). The modeler only programs the artificial physics of this simulated world. Each of these DH can then be seen as causally depending on lower level DHs (or elementary particles). The dependence on the efficient cause of the executing machine is only indirect. If the components of Rosen's (M,R) system— f , ϕ and B —are implemented as such DHs, then this would make a convincing case of a computational model of (M,R) systems, and thus of a complex system *sensu* Rosen. Admittedly, so far nobody has been able to implement CWEC in a computer model in this way, but this does not mean that it cannot be done.

We conclude that computer models cannot be CWEC themselves; this is a nearly trivial point. However, formal models can represent systems that are CWEC, which as modelers is all we want. This indicates that CWEC is not a particularly strong condition. It is also unclear how CWEC is relevant in the context of climate change or evolution, domains that are often associated with complexity. This leaves one now wondering why one should tie the notion of complexity to CWEC. CWEC may or may not be interesting as a description of organizational principles of living systems, but it is not clear that CWEC is in any fundamental way limiting the modeling process. There is no obvious reason why one should see it as the cause of the perceived difficulty to model/explain/predict real systems. There is nothing within the concept that points to limitations of formal models or limitations of the modeling process itself. Its relevance for complexity is therefore hard to see. For this reason, we reject Rosen's complexity as a useful notion of complexity.

Beyond CWEC: radical openness

"Closure with respect to efficient causation" means that everything within the system can be explained, in efficient causal terms, from within the system. We have seen that

this is in fact not a particularly strong condition and could be achieved within formal systems (although with some difficulty). While CWEC is therefore by itself not suitable as a basis for complexity, Rosen's approach is pointing into the right direction. Closure properties of systems are relevant in the context of models and their interpretation, and they crop up everywhere in modeling. For example, experimental physicists go to great lengths isolating their laboratory systems from interactions with the environment, which is in essence creating closure—one kind of it. Modelers in economics or biosciences will likewise place great emphasis on finding abstractions that are *approximately* closed, at least over certain length and time scales. Imposing this type of closure is a pre-condition for feasible models.

To probe into these closure properties of systems, let us extend Rosen's analysis of (M,R) system. Rosen himself was primarily interested in the idea of efficient causation that is how components are maintained, while he was not interested in other aspects of the (M,R) system, such as, for example, the input A . This is of course not the only possible focus. Instead of asking about components, say the component f , one could ask about the *mapping* f in the system and try to explain it. Why is it f rather than a different mapping f' ? This question has nothing to do with the efficient cause of the component f but addresses the function of f and why it is needed within the system. One could answer this question by referring to the role of f in the (M,R) system and the necessity for the element B for this system to work, the availability of A , the input requirements of ϕ , and so on. This kind of question can be asked for every component of the system, and the answer will always involve some kind of reference to the function of the system as a whole.

In addition to asking questions about the components of the (M,R) system one can also wonder about its input. Why A ? This is a question that Rosen dismissed as trivial, but we do not think it is. There are a number of possible ways to explain A . One possibility is to answer by reference to f and refer to the fact that f can only accept specific types of inputs. So, A is an input to the system because this is what f requires. This only explains why A is taken as an input, but not why the elements of A exist in the first place and where they come from. In order to understand the origins of A , it is necessary to extend the focus of the investigation beyond the (M,R) system.

Let us examine a few possible answers to the "Why A ?" question. It is conceivable that there are no further explanations for A because A can contain everything in the world that is not f , ϕ or B . While a theoretical possibility, it is not clear that the corresponding model would be very interesting. The elements f , ϕ , and B would need to be enormous and encapsulate nearly everything there is (except for A). The second possibility is to consider A as a limited

subset of what there is outside the system. In this case then we can ask another question about A , namely, where it comes from, and why it exists. One possible answer is that the material components that make A could just be in existence; they are a natural resource that does not require any further explanation. This would only be plausible if A is an elementary particle whose existence cannot be further questioned. This will be the case only in the rarest of modeling projects.

A more relevant scenario is that the resource A is produced by another system, let us call it $(M,R)'$ system. In this case, the amount of A will depend on the production rate of the resource, and its existence has to be explained with reference to $(M,R)'$. Yet, explaining A in terms of $(M,R)'$ does not solve the problem, it merely shifts the explanatory burden onwards, which leads to an infinite regress, similar to the one Rosen has identified. Only, in this case we see no natural way to close "the system" and avoid an infinite regress unless all included in the extended model.

The process of asking for more and more explanations dissolves the boundaries between the (assumed) (M,R) systems that we started with. We find that there is no natural justification for considering the (M,R) system as a system *in sich*, i.e., as having some special status of autonomy or independence from the environment. The original focus on (M,R) system is entirely an artifact of Rosen's specific interests in the organizational features of organisms. As this interest is shifted, or as the purpose of the model changes, so will the particular shape of the model. By asking questions about A , what was previously a system, is now only a component of a larger system. We achieved this simply by changing our interests, by deciding that we are interested in the origin of A and where it is produced, rather than, as Rosen did, focus on efficient causations. Both interests can be motivated, but ultimately both are arbitrary in the sense that none is, in a universal sense, superior to the other. At no point did we assume that the "world" changed, but the systems did.

More significantly, as we have extended the system, we may also choose to reformulate our questions completely and cease to insist on CWEC. We could develop very different interest and decide that components such as ϕ in our original (M,R) system no longer are interesting. By doing this, we re-draw the boundaries of the system in a way that may cut across boundaries of the original (M,R) and $(M,R)'$ systems, creating completely new systems in the process. How we re-draw these boundaries is solely dependent on our interests.

To reach this conclusion, we started off with (M,R) systems, but this was of course completely arbitrary. No matter what our starting point, we would always have come to the same insight that interests and the purposes of

the modeling primarily determine our choice of what is and is not included in the model. This is precisely what has previously (Chu et al. 2003) been called *radical openness*. Everything connects to everything else and the notion of a “system” only exists insofar as it is a pragmatic choice valid for a particular modeling project.

The intuition behind radical openness is perhaps not surprising, or particularly novel, but it certainly is an unusual perspective for working scientists. Whenever a researcher studies a particular system, she will instinctively ring-fence the elements of her curiosity and assume nearly everything within the domain of her interest as being fixed. She will narrowly focus her attention on one particular question. This is also what Rosen did who assumed that everything outside his (M,R) system is given. In the actual practice of scientific (or applied) modeling, there is no need to explain everything, and it is good practice to assume most aspects and elements of the model to be fixed, even though this is only an approximation.

In the context of complexity, things are different and the idea of radical openness is relevant because it is a force that is “opposed” to the process of choosing boundaries between system and ambient and between focal phenomena and the environment. This pragmatic process is governed by choices and purposes including the aims and objectives of the model, what kinds of questions it needs to address, but also tractability issues in the case of formal models and the availability of laboratory systems to test the model. There is however no guarantee that for every modeling purpose there is a clear choice of system that can be contained within a useful formal model. Moreover, even if there is one, there is no guarantee that for all purposes there is a procedure to determine whether or not a particular choice of formal model is suitable for the purpose of the modeling exercise. When containment fails, then everything connects up with everything else and formal modeling will become very hard.

Based on this, we conjecture that radical openness is a cause of complexity, in the sense that there are circumstances when this radical openness cannot be contained, and it is either not possible to identify a system, or it is not possible to be certain that a specific choice of system is fit for purpose. Radical openness is always present and not limited to certain “systems.” Indeed radical openness is an anti-thesis to the notion of “system.” However, whether or not it does affect the modeling depends on the aims and purposes of the model.

Contextuality

The idea of radical openness is that there are no natural boundaries that define systems; there are just semantic

models, all of them connected together. A major aspect of the abstraction process leading from semantic models to formal models is to find suitable boundaries of models, i.e., the modeler has to make a choice as to what to include and what to leave out from the model.

Drawing boundaries are not sufficient. The other task that the modeler needs to accomplish is to find a suitable encoding of the semantic models that is a representation of the semantic models as a formal model. This process of encoding has two components. First, a technical part that includes a choice of the appropriate modeling technique (i.e., whether to use differential equations, agent-based models, etc.). In practice this decision can be very important, but for our purposes more interesting is the second aspect of encoding. This second part is about deciding the level of abstraction that one requires. This step involves ignoring much, or even most of the semantic models and concentrating on a simple, yet suitable abstraction that allows feasible formal models, while at the same time representing sufficient detail of the semantic models to make the model credible. It is worthwhile stressing here again that the choice of what needs to be included into the formal model and what is left out is primarily dictated by the specific purpose of the model and the aims of the modeler, but is also influenced by considerations of tractability.

Each formal model defines a system, as discussed above. If one now considers the space of all systems/models that have been developed up to a certain point, then one will notice that there is an overlap between some of these systems. This observation motivates the notion of *contextuality*.

There are several ways in which models can overlap. One possibility is that two models (or rather the parts of the models that overlap) differ only in trivial ways and are otherwise the same. The differences may only be due to the requirements of the particular modeling methodology, i.e., choices such as the programming language or the formalism used, or they may be minor choices of the model design that do not in any way influence the behavior of the model.

The difference between two models may go deeper than that and there could be significant differences in the levels of detail used in the model. One model A may be a coarse grained version of the other model B . So, for example, model A could be a model of the DNA as a string of Watson–Crick pairs, whereas B is a detailed model of the same DNA that takes into account the chemical structure of the individual nucleotides. Both A and B will overlap in that they represent the same conceptual entity in roughly the same way, but to different degrees of detail. In this case, it is possible to compare both models and to decide whether or not the difference in representation is of practical significance. Such considerations are routine during every modeling project.

As far as contextuality is concerned, the most interesting possibility is when the two representations are *orthogonal*. This means that models *A* and *B* both contain abstractions from the same conceptual entity, but they are not in a clear refinement relation to one another. In this case, it is no longer meaningful to compare models *A* and *B*; they are incommensurable in that their abstractions are motivated by different purposes and goals. While these models are not in a refinement relation to one another, it may still be possible to have a model *C* that combines the orthogonal representations of the shared conceptual entity in *A* and *B*.

It is helpful to illustrate this using an example. We choose as an example various ways to model DNA molecules which can be conceptualized in different ways.

- *DNA as code* In many models, DNA is represented as a carrier of information or as a code. This representation is particularly popular in models of simulated evolution. In such models, DNA could be implemented as a binary string encoding some information [see, for example, (Mitchell 1997)]. More sophisticated models use a 4-letter alphabet to represent the string [see, for example, (Chu et al. 2005)]. Technically, DNA represented in this way could be implemented as an array, a data structure that is supported in many programming languages.
- *DNA as a physical object* Bio-physicists may be interested in DNA as an object in space and time which takes specific shapes and conformations depending on the ambient conditions [see, for example, (Sneppen and Zochhi 2005)]. One may ask questions about its properties, how fast it folds, how flexible it is, and so on. Such models will not need to represent the information carried by the DNA, but only the properties of the molecule as a polymer, its flexibility, its interactions with the environment, and so on. In these models, one could conceptualize DNA as a generic polymer consisting of a number of links that are allowed a certain degree of freedom relative to each other. Simulation models could take into account chemical properties of the individual base-pairs and how they interact with one another, but there is also some scope for equation-based approaches. Simple models would embed the DNA in 2D space and more sophisticated ones would take into account 3D movement of the individual base-pairs. The representation of DNA in this case cannot simply be an array, but it must contain spatial aspects of the individual monomeric elements of the DNA, and how they interact. While the representation is likely to be more involved than in the previous example of DNA as information, models of DNA as a polymer are not refinements of models of the DNA as information, because they focus on entirely different aspects.
- *DNA as information carrier* One may ask about the properties of DNA as carrier of information. What are the physical and chemical properties of DNA that enable it to store information? What are the properties of base-pairs and their interactions that allow the DNA to fulfill the role it has in organisms? Models to address this type of question need not take into account what the information contained on the DNA is, but they must consider the properties of the constituent parts that allow this information to be preserved. So, one needs to worry about the stability of DNA molecules, how mutations are prevented (see, for example, (Sneppen and Zochhi 2005; Eigen and Schuster 1979)), and so on. The flexibility of the DNA as a polymer is not of interest here. Formal models to address this question could use equations describing local interactions between base-pairs, and probabilistic models of the mutation rate, possibly also information theoretical models to understand the principles of information storage and transmission. These formal models would not represent spatial conformations of the DNA nor the information it contains.
- *DNA as a string of binding sites* Interaction of proteins with DNA require the former to find specific binding sites on the DNA (Chu et al. 2009; Wunderlich and Mirny 2008). The process of proteins locating their binding sites can be modeled as a random walk. DNA is then represented as a string of binding sites for transcription-factors. In these models, a binding sites corresponds to a sequence of nucleotides, and the strength of the binding between the nucleotide and the protein depends on how well this sequences matches an optimal binding motif of the protein. Again, one could represent DNA as an array in computer models, yet the interpretation of this array is very different to the above case of DNA as code. Again, such representations are orthogonal to the representation in the other models.

Each of these abstract representations of DNA can be ring-fenced, in the sense that one can exclusively focus on one while ignoring the possibility of others. The biophysical properties of DNA can be studied in isolation from the information the DNA carries, and the chemical properties can be investigated without having to worry too much about how resistant DNA is to mutations. Sometimes, there is partial encroachment of one aspect onto another one. For example, if we want to understand how proteins scan the DNA, then the spatial arrangement of the DNA as a molecule may be important. Yet, often this encroachment can be contained by crude approximations. For example, binding sites are often represented as rods distributed in 3D-space (Slutsky et al. 2004). This mimics the limited flexibility of DNA (individual rods) and the spatial arrangements of the molecule.

Some modeling projects may require that aspects are combined; for example, one may represent DNA as code and its spatial arrangement in the cell. In general though, focusing a model on one (or few) contexts is arguably essential to keep the scale of the modeling project within the limits of feasibility; this may not always be possible.

The real fundamental problem contextuality can pose is when it is unclear which particular abstraction is the correct one that is which simplifications are suitable and which ones are not, which contexts must be included and which ones are irrelevant. One situation when this could be the case is when there is no laboratory system available that can be used to test the model, although it should be stressed that this is neither sufficient nor necessary a condition. Yet, when there is no laboratory system, then it is often difficult to gauge the suitability of a given abstraction for the particular modeling purpose. A topical case in point is climate change modeling which suffers from tremendous uncertainties precisely because there is no way to check a model against a laboratory system. Interestingly, even in this case it will still be possible to check whether or not the results of the formal model are correct given any particular model definition. The process of developing a model is where complexity originates.

In the previous paragraph, we implicitly assumed that leaving out contexts may sometimes severely impact on the performance of a particular model relative to its specific purpose. Given the scope of this article, we will not be able to fully justify this claim in this contribution, although we will make a few remarks. A common problem of models, particularly, predictive models that can be traced back to missed contextuality are unforeseen side effects. Such side effects are usually a mismatch between a predictive model of an intervention into a system and the actually observed behavior. Another way in which contextuality makes itself felt are couplings between phenomena that have been ignored in the model. Seemingly unimportant contexts can prove to couple the system of interest to apparently unrelated phenomena; missing this can render the entire model invalid. This is no problem as long as the model can be corrected, but could be a big problem if there is no possibility to test the model before it is used. A few examples elucidate this:

- The development of novel drugs suffers from the lack of a suitable laboratory system to test their effects. To be precise, the interaction of drugs with their environment can be tested in laboratory systems and can to some extent also be modeled formally. It is not particularly challenging to understand how various proteins interact with specific target molecules in the laboratory. The results of *in vitro* research do not always translate well into the *in vivo* situation. When

drugs interact with the living organism then new aspects come into play. The active ingredient may take a number of different roles *in vivo* and lead to entirely unforeseen side effect. This is simply due to the fact that within the living organisms, there are potentially many more contexts than have been taken into account during *in vitro* testing. From a modeling point of view, the problem is that it is not known beforehand which contexts of an organisms may or may not be relevant and under which conditions. While there are laboratory systems to test drugs, these are not suitable for the purpose of understanding the effects of the drug *in vivo*.

- Climate change is a prime example of a modeling problem where the consideration of contexts is essential given the purpose of the exercise (i.e., prediction of how the climate is going to change in response to a given amount of CO₂ emission). As an example one can think of the role of the world's oceans. They are not only important regulators of the air temperature, they are also carbon sinks, reservoirs of humidity, they host entire climate-relevant eco-systems, and so on. Each of these roles is a different contexts, and in practical climate modeling tasks it is unclear which ones are important and how they should be represented. The main problem is not primarily the size of the climate models, but to know which effects to include and which ones to leave out.
- Contextuality is also apparent in economics modeling. One aspect of contextuality is that agents in economic systems have multiple roles, say, customers, stock owners, employees, consumers, but they are also voters (in democracies) and polluters of the environment. These multiple aspects of the entities in the system can have grave consequences when trying to predict or explain the behavior of economies.

Case study: evolution

In this section, we will briefly describe a few examples to indicate how complexity manifests itself in practice. All the examples in this section relate to evolution, but complexity will affect the various models in very different ways.

Natural evolution

There is a whole branch of theoretical biology trying to explain the observed traits of behaviors in organisms by reference to their evolutionary origin. An example is the problem of the origin of social behavior. This is often

formulated in terms of *dilemmas*. A behavior that benefits the group may be detrimental to the individual. Naively, one may assume the individual displaying such group beneficial behavior is selected against, with overall detrimental effects on the group. Siderophores in bacteria are a case in point. Bacteria secrete siderophores (Buckling et al. 2007; Griffin et al. 2004) into the environment to liberate iron from the host. Collectively, the bacteria can achieve high enough a concentration of siderophores to liberate adequate amounts of iron. There is a dilemma in this scenario in that synthesizing these siderophores comes at a metabolic cost. The benefits are shared amongst all bacteria. This “shared resource” creates an incentive for each cell to “cheat.” Any cell that stops producing siderophores would continue to benefit from the iron liberated, essentially at the expense of other cells. In the long run, cells that do not produce siderophores will grow faster and have more offspring (because they can divide faster), i.e., they are fitter in an evolutionary sense.

Social dilemmas of this sort are often studied using formal models. In these models, the evolving entities are typically abstracted to the point where they are represented as pure strategies; the concept of fitness is formalized as a pay-off that is allocated to strategies depending on the entire population of strategies and depending also on the environment. The aim of these formal models is, usually, to understand under which conditions the temptation to cheat can be overcome, i.e., under which conditions group-beneficial strategies can evolve. Considerable research effort has been invested into investigating this and similar scenarios to the effect that we now understand the general conditions that support the evolution of social behavior. At least one reason for the success of this research program is that the models ignore nearly all properties of the systems, except the (presumed) relative benefit of strategies for the bearer.

The question we want to ask is whether or not this type of group selection problems are complex or not. To begin with, one may think that they are. There is no real-world laboratory system which can be used to test the theoretical models. There may be instances of group selection in real biology, but it is difficult to corroborate specific models experimentally. After all, there may be many factors contributing to the evolution of siderophores, and the particular host environment may play an important role. This suggests that the problem is complex (which would chime well with the more general impression that evolution is a complex problem).

On the other hand, the debate on group selection is often made with reference to itself, rather than with reference to actual systems. In the mainstream literature one may find weak references to biology; however, fundamentally, the debate on group selection is a debate on theory. Based on a

relatively small number of theoretical scenarios [such as the prisoners’ dilemma (Fletcher and Zwick 2007; Worden and Levin 2007)] theoretical models are developed to investigate the emergence of cooperation. Typical questions are how the evolution of cooperation depends on the spatial structure of the evolving entities (Traulsen and Nowak 2006; Santos et al. 2006) or what precisely the mechanisms are that allow social dilemmas to be overcome [i.e., group selection versus kin selection versus inclusive fitness (Buckling et al. 2007; West et al. 2007; Wild et al. 2009)]. To what extent these conditions are actually realized in real systems is secondary. The purpose of these models is not (at least not directly) to predict or explain any *particular* real world phenomenon, but rather to settle a theoretical question. Contextuality and radical openness are not a problem because the models are designed to avoid it. Hence, there is no complexity in this research question, at least not in the way it is commonly approached.

It is instructive to contrast this with a different type of problem in evolution: the evolution of bipedalism in humans, i.e., the question as to what our walking on two legs (rather than on four) is an adaptation to. One would expect that there was at some point in the evolution of humans a set of circumstances that made it favorable to walk upright. What could those circumstances be? There is no shortage of possible conjectures. For example, one could speculate that bipedalism evolved as humans moved out into the savannah which is dominated by high grass; standing upright enables one to see predators coming from far away, which gives increased fitness (via a reduced chance of being eaten). Upright walkers had more offspring, hence bipedalism evolved. This is one possible narrative (or at least a caricature of one). The problem is there are many others that are at least as plausible.

Coming up with apparently plausible stories is easy. The problem is to decide between competing explanations or even to get a rigorous understanding of the plausibility of any particular scenario in the first place. At the root of this difficulty, one may conjecture, is the lack of a laboratory system to test the plausibility of any account of the evolutionary origin of a feature. This was also the case in the group selection scenario. The key difference between these cases is 2-fold. For one, the question about bipedalism is about a specific biological system, and much less is it a theoretical question that can be considered independently of any empirical context. Moreover, bipedalism cannot be regarded purely as a strategy that is separated from other aspects of the organisms. The transition from walking on four legs to walking on two legs requires rather significant anatomical changes that may impact on a number of other aspects of the life strategy of the organism. For instance, in order to walk on two legs, the hips need to be reconstructed. Apart from being a necessary part in the basic

stability of the organisms, hips also constrain the birth canal in females, which may affect reproductive strategies. Bipedalism frees up front-legs for other tasks, which may impact on foraging/hunting strategies and also social interactions. Bipedalism impacts on the heat economy of the organism in that it changes the area that is directly exposed to solar radiation. This could explain the loss of fur in humans, which in turn enables novel ways of heat regulations (i.e., sweating) that are not effective in animals with fur and there may be further ranging impacts on many other aspects of the species (i.e., sexual behaviors, social structures, and so on).

This example illustrates how contextuality and radical openness together can create uncertainty about what needs to be included in a model and what can be left out. The radical openness manifest itself through the uncertainty about what is and what is not important. Does the effect on the heat economy need to be modeled, or can it be left out? Contextuality enters the modeling problem through the various roles that elements in the models play. In our (somewhat oversimplified) example above, it all starts with the hip that has at least two roles; from there it spirals out, bipedalism affecting phenomena that are seemingly unrelated to upright walking. Maybe bipedalism and its evolution can be modeled within a single context, but we cannot be sure. We have no way to tell. The various contexts of a hip couple, at least potentially, the mode of locomotion to the mode of procreation to the mode of regulating the organism's temperature and maybe to many other things that we have not considered; yet maybe it does not. The problem of modeling the evolution of bipedalism is that there is no laboratory system to help the modeler to draw the boundaries and define a particular system. We do not know which aspects are important and which ones are not. We therefore cannot be certain about the quality of the models. As a result, many explanations are proposed, and one is as good or bad as another one. This is complexity.

It is instructive to contrast the specific problem of how bipedalism evolved in humans with the more general problem of how bipedalism evolved. An example is a model by Hase and Yamazak (2007) who use evolutionary algorithms to evolve the ability to walk upright in robots. Their model or set-up is successful because its purpose is very narrowly focused on evolving a control system for robots, and it is not about understanding how bipedalism evolved in humans. Superficially, their research question seems closely related to the question of the evolution of bipedalism in humans, but fundamentally it is not. The purpose of one is the reconstruction of the evolutionary origins of bipedalism in a real system, which leads to problems of radical openness and contextuality. By shifting the attention away from the evolutionary origin of a feature of a real species to problems of bipedalism as a control-

system a change of focus happened that removed all the complexity. The robots can be confined into a laboratory environment. Hase and Yamazak effectively draw the boundaries and create a system by defining a specific problem which is tractable and most of all whose solution is testable. The cost is that their results are no longer relevant for how upright walking evolved in humans. This stresses again: It is the purpose that is complex, not the "system" itself.

Artificial evolution: tierra

In our second example, we will now consider the inverse problem of evolution. Rather than attempting to construct an evolutionary narrative to explain a specific observed trait, there have been many attempts to harness the power of computers to re-create evolution *in silico* (Adami 1998; Langton 1989). One of the earliest attempts to do this is Ray's Tierra system (Ray 1996). It consists of self-replicating computer programs that compete with one another for access to process time within a (simulated) processor. Each individual program is a number of instructions written in a purpose designed machine language. The environment in which these digital "creatures" live is an addressable memory space. In this sense, the Tierra organisms live in a spatial world and they are spatially extended in their environment. Upon replication, random errors may be introduced into the offspring with a low probability which simulates mutation events. When seeded with a hand-programmed ancestral replicator then an evolutionary arms race kicks off in Tierra where more efficient replicators replace their hand-programmed ancestor at first, and are then eventually themselves replaced by even more efficient individuals. The evolution in Tierra is driven by limited access to CPU time. Hence, there is an adaptive pressure towards fast self-replication.

Quite deservedly, Tierra has initially been received with enthusiasm by the community. One of the reasons for this excitement is that Tierra creatures evolved unexpected strategies to increase their fitness. One such strategy is parasitism, where replicators use the code of other replicators. Parasites then triggered an evolutionary arms race in the course of which one observes the evolution of immunity to parasites and hyper-parasites amongst other strategies. It is notable that all these strategies rely crucially on each individual Tierra organism being spatially extended, which allows other organisms to encroach on each other.

Undoubtedly, the evolution of the digital creatures in Tierra is fascinating. However, it is also limited. Tierra does not produce new variants, novel niches, and strategies indefinitely. In every run, after not too long, evolution reaches a creative plateau. From then on no more interesting variants are created. Mutation events continue to

introduce genotypic variety; yet after the evolutionary plateau has been reached these new variants only enjoy a brief spell of life, before they die out again. This leveling off of evolutionary activity in Tierra is not a consequence of limitations of genotype or phenotype space. The programming language in which Tierra's individuals are written, so to speak the code of its genome, is Turing complete and can be used to express any conceivable algorithm. Yet, this practically unlimited space of possibilities is never explored during the evolution of the system.

This observation of limited evolution is not specific to Tierra, but generalizes to all other computer-simulation of evolutionary systems. Many variants of and improvements to Tierra have been proposed since [see, for example, (Adami 1998)]. Yet, no matter what the sophistication of the system, they all share the basic effect of evolution that is leveling off eventually (Bedau et al. 1997). One may speculate about the reason for this limitation. One explanation is the lack of complexity or the lack of both radical openness and a rich contextuality within all these simulation models. This needs to be understood in the following sense: As far as modeling of real world systems is concerned, complexity depends on the particular purpose of the modeling exercise. When we say that Tierra is not complex then we mean that there is no model of Tierra (i.e., no model of the model) that suffers from limitations due to radical openness or contextuality.

Regarding the radical openness, this is nearly trivial. Computer programs are their own largest models and themselves fully specified by their program code. It is slightly more difficult to argue that computer programs are of low contextuality. In essence this means that, if one modeled artificial computer worlds (i.e., if one attempts to model the computer model), then there would be only very few orthogonal abstractions. This is for a good reason: Every additional context significantly complicates the effort of programming the model and increases the likelihood of introducing bugs and errors. It is therefore instinctively avoided both by modelers (who seek simplicity by nature) and computer programmers (who do not wish to over-complicate their task).

Arguably, in Tierra, there are only two contexts. One context is to model Tierra as a population of self-replicating entities. Each individual is assigned a random number that determines the rate with which it copies itself. Offspring inherit this number, but mutations may slightly adjust the replication rate, and so on. This model would miss out many of the interesting facets of the Tierra world, but it would represent others quite well.

The other context that can be modeled is the spatial extendedness (in the sense that each Tierra organism occupies a certain area of the memory space of the

simulated Tierra-world.⁴) In order to be able to model the parasitism as it occurs in Tierra, this spatial aspect must be taken into account, but for other aspects of the system it is less relevant. For each of those contexts, there could be a multitude of models with different degrees of details. Ultimately, these two contexts exhaust all there is in Tierra. If regarded as a system *an sich* in its own right, Tierra is therefore extremely impoverished as far as contextuality is concerned.

We conjecture that it is this lack of contextuality that severely limits evolution, in that it limits the possible side-effects and repercussions of new mutations, how novel strategies can spread across the system and create new niches. Once a Tierra organism replicates at the maximally possible rate, then there is no room left for improvement. There is no niche creation. The optimization of the replication aspect does not interfere with anything else in the simulated organism. In the real world, evolutionary adaptations may spread far beyond their immediate origins, to unexpected places.

Discussion and conclusion

In this article, we assumed that complexity is the difficulty to generate models. We identified two main causes for complexity in this sense, namely, radical openness and contextuality. Both of these ideas entail that there is no such thing as a system, at least not in the sense that the world is partitioned into well separated entities that are just waiting for a modeller to describe them. Elements of radical openness and contextuality have been in various forms described before in different contexts. It is worthwhile to conclude this article by pointing to some of these predecessor ideas, and to briefly indicate how they relate to the present framework.

The notion of radical openness has elements of hierarchy theories that have been proposed over the years. Two of the best known ones are Salthe's "Evolving Hierarchical Systems" (Salthe 1985) and Ahl and Allen's "Hierarchy Theory" (Ahl and Allen 1996). The idea of hierarchy theories is that nature is organized into different strata that are in a non-symmetric relationship to one another. So, one stratum could be that of atoms, another one that of cells, yet another one is society. Clearly societies contain organisms that contain cells, and cells consist of atoms, hence there is a hierarchy of levels. Understanding levels are important, both philosophically, but also from a

⁴ This is not to be confused with the memory space the simulated objects take in the physical computer on which Tierra is implemented. This memory space, while practically of importance is of no consequence for the behavior of the model.

pragmatic scientific point of view. Philosophically, one incarnation of hierarchies is the reductionism debate. Scientifically hierarchies make themselves felt, for example, in the difficulty to simulate multiple scales concurrently.

While we do acknowledge that studying levels are essential, in this article, we have quite consciously avoided introducing scale and hierarchical levels into the discussion. This is not because we think it is irrelevant, and not because we think that there is nothing left to be said about this topic. The reason we avoided talking about hierarchies is that the concept is so difficult and invites a discussion that can easily sidetrack from the main point of this contribution. Once we talk about levels, then we also have to say, what these levels are, where they begin, and where they end. Is the mitochondrion at the same hierarchical level as the cell, or is it at the level of the cell wall or a rather at that of proteins? The more one attempts to drill down questions of this sort the more one will need to refine the notational toolkit to be able to say anything useful about complexity and modeling.

While talking about hierarchy in natural systems is difficult, introducing the idea is not even necessary for our purpose. In this article, we have been concerned with the question of what makes modeling difficult, rather than asking about the structure of natural phenomena. We also acknowledged that modeling is an inherently pragmatic activity. The structure of a model depends primarily on the aims and objectives of the modeller, rather than on some assumed structure of the world. Hence, how a mitochondrion relates to the cell as a whole and how it relates to proteins depends on the aims of the modeller. It is not an inherent property of the world.

The notion of hierarchy as we intuitively understand it is encapsulated in the concepts of radical openness and contextuality, in that scale is an important reason why models have to be truncated and artificially closed. At the same time, contextuality is more than hierarchy. Hierarchical organization is only one of the factors that force the truncation of models. Even at the same scale there may be couplings between phenomena that will be left out by modelers (or not, depending on the question at hand). Conceptualising truncations in terms of hierarchical levels are not always insightful. As an example, think of a hypothetical modeling problem of how the economy of country *X* couples to climate change and how a changed climate couples back to the economy. For the modeller, it is irrelevant whether the economy and the climate are the same hierarchical level. What matters is how they are coupled together and how the semantic understanding of a connection between these phenomena can be represented in a formal model. Hence, while hierarchy theory provides some of the ideas that are set forth in this contribution, for the purpose understanding modeling the question of hierarchies is secondary.

Another apparently related issue is that of meaning or semantics and its role in biological systems in the way it is often studied in biosemiotics [see, for example, (Hoffmeyer 2009)]. Again, we would not deny the usefulness of understanding the meaning of meaning in biology or any other aspect of science. However, the role of “semantics” as in semantic models is a completely different one. Semantic models are about the interpretation of a formal model in the head of a scientist. It is the idea that in order for models to become part of the scientific corpus they must add to scientific understanding, where understanding is intrinsically semantic. It is crucial to stress that accepting the distinction between semantic and syntactic models does not entail any assumptions about the contents of semantic models. These semantic models could be about the ideas of biosemiotics, but they could equally be the semantic models of an ultra-reductionistic scientist who denies the significance of meaning altogether.

Acknowledgments The research for this article was partially funded by the Norwegian Research Council. The author gratefully acknowledges the hospitality of the Senter for Vitskapsteori at the University of Bergen (Norway) where parts of this manuscript were conceived and drafted.

References

- Adami C (1998) Introduction to artificial life. Springer, New York
- Ahl V, Allen T (1996) Hierarchy theory: a vision, vocabulary and epistemology. Columbia University Press, New York
- Bak P (1997) How nature works. Oxford University Press, Oxford
- Bedau M, Smith R (1999) Is echo a complex adaptive system?. *Evol Comput* 7(1):45–68
- Bedau M, Snyder E., Brown T., Packard N. (1997) A comparison of evolutionary activity and artificial evolving systems in the biosphere. In: Husbands P, Harvey I (eds) In: Proceedings of the fourth European conference on artificial life.
- Bennett C (1986) On the nature and origin of complexity in discrete, homogeneous, locally-interacting systems. *Found Physics* 16(6):585–592
- Buckling A, Harrison F, Vos M, Brockhurst M, Gardner A, West S, Griffin A (2007) Siderophore-mediated cooperation and virulence in *Pseudomonas aeruginosa*. *FEMS Microbiol Ecol* 62(2):135–141
- Chu D, Ho W (2006) A category theoretical argument against the possibility of artificial life. *Artif Life* 12(4):117–135
- Chu D, Ho W (2007) Computational realizations of living systems. *Artif Life* 13(4):369–381
- Chu D, Ho W (2007) The localization hypothesis and machines. *Artif Life* 13(3):299–302
- Chu D, Strand R, Fjelland R (2003) Theories of complexity. *Complexity* 8(3):19–30
- Chu D, Lee H, Lenaerts T (2005) Evolution of DNA uptake signal sequences. *Artif Life* 11(3):317–338
- Chu D, Zabet N, Mitavskiy B (2009) Models of transcription factor binding: sensitivity of activation functions to model assumptions. *J Theor Biol* 257(3):419–429
- Edmonds B (1999) Syntactic measures of complexity. Ph.D. thesis, University of Manchester. <http://www.cpm.mmu.ac.uk/~bruce/thesis>

- Eigen M, Schuster P (1979) *The hypercycle: a principle of natural self-organization*. Springer, Berlin, Heidelberg and New York
- Stanley HS et.al (1996a) Scaling and universality in animate and inanimate systems. *Physica A* 231: 20–48
- Stanley HS et.al (1996b) Scaling and universality in living systems. *Fractals* 4: 427–451
- Fletcher J, Zwick M (2007) The evolution of altruism: game theory in multilevel selection and inclusive fitness. *J Theor Biol* 245(1): 26–36
- Griffin A, West S, Buckling A (2004) Cooperation and competition in pathogenic bacteria. *Nature* 430(7003):1024–1027
- Gross D, Lenaerts T (2003) Towards a definition of dynamical hierarchies. In: Bilotta E, Smith DGT, Lenaerts T, Bullock S, Lund H, Bird J, Watson R, Pantano P, Pagliarini L, Abbass H, Standish R, Bedau M (eds) *Workshop proceedings of the eighth international conference on artificial life*. University of New South Wales Press, pp. 45–55
- Gross D, Strand R (2000) Can agent-based models assist decisions on large-scale practical problems? A philosophical analysis. *Complexity* 5(5):26–33
- Hase K, Yamazaki N (2007) Computational evolution of human bipedal walking by a neuro-musculo-skeletal model. *Artif Life Robot* 3(3):133–138
- Hoffmeyer J (2009) *Biosemitotics: an examination into the signs of life and the life of signs*. University of Chicago Press, Chicago
- Holland J (1995) *Hidden order*. Addison-Wesley, Reading
- Holland J (1998) *Emergence*. Oxford University Press, Oxford
- Langton C (ed.) (1989) *Artificial life—an overview*. MIT Press, Cambridge
- Letelier J, Marin G, Mpodozis J (2004) Autopoietic and (M,R) systems. *J Theor Biol* 222:261–272
- Letelier J, Soto-Andrade J, Abarzua FG, Cornish-Bowden A, Cardenas M (2006) Organizational invariance and metabolic closure: analysis in terms of (M,R) systems. *J Theor Biol* 238: 949–961
- McMullin B, Varela F (1997) Rediscovering computational autopoiesis'. In: Husbands P, Harvey I (eds) *Proceedings of the fourth European conference on artificial life*. MIT Press, Cambridge, <http://www.eeng.dcu.ie/~alife/bmcm-ecal97/>
- McShea D (1996) Metazoan complexity and evolution: is there a trend?. *Evolution* 50:477–492
- Mitchell M (1997) *An introduction to genetic algorithms*. A Bradford Book. 3rd edn. MIT Press, Cambridge, pp 378–387
- Noble D (2009) *The music of life*. Oxford University Press, Oxford
- Pickering A (1995) *The mangle of practice*. The University of Chicago Press, Chicago
- Ray T (1996) *An approach to the syntheses of life*, Oxford Readings in Philosophy. Oxford University Press, Oxford, pp 111–145
- Rescher N (1998) *Complexity—a philosophical overview*. Transaction, New Brunswick, p 219
- Rosen R (1991) *Life Itself*. Columbia University Press, New York
- Rosen R (1999) *Essays on life itself*. Columbia University Press, New York
- Salthe S (1985) *Evolving hierarchical systems*. Columbia University Press, New York
- Santos FC, Pacheco JM, Lenaerts T (2006) Evolutionary dynamics of social dilemmas in structured heterogeneous populations. *Proc Natl Acad Sci USA* 103(9): 3490–3494
- Slutsky M, Kardar M, Mirny L (2004) Diffusion in correlated random potentials, with applications to DNA. *Phys Rev E* 69(6):061903
- Sneppen K, Zocchi G (2005) *Physics in molecular biology*. Cambridge University Press, Cambridge
- Traulsen A, Nowak M (2006) Evolution of cooperation by multilevel selection. *Proc Natl Acad Sci USA* 103(29): 10952–10955
- West S, Diggle S, Buckling A, Gardner A, Griffin A (2007) The Social lives of microbes. *Annu Rev Ecol Evol System* 38:53–77
- Wild G, Gardner A, SA West (2009) Adaptation and the evolution of parasite virulence in a connected world. *Nature* 459(7249): 983–986
- Wolkenhauer O (2002) Systems biology: the reincarnation of systems theory applied to biology?. *Brief Bioinform* 2(3):258–270
- Wolkenhauer O, Hofmeyr J (2007) An abstract cell model that describes the self-organization of cell function in living systems. *J Theor Biol* 246(3):461–476
- Worden L, Levin S (2007) Evolutionary escape from the prisoner's dilemma. *J Theor Biol* 245(3):411–422
- Wunderlich Z, Mirny L (2008) Spatial effects on the speed and reliability of protein-DNA search. *Nucleic Acids Res* 36(11): 3570–3578
- Wynne B (2005) Reflexing complexity: post-genomic knowledge and reductionist returns in public science. *Theory CultureSoc* 22(5):133–138