

Kent Academic Repository

Full text document (pdf)

Citation for published version

Chu, Dominique (2008) Criteria For Conceptual And Operational Notions of Complexity. *Artificial Life*, 14 (3). pp. 313-323. ISSN 1064-5462.

DOI

Link to record in KAR

<https://kar.kent.ac.uk/23992/>

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>

Criteria for Conceptual and Operational Notions of Complexity

Dominique Chu*
University of Kent

Keywords

Complexity, postnormal science, information, physical complexity

Abstract While complex systems have been studied now for more than two decades, there still is no agreement on what complexity actually is. This lack of a definition might be a problem when asking questions about the evolution of complexity. In this article criteria against which candidate measures of complexity can be assessed are discussed. The main conclusion of this article is that because of the absence of a basic consensus on what complexity is, there is no criterion that can be used to decide whether or not a proposed measure actually measures complexity. The main recommendation is to abandon complexity as a formal notion; instead, research into the evolution of complexity should use well-understood proxy notions (as is sometimes done in the literature). For the time being “complexity” should remain an informal notion. Research into evolutionary trends of these proxy notions might eventually lead to an emergent community consensus on what complexity is.

I Introduction

While most will have a strong intuition that the complexity of organisms grows over evolutionary time, as yet there is no strong empirical support for this proposition (although there is a significant amount of literature attempting to tackle various aspects of this topic; examples are [23, 24 29, 4, 31, 22, 12, 26]). At present, the main obstacle to progress on this topic is the lack of a practical and generally agreed-on definition of complexity. There is no agreement in the community of researchers on what complexity should mean and what its essential properties are. As will be argued below, this disagreement is not a disagreement about how to fill some gaps in an already erect building of complexity, but really about how to build the foundation. The various concepts of complexity that have been proposed range from complexity as an essentially information theoretic measure, via complexity as a property of models, to complexity as emerging from the system-environment interaction.

This article has two aims. Firstly, it will provide a brief survey of existing notions of complexity in Section 3. This will not be an exhaustive review of complexity, but merely aim to illustrate various types of complexity as they have been proposed in the literature. It will become clear that not only are the notions of complexity presented there are not coextensive, but in fact some of them exclude one another. This observation of diversity then leads to the question of how one can or should choose between all these different notions of complexity. In order to approach this, the distinction between a conceptual and an operational notion will be introduced in Section 2. This distinction will

* Computing Laboratory, University of Kent, CT2 7NF Canterbury, UK. E-mail: D.F.Chu@kent.ac.uk

be useful in Section 4, where two properties of a useful notion of complexity are discussed: One property is that it needs a well-defined experimental protocol (this will be described in more detail below). The second property is conceptual soundness. While the first one is problematic in many (if not most) of the currently proposed notions of complexity, the conclusion in Section 5 is that it is this second property that is the real difficulty at present; and it is unlikely that it can be resolved anytime soon. In the light of this, the article will end with a perhaps unsatisfactory (to some) recommendation: In the absence of a consensus on complexity, the most efficient way forward is to retain complexity as an informal, intuitive concept, while abandoning (at least for the time being) attempts to condense the contradicting ansatzes to a least common denominator. In the context of studying the evolution of complexity this strategy is already, at least partially, followed (see for example [23, 24, 29]). By giving up the ambition to find a universally accepted concept of complexity now, research efforts can be diverted into empirical studies of the evolution of “complexity-related” properties of systems, possibly leading to an emerging understanding of complexity in the future.

2 On the Nature of Scientific Notions

Before entering the main discussion on complexity some preliminaries are necessary. The basic distinction between the *operational* and the *conceptual notion* will feature heavily in this article, and it is therefore worthwhile to introduce it here. A word of caution, though. While this distinction will turn out to be very useful in the present case of complexity, no claim is made of its general applicability to all scientific notions, nor that it would withstand a thorough scrutiny by philosophers of science.

The *operational notion* is a part of a scientific notion that allows one to derive a procedure for measuring a relevant quantity. So, for example, $S = Q/T$ is an operational notion if we know how to measure Q and T . Any scientific notion can have several operational notions (or none). (In the case of several notions one needs to make sure that these are consistent, i.e., lead to the same measurement results.)

While the operational notion specifies how to measure a scientific quantity, the *conceptual notion* contains the semantic contents of a notion. Conceptual notions are essential for the overall understanding of scientific concepts but are not necessarily formulated as a set of dogmas or theorems. For example, in thermodynamics the concept of entropy will be associated with the fundamental laws of thermodynamics, the concept of order and disorder, equilibrium states, and so on. All these aspects form part of the conceptual notion. Note that the operational notion will often contribute to the conceptual notion. For example, knowing that $S = Q/T$ gives some idea about what S means, if it is known what Q and T mean. In well-established scientific notions the conceptual notion will normally be much deeper than what can be purely derived from the operational notion.

An operational notion by itself is (mostly) not very helpful in generating understanding about a system. Given one or several operational notions, it is possible to form an arbitrary number of new ones. However, this will normally not lead to new insights. For example, assuming T and Q are operational notions, then so is $v = Q^2 T^{1/2}$. While it is possible to measure v , there is little conceptual depth in v . It is unlikely that measuring v will uncover anything new or useful about the system. In most cases, a scientific notion without conceptual depth is of little use.¹ On the other hand, a conceptual notion might still be useful even though there is no way to measure a quantity associated with it.

3 Notions of Complexity

The preliminaries having been dealt with, this section will survey some notions of complexity, as they have been proposed in the literature. The first and foremost aim of the following review is to sketch

¹ There are conceivable exceptions to this. One might, for example, be able to correlate the occurrence of certain practically relevant events with quantitative values of an otherwise poorly understood observable. Such complications will be ignored here, because they are not relevant for the present discussion.

(with a broad brush) various conceptual approaches to complexity; it will not provide a detailed review of all complexity measures that have been proposed so far. (The reader interested in such a review is referred to the existing literature on this topic [12, 26, 30].)

3.1 Information-Based Measures

There is a whole family of complexity measures that are based on information theoretical measures (henceforth collectively referred to as *I complexity*). The reader interested in the details is encouraged to consult the literature [7, 26, 12, 8]. There is a great variety of I-complexity measures, differing from one another in various aspects. Common to all I-complexity notions is that they associate complexity with information (in one way or another). Measuring I complexity usually requires a string of some sort, such as a computer program, nucleotide sequences, or strings of amino acid or the like. Characteristic for I complexity is that, as long as such a string can be found, every system is complex to some degree. So, unlike other notions of complexity, I-complexity notions do not imply the existence of a class of “simple systems” that are in some qualitative way different to “complex systems.”

Many I-complexity measures have been proposed in the literature. For example, Bar-Yam [7, Chap. 1] equates complexity with the information contents of a message: A message (or string) is more complex if it is more unexpected. A string of numbers is maximally complex if it is random. A similar measure is *algorithmic information contents* [19]. The idea here is that the complexity of a system is the length of the shortest algorithm that generates it. Again, strings of random numbers would be of high complexity, whereas the complexity of time series of chaotic systems would be the size of the set of equations that describe the system. Another example of an I-complexity measure has been proposed by Adami and coworkers [4] to investigate the growth of I complexity in their artificial life system Avida (they call this measure *physical complexity*). Physical complexity is defined as the information organisms have about their environment. By “information” the authors really mean genes coding for phenotypical traits that help survival by (implicitly) providing information about the conditions of their niche. In practice the authors propose to determine physical complexity by measuring the degree to which genes have converged across the population; according to this measure the complexity is maximal if all genomes in the population are equal.

In summary: I-complexity notions measure properties of strings (or similar structures); they do not imply a partition of systems into a complex and a simple class. I-complexity measures are particularly easy to apply to formal objects such as computer programs, texts, sequences, or graphs.

3.2 Rosen Complexity

A different notion of complexity is due to Robert Rosen [27, 28, 10]. A system is Rosen complex (or R-complex) if there is no finite formalism that can model all of its aspects. Rosen showed that there is at least one class of systems that is R-complex, namely the class of systems that are “closed with respect to efficient causation.” It would take us beyond the scope of this article to discuss in detail what this means. Suffice to say that closure *sensu* Rosen is closely related to the better-known closure property of autopoietic systems [21, 32, 20]. Roughly, closure means that the processes going on in living systems reconstitute one another in a circular fashion, that is, process *a* maintains process *b* maintains process *c* ... maintains process *a*. According to Rosen, all (real) living systems are closed with respect to efficient causation and are thus R-complex. On the other hand he claims that computer-generated lifelike forms cannot possess this closure property and will not be R-complex.

The notion of R complexity is not entirely uncontroversial. There are reviewers who argue that Rosen’s arguments (that simulated systems cannot be R-complex) are flawed [10, 33]. A detailed discussion of this would go beyond the scope of this article; suffice to say that the disagreement can only be about which classes of systems are R-complex. The notion of R complexity itself is as legitimate as information-theory based ideas about complexity, but it is also very different to I complexity: R complexity is a dichotomy, that is, it is present or not, whereas (at least conceptually) all systems are I-complex to some degree. Moreover, the classes of systems that are particularly amenable to operational notions of I complexity (namely, computational systems) are by definition

not R-complex. Vice versa, it will become clear below that the operational notion of I-complexity is problematic in the context of R-complex systems such as organisms.

3.3 Complexity and Laboratory Systems

Complex systems have often been (and still are) associated with so-called *complex adaptive systems*, or CASs [9, 17, 18]. The details of what is and what is not a CAS vary somewhat from author to author. Common to most definitions are these features:

- CASs consist of a middle-sized number of autonomous entities.²
- These entities have the ability to adapt to one another and to their environment.
- The population of entities is heterogeneous.
- Interactions between entities are nonlinear.

CASs in this sense characterize a class of simulated and real systems; in the literature (see, for example, [6, 9]) these systems are often associated with complex systems, thus distinguishing them from simple systems as they are studied in *sciences of the simple* such as physics and chemistry. These sciences of the simple are characterized by general theories that unify large sets of phenomena and typically have strong predictive aspects, whereas a CAS (and hence a complexity science) does not have unifying and predictive theories [11] (at least, none have been found so far).

The notion of CAS is very common, but often the connection between CASs and complexity is left unclear. There have been attempts to clarify this connection, for example, in [16, 11] the authors attempt to make the connection between complexity and CASs via the notion of a laboratory system. A laboratory system (LS) is the real-world counterpart of an idealized system in physics. This is best explained by means of an example. Consider the following statement:

All bodies fall at the same speed.

This is what most pupils learn in one of their first lessons on elementary physics. However, as the reader can easily convince herself, it is wrong: Simultaneously dropping this issue of *Artificial Life* and a single sheet of paper will show that the bound journal issue will hit the floor earlier than the loose sheet of paper. The discrepancy between the above statement (prediction) and the experiment is of course explained by the fact that the experimental conditions were not conducted under ideal conditions, that is, in vacuum.

The validity of most statements, predictions and theories in sciences of the simple is restricted to ideal conditions of some sort. Testing them requires that experimental scientists build LSs that approximate these idealized conditions sufficiently well to allow meaningful tests of the theoretical predictions. In real (natural or engineered) systems the idealized laws of physics are practically never strictly valid; as it turns out, however, the insights gained from idealized theories are nonetheless applicable to a wide range of practically relevant circumstances and systems as long as one factors in perturbations resulting from non-ideal circumstances. The range of systems where such a LS-based approach works is of high practical importance and in particular includes many engineered systems. Yet, there are systems where LS-based approaches fail. One way to characterize LS complexity is to say that a system is LS-complex if a LS-based approach to studying it no longer yields useful results (CAS would be an example of such a class).

The intuition behind LS complexity is often expressed by contrasting a jumbo jet and a bacterium: The former is then said to be merely *complicated*, whereas a bacterium is genuinely *complex*. This distinction is based on the insight that the engineering of a jumbo jet, while by no means an easy

² The idea behind this is that for a small number of entities one could use exact methods to solve the system; for a very large number one could treat the system as continuous, hence significantly simplifying the mathematics.

task, is very much based on LS approaches (in the sense that the engineering going into the jet is tractable and modularized).

Conceptually LS complexity is first and foremost a dichotomy (although one could attempt to differentiate LS-complex systems according to their degree of complexity; note that LSs would always have zero complexity). This highlights a crucial difference to I-complexity measures that assign some degree of complexity to all systems, including laboratory systems. Even if one based the operational notion of LS complexity on I complexity, there would still be a conceptual incompatibility between LS complexity and I complexity. Similarly, there is a conceptual incompatibility between LS complexity and R complexity; any computer-simulated CAS is certainly not R-complex (by virtue of its being computer simulated), but would be LS-complex.

3.4 Funtowicz-Ravetz Complexity

The final notion of complexity discussed here is due to Funtowicz and Ravetz [14, 15, 13]. These authors are mainly interested in the role complexity plays in the context of public decision making, but, as will be indicated below, their notion of complexity is still potentially relevant in other contexts as well. Funtowicz and Ravetz distinguish between ordinary complexity (this corresponds roughly to LS complexity) and what they call “emergent³ complexity.” Emergent complexity according to Funtowicz and Ravetz (or FR complexity) is based on the observation that scientific activity and its products do not exist in isolation from the rest of the world but affect and interact with it in various ways. For example, new technologies can have great impact on society, the environment, and science itself; they are not the neutral products of scientific activities, but themselves catalysts of (possibly large) changes. This introduces complexities into the scientific process that are not visible if one limits one’s attention to the system under consideration itself.

Take as an example the discovery of nuclear fission. This discovery took place within the realm of LS-based science. The results of this research (such as the nuclear bomb) propagate out far beyond the laboratory of the scientists, having impacts on a global scale, possibly threatening the existence of mankind. Other fields of science such as genetics, food science, and economics might give rise to similar FR complexities. The reader interested in a more detailed account of FR complexity is referred to the original literature [14]. For the present purpose it is sufficient to point out that FR complexity is not a property that is inherent in a particular system, but emerges from the way science and its products are used and interact with the rest of the world. This makes the conceptual notion of FR complexity fundamentally different to the other notions of complexity that have been described above. FR complexity is also mainly a conceptual notion and does not offer a measure of complexity, although Funtowicz and Ravetz do describe a hierarchy of various forms of complexity [14].

Funtowicz and Ravetz are mainly interested in complexity in the context of public decision making; their notion of complexity might at first seem to be irrelevant to the other notions of complexity described here. This first impression is unjustified, however. The essence of FR complexity is that complexity is not inherent in the system but emerges from the interaction of the system with its ambience. Once you take away the context from a complex system, all that remains is complicatedness. Measuring the intrinsic complexity of a system is, if one subscribes to the concept of FR-complexity, rather meaningless. This aspect is perhaps not too dissimilar from the motivation behind Adami’s concept of physical complexity (which equally sees complexity as a relation between a system and its environment). Particularly in the context of evolution (a process that is very much about the individual-environment interaction) the idea of including the environment in a definition of complexity should at least be worth some consideration.

3.5 Summary

This brief survey indicates that there is a great variety of ideas about complexity. The difference between these ideas is not a matter of detail but a deep conceptual incompatibility. That these

³ The word “emergent” here should not be associated with the usual meaning of the word in the context of artificial life, but rather be taken as a name.

notions can coexist side by side reveals that there is a fundamental uncertainty about what complexity is. The question now is: Given all these notions of complexity (and many others that might exist), how can one decide which of those is the correct one, or at least the proper one to use?

4 Criteria

This section will define two criteria that can be used to evaluate the usefulness of a scientific notion. The first criterion is that the experimental protocol of the operational notion is well defined. As will be argued below, this criterion is not met by most complexity measures, suggesting that, at the very least, more work is needed. The second, arguably more important criterion (discussed in Section 4.4) is conceptual soundness: How can one determine whether or not a specific notion of complexity actually captures complexity?

4.1 Experimental Protocol

For any scientific notion it is essential that the *experimental* protocol of an operational notion be well defined. This means the following: The experimental protocol is well defined if there is a set of criteria that allows a competent observer to judge whether or not a specific measurement procedure has been applied correctly. Think of measuring the temperature of a system: For a correct measurement the apparatus and the system need to be brought into thermodynamical equilibrium. Briefly touching the system with the thermometer will normally not be a valid operational procedure and would lead to meaningless measurements. Another condition is that the system itself is of macroscopic size; it is meaningless to measure the temperature of (for example) an electron. One must also exclude external perturbations that affect the measurement result (a common mistake is to measure air temperature in bright sunshine).

If the experimental protocol is well defined, then a competent observer should be able to decide from observations whether or not a specific measurement was valid or not; two independent observers of the same measurement should agree in their judgment. At the very least there must be possibilities to perform control experiments if the validity of measurements is unclear.

4.2 Problems of Measures of Complexity

The experimental protocol is unclear in some proposed measures of complexity. Instead of demonstrating the shortcomings of each of the above-described notions of complexity, the type of argument that is relevant will only be demonstrated using the particular case of I-complexity based measures; similar arguments can be made for each of the other notions of complexity.

Consider the example of algorithmic information contents as complexity. As long as one is in the realm of computer programs, valid experimental protocols to at least find upper bounds for complexity are fairly clear (i.e., search for the smallest possible program that reproduces a system or set of data). The situation changes if one attempts to apply the measure to real organisms: What is the algorithm that generates an organism? One could resolve this by using a description of the organism (rather than the system itself); such a description can be converted into computer readable formats convenient for measuring various forms of I complexity. In practice, though, this approach is problematic because it requires a complete or near-complete description of the system; these are nearly never available in the case of organisms. Another option is to view developmental or generative processes of organisms as executions of algorithms. One might, for example, base an estimate of the algorithmic information contents on the length (or compressed length) of the genome. Unfortunately, to what extent and under which circumstances the genome really can be seen as an algorithm or at least as a reliable indicator for the algorithmic information contents is controversial. In summary: There are ways to determine the algorithmic information contents of an organism (or any other I-complexity measure), but each of those ways will lead to a different answer. None is a dominant contender for the “true” measure that reflects the complexity of the organism; the lack of a valid experimental protocol severely restricts the usefulness of these complexity measures. This

should not be understood as a criticism uniquely applying to I-complexity, but rather as a reminder that any relevant operational notion of complexity needs to be applicable to the entire range of systems that one wishes to include in the investigation (so, in the present case of the evolution of complexity, typically at least to organisms and computational structures).

4.3 A Specific Example

It is useful to illustrate this using a more specific example. Adami's notion of physical complexity was described above; it was originally intended as an improvement over many existing I-complexity measures.⁴ What follows will contain some objections to the usefulness of physical complexity as a *universal* notion of complexity based on the lack of a well-defined experimental protocol. These arguments, however, are not meant to imply that there is no scope for useful applications of physical complexity measures to genomes of real and artificial organisms. Quite to the contrary, physical complexity has already been successfully applied (see, for example, Adami [3]).

Adami illustrates the use of physical complexity in (perhaps) the simplest conceivable evolutionary system: His artificial ecology of evolving computer programs called Avida [1]. The advantage of a computational system such as Avida is that every aspect of the evolving system is known. However, even in this ideal case physical complexity cannot be measured directly. Instead of measuring the information the evolving entities have about their environment, it is necessary to fall back onto an approximate measure (similarity of the DNA sequence in the population, as described above). This does not introduce a big error, but it foreshadows the difficulties of measuring physical complexity in real systems where the difficulties in determining an experimental protocol to measure it become more pronounced.

Adami himself cites several such difficulties, including hitchhiking and sexual recombination (although he also expresses confidence that these problems can be overcome). In a more general context, there are a number of questions regarding how to consistently measure physical complexity in real organisms. Again, it should be stressed that there is no doubt that there are specific applications of physical complexity that can be fruitful; yet if physical complexity is to be a universal measure of complexity, then it must be applicable in a very wide range of circumstances.

For example, one might be interested in comparing the complexity of two organisms—no assumptions made about the relatedness of the species these organisms are a member of. To do this one could determine the complexity of each, that is, measure the information each organism has about its environment; this, as Adami points out, is not practical, so instead what is measured is the degree to which the genomes in the population have converged. There are two obvious problems of this measure.

Firstly, the maximally attainable physical complexity, as proposed by Adami, does depend on the length of the sequence that is used to determine the complexity. This is at the very least a minor problem when dealing with two species whose genome is of different length. There are two possible strategies to avoid this. One could choose two subsequences of the particular genomes, or alternatively one could normalize the measure over the length of the genome. Both approaches are problematic. In the first case, there needs to be a rule to determine how the subsequences should be chosen; the correct choice is crucial because intrapopulation variability might vary across the genome. In the second case it would be difficult to avoid noncoding parts of the genome (with potentially much higher variation) influencing the result; generally species with a higher proportion of such noncoding DNA will tend to come out with lower physical complexity. Hence, at the very least a clear rule is required for how complexity should be measured.

Assuming these questions can somehow be answered, there is at least a second problem. Assume that one wants to compare the complexity of a bacterium with that of an Avida organism. Assume further that hypothetically both attain the maximum complexity value L . What would that tell us? In the first instance it would mean that the populations to which both organisms belong are genetically

⁴ Note that Adami himself seems to find most I-complexity measures unfit. Referring to I-complexity measures discussed in a review by Badii and Politi [5], Adami writes: "Most of them, however, do not appear satisfactory from an intuitive point of view" [2, p. 1086].

uniform, or that they have maximal information about their respective environments. It is not clear, however, to what extent a shared value of L is indicative of a commonality between the bacterium and the Avida organism (does it mean they are equally complex?).

There are a number of other problems along similar lines. In particular it is unclear how one could measure the physical complexity of engineered objects (that do not have a genome), such as for example jumbo jets; how one could measure the complexity of parts of organisms that are not directly coded for by a specific set of genes (what is the complexity of the human eye?), and so on. It cannot be excluded that these difficulties will be overcome by a proper specification of the experimental protocol. Yet so far, such a protocol does not exist. Issues that need to be resolved by such a protocol include:

- How should the population (used to determine physical complexity) be chosen?
- What sample size is necessary to obtain meaningful results?
- How should the portion of the genome used to determine the physical complexity be chosen?
- How can one systematically discount for contingencies such as population bottlenecks during speciation, high rates of neutral mutations, and so on?
- How can one measure the complexity of engineered objects?

This section should not be concluded without stressing once again that the lack of an experimental protocol is by no means unique to I-complexity measures, but extends to all of the above-described notions of complexity. A detailed analysis would distract too much from the main focus of this article and will be left for future publications.

4.4 Conceptual Soundness

A well-defined experimental protocol by itself is not a sufficient criterion for an operational notion of complexity; there are many measures that have a well-defined protocol, but do not measure complexity. Hence, at least one more criterion is required: Conceptual soundness. This means (a) that the operational notion is consistent with the conceptual notion (i.e., the measure of complexity measures what the conceptual notion describes) and (b) that the conceptual notion of complexity captures what complexity “really is.”

The above formulation suggests an essentialist position on complexity, that is, that there is a complexity property out there in the world that needs to be discovered. However, such an essentialism is not to be implied above (though not excluded either). The phrase “...what complexity really is...” is only to be understood as a convenient shorthand. It can mean either that there indeed is an essence to the property of complexity, or alternatively, that there is a community consensus on the concept of complexity.

The idea of conceptual soundness is best illustrated by using the example of the concept of life: Assume one conceptualizes life as autopoiesis (i.e., a system is alive if (and only if) it is autopoietic). This conceptual notion of life might or might not be sound, and there is a clear procedure to test this: (a) Look at a specific example of a living system (autopoietic system). (b) Determine whether it is autopoietic (living). (c) If not, then this particular notion of life is unsound; otherwise, continue looking. Following a similar procedure, many definitions of life have been found unsound because they misclassified particular examples. In the case of life this method can work because there is broad agreement across scientists as to which systems are living and which are not. The question now is: Is there a similar way to decide whether or not particular notions of complexity are sound?

As it turns out, this procedure cannot be carried over to the case of complexity. Firstly, there is no set of known complex systems that could be used to benchmark candidate notions of complexity. Secondly, many notions of complexity are not dichotomies, and more subtle approaches are required.

An alternative possibility is to assess candidate notions of complexity against the intuition (rather than directly against known complex system): Assuming the relevant concept to be tested has an operational notion of complexity, one could apply the measure to a number of systems S_1, S_2, \dots, S_n whose ranking with respect to their complexity is intuitively clear. If the ranking suggested by the measured complexity values grossly contradicts the intuitive ranking, then one would be inclined to reject this measure as unsound. Unfortunately, such a consensus ranking does not exist. There is perhaps a broad agreement on trivial rankings such as the relative complexity of cats and screwdrivers; once the comparisons become more subtle, any basis for a community agreement vanishes. Appealing to intuition as a judge is likely to lead to disagreements between researchers. These cannot be resolved without some additional criterion that decides whose intuition is correct and whose is not. On a more fundamental level, it is doubtful that intuition is a good foundation for any scientific endeavor, for it introduces an element of subjectivity into the scientific process; this is usually not desirable and indeed unnecessary. Many scientific notions do not correspond to any intuitive concept (for example, enthalpy). Finally, intuitively some might confound the complexity of a system with (what was called above) the complicatedness of a system. It might well be that complicatedness and complexity turn out to be intimately linked; yet at present this is not clear, and there is at least a sizable minority opinion holding that they are distinct. Hence, a danger when relying on intuition is that the eventual outcomes of the discourse get prejudiced by some tacit assumptions.

What about other criteria to decide the soundness of notions of complexity? The claim made here is that no such criterion can exist. In order to see this, assume that there is such a criterion. This criterion will then allow us to decide whether or not a specific notion of complexity is conceptually sound, that is, whether or not the conceptual and the operational notion are consistent with one another and with the “true” meaning of complexity. This means that the criterion will (at least implicitly) have to contain this “true” notion of complexity. How can it be determined that the notion of complexity it implies is indeed correct? Yet another criterion would be required, thus leading into an infinite regress.

The problem seems to be severe indeed: We do not know how to measure complexity, and neither do we know what precisely it is that we do not know how to measure. In this situation of double uncertainty, asking about the evolution of complexity is a bit like asking about the evolution of Ξ . Complexity and Ξ have in common that it is unknown how to measure them and what they actually are.

5 Conclusion

Is there any point in continuing to research the evolution of complexity? Perhaps there is, if one ceases to insist on a well-defined notion of complexity. After all, the field of complex systems has been studied now for over 20 years, and it is fair to say that it has been successful in many ways. All this time, complexity researchers have managed to study complex systems without being able to pin down what complexity actually is. In the light of the success of the science of complexity, it is inconceivable to abandon the notion; this would be like abandoning the concept of life in biology, simply because nobody can properly define it.

The mere fact that a concept is useful does not mean that it is necessarily well defined. A possible conclusion from the preceding discussion is that complexity is simply not a scientific notion on a par with other, very well-defined and measurable notions such as temperature or entropy. At least at present, attempts to make the concept sufficiently precise to lift it to the status of such concepts results in trivializing it in its scope in exchange for a stronger operational notion. That a quantity is measurable does not mean it is useful; that an aspect of a complex system can be measured does not mean this aspect is indicative of “complexity.” Over time it might become clear that complexity is itself a complex notion (i.e., there is no single measure that captures it, but new approaches of characterization, potentially involving both quantitative and qualitative aspects, are necessary).

If one leaves complexity as a powerful intuitive notion, however, without a precise definition, does that mean that a research program on the evolution of complexity needs to be abandoned? Hardly!

The fuzziness of the notion of complexity will only hamper this program as long as one insists that there is one true and well-defined notion of complexity. On the other hand, if one accepted that complexity is not (at least at present) a scientific notion but rather an informal experience, a research program into the evolution of complexity would still be fruitful.

So what is a possible solution? Instead of putting research efforts into prolonged (and most likely ultimately fruitless) discussions of what complexity really is, one should attempt to characterize aspects of systems that (a) are complexity-related and (b) are well defined and (c) have a clear operational protocol. The advantage of this approach is that it allows the measurement of well-defined quantities and how they change over evolutionary times without having to commit to a specific notion of complexity. Some researchers have taken precisely this route. Daniel McShea [23, 25, 24] attempts to determine a number of different indicators of complexity instead of postulating just one way to measure it. He measures aspects such as the number of functional parts of a system, the nestedness of a system, the level of interaction between various hierarchical levels, and so on. Other possible complexity-related properties of systems are the number of protein species of an organism and the degree distributions of interaction networks, to mention but two.

All these are properties of systems that are clearly relevant to but certainly not equivalent with complexity. Such an approach might seem unsatisfactory in that it does not immediately put a number to complexity. A simple value describing the complexity of a system would be ideal, of course, in order to rank complex systems and to easily investigate simple trends. Yet, concentrating on complexity-related, but well-defined, measures might in the long run be more fruitful in creating a consensus on what complexity is and how it should be measured. It could also be, however, that such a consensus will never arise, at least not in the sense of convergence to a scalar-valued complexity measure. This article is perhaps best concluded by a quotation from McShea:

Is a human more complex than a trilobite overall? The question seems unanswerable in principle [...] [23, p. 480]

Acknowledgments

Some of the ideas of this article were developed during a visit to the Senter for Vitskapsteori at the University of Bergen, Norway in October 2006. I am grateful to the Senter for their hospitality and to Prof. Roger Strand for discussions that helped to focus my thinking on the subject of complexity. I also thank the anonymous reviewers for constructive and helpful comments that helped to improve the manuscript.

References

1. Adami, C. (1998). *Introduction to artificial life*. New York: Springer-Verlag.
2. Adami, C. (2002). What is complexity? *Bioessays*, 24(12), 1085–1094.
3. Adami, C. (2004). Information theory in molecular biology. *Physics of Life Reviews*, 1, 3.
4. Adami, C., Ofria, C., & Collier, T. C. (2000). Evolution of biological complexity. *Proceedings of the National Academy of Science of the USA*, 97(9), 4463–4468.
5. Badii, R., & Politi, A. (1997). *Complexity: Hierarchical structures and scaling in physics*. Cambridge, UK: Cambridge University Press.
6. Bak, P. (1997). *How nature works*. Oxford University Press.
7. Bar-Yam, Y. (1997). *Dynamics of complex systems*. Reading, MA: Addison-Wesley.
8. Bennett, C. (1986). On the nature and origin of complexity in discrete, homogeneous, locally-interacting systems. *Foundations of Physics*, 16(6), 585–592.
9. Casti, J. (1997). *Would-be worlds*. New York: John Wiley & Sons.
10. Chu, D., & Ho, W. (2006). A category theoretical argument against the possibility of artificial life. *Artificial Life*, 12(4), 117–135.

11. Chu, D., Strand, R., & Fjelland, R. (2003). Theories of complexity. *Complexity*, 8(3), 19–30.
12. Edmonds, B. (1999). *Syntactic Measures of Complexity*. Ph.D. thesis, University of Manchester. Available at <http://www.cpm.mmu.ac.uk/~bruce/thesis>.
13. Funtowicz, S., & Ravetz, J. (1990). *Uncertainty and quality in science for policy*. Dordrecht: Kluwer.
14. Funtowicz, S., & Ravetz, J. (1994). Emergent complex systems. *Futures*, 26(6), 568–582.
15. Funtowicz, S., & Ravetz, J. (1994). Uncertainty, complexity and post-normal science. *Environmental Toxicology and Chemistry*, 13(12), 1881–1885.
16. Gross, D., & Strand, R. (2000). Can agent-based models assist decisions on large-scale practical problems? A philosophical analysis. *Complexity*, 5(5), 26–33.
17. Holland, J. (1995). *Hidden Order*. Reading: Addison-Wesley.
18. Holland, J. (1998). *Emergence*. Oxford, UK: Oxford University Press.
19. Kolmogorov, A. (1965). Three approaches to the quantitative definition of information. *Problemy Peredachi Informatsii*, 1(1), 3–11.
20. Letelier, J., Marin, G., & Mpodozis, J. (2004). Autopoietic and (M,R) systems. *Journal of Theoretical Biology*, 222, 261–272.
21. Maturana, H., & Varela, F. (1980). *Autopoiesis and cognition: The realization of the living*. Boston studies in the philosophy of science. Dordrecht: Reidel.
22. McMullin, B. (2000). John von Neumann and the evolutionary growth of complexity: Looking backwards, looking forwards. *Artificial Life*, 6(4), 347–361.
23. McShea, D. (1996). Metazoan complexity and evolution: Is there a trend? *Evolution*, 50, 477–492.
24. McShea, D. (2002). A complexity drain on cells in the evolution of multicellularity. *Evolution: International Journal of Organic Evolution*, 56(3), 441–452.
25. McShea, D. (2005). The evolution of complexity without natural selection, a possible large-scale trend of the fourth kind. *Paleobiology*, 31(2), 146–156.
26. Rescher, N. (1998). *Complexity—a philosophical overview*. New Brunswick, NJ: Transaction.
27. Rosen, R. (1991). *Life itself*. New York: Columbia University Press.
28. Rosen, R. (1999). *Essays on life itself*. New York: Columbia University Press.
29. Saunders, W., Work, D., & Nikolaeva, S. (1999). Evolution of complexity in paleozoic ammonoid sutures. *Science*, 286(5440), 760–763.
30. Standish, R. (2002). On complexity and emergence. *Complexity International*, 9.
31. Standish, R. K. (1999). Some techniques for the measurement of complexity in Tierra. *Lecture Notes in Computer Science*, 1674, 104–110.
32. Varela, F., Maturana, H., & Uribe, R. (1974). Autopoiesis: The organisation of living systems, its characterization and a model. *Biosystems*, 5, 187–196.
33. Wells, A. (2006). In defense of mechanism. *Ecological Psychology*, 18(1), 39–65.

AUTHOR QUERY

AUTHOR PLEASE ANSWER QUERY

During the preparation of your manuscript, the question listed below arose. Kindly supply the necessary information.

1. Please check if the proposed running head is okay.

END OF QUERY

Author Proof---not final version.