

What is a Complex System?

James Ladyman, James Lambert

Department of Philosophy, University of Bristol, U.K.

Karoline Wiesner

Department of Mathematics and Centre for Complexity Sciences, University of Bristol, U.K.

(Dated: February 27, 2011)

Complex systems research is becoming ever more important in both the natural and social sciences. It is commonly implied that there is such a thing as a complex system across the disciplines. However, there is no concise definition of a complex system, let alone a definition that all disciplines agree on. We review various attempts to characterize a complex system, and consider a core set of features that are widely associated with complex systems by scientists in the field. We argue that some of these features are neither necessary nor sufficient for complexity, and that some of them are too vague or confused to be of any analytical use. In order to bring mathematical rigour to the issue we then review some standard measures of complexity from the scientific literature, and offer a taxonomy for them, before arguing that the one that best captures the qualitative notion of complexity is that of the statistical complexity. Finally, we offer our own list of necessary conditions as a characterization of complexity. These conditions are qualitative and may not be jointly sufficient for complexity. We close with some suggestions for future work.

I. INTRODUCTION

The idea of complexity is sometimes said to be part of a new unifying framework for science, and a revolution in our understanding of systems the behaviour of which has proved difficult to predict and control thus far, such as the human brain and the world economy. However, it is important to ask whether there is such a thing as complexity science, rather than merely branches of different sciences, each of which have to deal with their own examples of complex systems. In other words, is there a single natural phenomenon called complexity, which is found in a variety of physical (including living) systems, and which can be the subject of a single scientific theory, or are the different examples of complex systems complex in ways that sometimes have nothing in common?

Hence, the fundamental foundational question in the domain of the complexity sciences is: What is complexity? Assuming that there is an answer to the latter question and that ‘complexity’ is not

just an empty term, we should then ask more particularly whether there is one kind of complexity for all the sciences or whether complexity is domain-specific?

In the next section, we review various attempts to characterize a complex system, and consider a core set of features that are widely associated with complexity by scientists in the field. We argue that some of these features are not necessary for complexity after all, and that some of them are too vague or confused to be of an analytical use. In Section III we briefly discuss the connections between complexity and probability and information, before explaining the Shannon entropy and the algorithmic complexity, and arguing that neither is adequate as a measure of complexity. Having introduced a distinction between ‘deterministic’ and ‘statistical’ complexity measures, in Section IV we review some standard measures of complexity from the scientific literature, and argue that the one that best captures the qualitative notion of complexity is that of the statistical complexity. In Section V, we scrutinise the alleged necessity and sufficiency of a number of conditions that have been proposed in the literature. We offer our own list of necessary conditions for complexity. These conditions are qualitative and may not be jointly sufficient for complexity. We then briefly discuss the very important idea of a hierarchy of organization in complex systems and close with some suggestions for future work. We do not arrive at a definition of a complex system in the sense of necessary and sufficient conditions but engage in explication of the notion of complex system, so that what is implicit in the science is made more explicit though not completely so. We can hardly think of any important scientific concepts that admit of an analysis into necessary and sufficient conditions but that ought not to dissuade us from introducing such clarity and precision as we can.

II. COMPLEX SYSTEMS AND THEIR FEATURES

Complexity science has been comparatively little studied by analytic philosophers of science, however, it has been widely discussed by social scientists and philosophers of social science. Rather than begin our investigations with what they have said about it we think it best to start with what practicing complexity scientists say about what a complex system is. The following quotations (apart from the last one) come from a special issue of *Science* on “Complex Systems” featuring many key figures in the field (*Science* 2 April 1999).

1. “To us, complexity means that we have structure with variations.”[16, p. 87]
2. “In one characterization, a complex system is one whose evolution is very sensitive to initial

conditions or to small perturbations, one in which the number of independent interacting components is large, or one in which there are multiple pathways by which the system can evolve. Analytical descriptions of such systems typically require nonlinear differential equations. A second characterization is more informal; that is, the system is "complicated" by some subjective judgment and is not amenable to exact description, analytical or otherwise." [49, p. 89]

3. "In a general sense, the adjective "complex" describes a system or component that by design or function or both is difficult to understand and verify. [...] complexity is determined by such factors as the number of components and the intricacy of the interfaces between them, the number and intricacy of conditional branches, the degree of nesting, and the types of data structures." [47, p. 92]
4. "Complexity theory indicates that large populations of units can self-organize into aggregations that generate pattern, store information, and engage in collective decision-making." [34, p. 99]
5. "Complexity in natural landform patterns is a manifestation of two key characteristics. Natural patterns form from processes that are nonlinear, those that modify the properties of the environment in which they operate or that are strongly coupled; and natural patterns form in systems that are open, driven from equilibrium by the exchange of energy, momentum, material, or information across their boundaries." [48, p. 102]
6. "A complex system is literally one in which there are multiple interactions between many different components." [36, p. 105]
7. "Common to all studies on complexity are systems with multiple elements adapting or reacting to the pattern these elements create." [1, p. 107]
8. "In recent years the scientific community has coined the rubric 'complex system' to describe phenomena, structure, aggregates, organisms, or problems that share some common theme: (i) They are inherently complicated or intricate [...]; (ii) they are rarely completely deterministic; (iii) mathematical models of the system are usually complex and involve non-linear, ill-posed, or chaotic behavior; (iv) the systems are predisposed to unexpected outcomes (so-called emergent behaviour)." [12, p. 410]
9. "Complexity starts when causality breaks down" [10]

The last citation well illustrates the difficulties of this field. Clearly many people will have a sufficiently permissive idea of causality to allow that there are causal relationships in complex systems, indeed many people will claim that complexity science's main job is to understand them. (1) may be true but is hardly informative unless we define what we mean by structure and variations. (2) asks us to chose between the conflation of complexity science with chaos and nonlinear dynamics, or the conflation of complexity with having a lot of components, or the conflation of complexity with a system with different possible histories on the one hand, and a completely subjective answer to our question. (3) and (4) takes us to more interesting territory. The computational notions of data structures, conditional branches and information processing are central to complexity science, and they will be of central importance in sections 3, 4 and 5. (5) introduces the central idea of nonlinearity. We argue in the next section that while many complex systems are subject to nonlinear dynamics, this is neither a necessary nor a sufficient condition for complexity. We endorse the view of (6) and (7) that a system cannot be complex unless there are many components interacting within it, but we argue this condition is not sufficient, and that it is of limited interest in so far as it is left vague what 'many' means. (8) introduces the idea of emergence that we argue is too confused to be part of an informative characterization of a complex system.

Abstracting from the quotations above and drawing on the culture of complexity science as expressed through a wide range of popular as well as academic sources, we arrive at the following list of properties associated with the idea of a complex system:

A. Nonlinearity

Nonlinearity is often considered to be essential for complexity. A system is linear if one can add any two solutions and obtain another, and multiply any solution by any factor and obtain another. Nonlinearity means that this superposition principle does not apply.

The interesting consequences of nonlinearity are often when the divergence of the system from the superposition principle is particularly extreme with respect to properties other than those specified by the microstate, such as, for example, the property of being alive or dead of an organism whose fate is determined by a nonlinear dynamical system like a skidding car. We often consider systems where we take fine-grained states, such as the positions and momenta of particles for example, as the inputs for dynamical equations, but we are really interested in the values of physical quantities that are coarse-grained with respect to the microstate. Non-linearity in the equations of motion may have the consequence that small differences in the values of the initial

conditions may make for radically different macrostates.[51]

In the popular and philosophical literature on complex systems a lot of heat and very little light is liable to be generated by talk of linearity and non-linearity. For example, Klaus Mainzer claims that “[l]inear thinking and the belief that the whole is only the sum of its parts are evidently obsolete” [30, p. 1]. It is not explained what is meant by linear thinking nor what non-linearity has to do with the denial of ontological reductionism. Furthermore, an obvious response to this kind of claim is to point out that it is perfectly possible to think in a linear way about systems that exhibit non-linear dynamics. Unfortunately the discussion of complexity abounds with non-sequiters involving nonlinearity. However, nonlinearity must be considered an important part of the theory of complexity if there is to be one since certainly many complex systems are also non-linear systems. Nonetheless, being subject to non-linear dynamics is not a necessary condition for a complex system. For example, there are structures involving linear matrices that describe networks which are a substantial part of what is studied by the complex sciences, and there are complex systems subject to game-theoretic and quantum dynamics all of which are subject to linear dynamics [38]. In general, feedback can give rise to complexity even in linear systems. Neither non-linear dynamics or linear dynamics can be necessary conditions for complexity since complexity scientists also study static structures. One may of course argue that such a complex synchronic structure could only come about through a dynamics that is nonlinear. This is the motivation for some of the conceptions of complexity discussed in Section IV. Non-linearity is also not sufficient for complexity not least because a simple system consisting of say a single chaotic pendulum can be subject to non-linear dynamics but is not a complex system.

Complexity is often linked with chaos and as noted above it may be conflated with it. There are systems that exhibit complexity in virtue of being chaotic. On the other hand, a completely chaotic system is undistinguishable from one behaving randomly. R.S. MacKay argues for a definition of complexity as the study of systems with many interdependent components and excludes low-dimensional dynamical systems [38] and hence many chaotic systems. Furthermore, since chaotic behaviour is a feature of deterministic systems, any dynamical system that is stochastic will not be chaotic, and complexity scientists study many such systems. So it seems that chaos and nonlinearity are both neither necessary nor sufficient for complexity.[52]

However, we may suppose perhaps that in many cases non-linearity in some guise, usually of dynamics, is at least a necessary part of some set of conditions that are jointly sufficient for complexity. (But there may be more than one such set.)

B. Feedback

Feedback is an important necessary condition for complex dynamical systems. A part of a system receives feedback when the way its neighbours interact with it at a later time depends on how it interacts with them at an earlier time. Consider a flock of birds. Each member of the group takes a course that depends on the proximity and bearing of the birds around it, but after it adjusts its course, its neighbours all change their flight plans in response in part to its trajectory; so when it comes to plan its next move, its neighbours' states now reflect in part its own earlier behaviour.

The presence of feedback in a system is not sufficient for complexity because the individuals need to be part of a large enough group to exhibit complexity, and because of how the feedback needs to give rise to some kind of higher level order, such as, for example, the behaviour of ants who are able to undertake complex tasks such as building bridges or farms even though no individual ant has any idea what they are doing, and left to their own they will exhibit much simpler behaviour. The ants behave as they do because of the way they interact with each other.

An abstract way of representing the prevalence of feedback in a complex system is provided by the theory of causal graphs. A chain of causal arrows indicates no feedback while a graph with loops of causal arrows shows feedback. In many contexts feedback is used by a control system, the paradigm of which is the Watt steam regulator, where the speed of rotation of the device interacts in a feedback loop with the steam engine to control the speed of the engine. However this is not a complex system because it has a central controller who sets up the machine. Control theory is importantly related to complexity theory because another central idea associated with complex systems is that of order, organization and control that is distributed and locally generated (as with the ants) rather than centrally produced (as with the steam regulator). Feedback can also be used for error correction, for example, in motor systems in the brain. We return to this in Section III.

C. Spontaneous order

Given the above it is clear that a fundamental idea in complex systems research is that of order in a system's behaviour that arises from the aggregate of a very large number of unco-ordinated interactions between elements. However, it is far from easy to say what order is. Notions that are related include symmetry, organization, periodicity, determinism and pattern. One of the most confusing issues is how order in complex systems relates to the information content of states and dynamics construed as information processing. The problem is that the interpretation of states

and processes as involving information may be argued to be of purely heuristic value and based on observer-relative notions of information being projected onto the physical world. We return to the role of information theory in answering our central questions in Section III. For now we note that the notion of order may mean so many things that it must be carefully qualified if it is to be of any analytical use in a theory of complex systems, but we note also that some such notion is central because pure randomness is sufficient for no complexity whatsoever. On the other hand, total order is also incompatible with complexity. The fact that complex systems are not random but also not completely ordered is of central importance in what follows. Nonetheless, it is a necessary condition for a complex system that it exhibit some kind of spontaneous order.

D. Robustness and lack of central control

The order in complex systems is said to be robust because, being distributed and not centrally produced, it is stable under perturbations of the system. For example, the order observed in the way a flock of birds stay together despite the individual and erratic motions of its members is stable in the sense that the buffeting of the system by the wind or the random elimination of some of the members of the flock does not destroy it. A centrally controlled system on the other hand is vulnerable to the malfunction of a few key components. Clearly, while lack of central control is always a feature of complex systems it is not sufficient for complexity since non-complex systems may have no control or order at all. A system may maintain its order in part by utilizing an error-correction mechanism (we return to this in Section IV). Robustness seems to be necessary but not sufficient for complexity because a random system can be said to be robust in the trivial sense that perturbations do not affect its order because it doesn't have any. A good example of robustness is the climatic structure of the Earth's weather where rough but relatively stable regularities and periodicities in the basic phenomena of wind velocity, temperature, pressure and humidity arise from an underlying non-linear dynamics. Note that these latter properties are course-grainings relative to the underlying state-space. That such properties exist and enable us to massively reduce the number of degrees of freedom that we need to consider is the subject of the next section.

E. Emergence

Emergence is a notoriously murky notion with a long history in the philosophy of science. People talking about complexity science often associate it with the limitations of reductionism. A strong, perhaps the strongest, notion of emergence is that emergent objects, properties or processes exhibit downwards causation. Upwards causation is uncontroversial in the following sense: a subatomic decay event may produce radiation that induces a mutation in a cell that in turn causes the death of an organism. The biological, chemical, economic and social worlds are not causally closed with respect to physics: economic effects may have physical causes. On the other hand, many people take it that the physical world is causally closed in the sense that all physical effects have physical causes. This immediately raises the question as to how complexity relates to physicalism, whether the latter is understood in terms of causal completeness or merely in terms of some kind of weak asymmetric supervenience of everything on the physical.

There is a sense in which approximately elliptical orbits emerge over time from the gravitational interaction between the sun and the planets. This is not the notion of emergence at issue here. Rather we are concerned with the kind of emergence exemplified by the formation of crystals, the organisation of ant colonies and in general, the way that levels of organisation in nature emerge from fundamental physics and physical parts of more complex systems. There is much controversy about how this happens and its implications. Again we conclude that the notion of emergence would need to be very precisely characterized to avoid simply adding to our confusion about the nature of complex systems. It is coherent to define emergence as an increase in complexity, and so it is preferable not to use the concept in the definition of complexity itself as both concepts seem to be at a similar level of generality.

Emergence is either purely epistemological, in which case it can be strong or weak depending on whether the lack of reduction is in principle or merely in practice, or it is ontological. In its latter form no consensus exists about how to understand it, although it is truisitic to say that there is a very important sense in which the physical interaction of atoms and molecules with light and electricity and magnetism and all other physical entities have led to the emergence of the immensely complex and structured system of life on earth including the human brain and the complexity of human culture and social life. Unless one is prepared to say with some metaphysicians that all this enjoys only some kind of second grade existence or even none at all, then it seems one must embrace ontological emergence in some sense. The problem then is with resisting or avoiding the argument that one's position requires that one accepts downwards causation at pain of making

most of reality causally inert or abstract rather than concrete. While non-reductive physicalism without the violation of physicalism may be defensible these issues are not to be solved en passant in characterizing complexity. Certainly we must say that emergence in all epistemological senses is necessary for complex systems. If a system doesn't exhibit higher-level order as discussed above, then it is not complex. However, emergence is not sufficient because, for example, an ideal gas exhibits emergent order but is not a complex system.

F. Hierarchical organisation

In complex systems there are often many levels of organization that can be thought of as forming a hierarchy of system and sub-system as proposed by Herbert Simon in his famous paper 'The Architecture of Complexity' [42].

The ultimate result of all the features of complex systems above is an entity that is organized into a variety of levels of structure and properties that interact with the level above and below and exhibit lawlike and causal regularities, and various kinds of symmetry, order and periodic behaviour. The best example of such a system is an ecosystem or the whole system of life on Earth. Other systems that display such organization include individual organisms, the brain, the cells of complex organisms and so on. A non-living example of such organization is the cosmos itself with its complex structure of atoms, molecules, gases, liquids, chemical kinds and geological kinds, and ultimately stars and galaxies, and clusters and superclusters.

G. Many more is different

There is a famous paper about complexity called 'More is Different' [35] but it is implicit that many more than a handful of individual elements need to interact to generate complex systems. The kind of hierarchical organization that emerges and gives rise to all the features we have discussed above, only exists if the system consists of a large number of parts, usually engaged in many interactions. We return to this in Section V.

The above discussion makes clear that the definition of complexity and complex systems is not straightforward and potentially philosophically interesting. The notions of order and organization introduced above and the idea of feedback are suggestive of an information-theoretic approach to complexity, since complex systems can be construed as maintaining their order and hierarchical organization by the exchange of information among their parts. Many people think that it is

fruitful to think of complex systems as characterized by the way they process information, as well as by the information theoretic properties of the data we obtain by sampling them. Since the notion of information is itself a philosophically problematic one, in the next section we explain the fundamentals of information theory to approach a mathematical theory of complexity.

III. COMPLEXITY, INFORMATION AND PROBABILITY

Complex systems are often characterized in information theoretic terms. Whether or not it makes sense to talk about physical systems containing and processing information independently of our representations of them is a vexed issue. It has recently been argued by Chris Timpson in the context of quantum information theory that talk of information flowing in a system is based on a confusion between abstract and concrete nouns [43]. His view is that information is not a physical quantity. On the other hand, there are radical proposals such as that popularized by the great physicist John Wheeler to take information as the primitive component of reality from which other physical properties are derived. We will not take a stand on these issues here but rather address the question of whether or not the standard measure of information content can be used to measure complexity. We remain neutral about the ontological status of information itself (and in particular about how information theory relates if at all to the semantic conception of information).

Information theory is closely connected with probability. Shannon was concerned with the statistical reliability of communication channels from an engineering perspective and invented a measure of information that can be applied to any probability distribution. At the heart of the association between complexity and information is the idea that complex systems can be understood as systems that are non-trivially structured without being completely ordered, and the order that they do possess can sometimes be understood as maintained by the internal processing of information. The question is whether this can all be made precise. One thing that seems certain is that complex systems are not ones in which all the parts are performing random walks, and this is enough to distinguish complexity from being complicated. A gas at equilibrium is an incredibly complicated system but it is not complex just because the behaviour of its parts is effectively random in the sense that the behaviours of individual molecules are largely uncorrelated. The mathematical theory that makes precise the notion of a random number is called algorithmic complexity theory. It has been claimed to be more fundamental than information theory in so far as the key result of the latter can be derived from it [24]. In Section III A we review the Shannon entropy, before explaining the algorithmic complexity in Section III B. In Section III D we make

the first of two distinctions among complexity measures and point out what all the above measures have in common, before arguing that this feature makes none of them suitable as measures of the complexity of physical systems.

A. The Shannon entropy

The Shannon entropy, usually denoted by H , was introduced to characterize different sources of information in relation to how much capacity a communication channel would need to carry messages from them. The mathematical form of the Shannon information is based on the application of the law of large numbers to very long sequences of letters in an alphabet. Since some possible sequences are typical in the sense of having the average number of occurrences of each character of the alphabet, and since in the limit almost all sequences will be typical sequences, it is possible to compress messages by replacing every typical message by a code number (that is shorter than the message itself). Shannon's noiseless coding theorem states that the optimal compression is measured by the Shannon entropy.

The Shannon entropy is formally defined as the sum of the probability of each letter of the alphabet being sent by the source multiplied by the logarithm to base two of that probability.

$$H = - \sum_{i=0}^n p_i(x) \log p_i(x)$$

For a source where every letter has the same probability of occurring, the Shannon entropy is maximal. For example, a string of zeros and ones in exactly equal proportions means a source that has a Shannon entropy of one. Receiving a particular message from a low-entropy source means receiving a message from a source in which some letters are much more likely than others. In a sense the message carries less information than a message from a source that has a higher Shannon entropy. The Shannon entropy is often used as a measure of information for this reason since it says how much information a source is sending in respect of how much the average message can be compressed, however, it does not measure the information content of any particular message.

Nonetheless, given a large single string of letters, we can use it to calculate the Shannon entropy of the source that produced it by assuming that the relative frequencies of the letters in the message match their overall probabilities across all the messages the source sends. If we consider a string that has a random assortment of letters then it will have a high Shannon entropy and if it is highly ordered its Shannon entropy will be low. In general we can consider the Shannon entropy of any probability distribution. A uniform probability distribution has maximal entropy and in this sense

a higher information content than a very specific probability distribution. Hence, the Shannon entropy is a measure of randomness not of complexity.

B. The algorithmic complexity

The algorithmic complexity, usually denoted by K , of an object (which was formulated independently by Solomonoff, Kolmogorov and Chaitin in the mid-sixties) is a measure of the length of the shortest computer programme that prints out the standard enumeration of the object [22, 24]. It is easiest to think in terms of binary strings and to associate every mathematical object (numbers, vectors, functions, and so on) with the string that describes it. In the case of numbers this will be just the binary expansion of the number, in the case of vectors it will be a string giving the binary expansions of the components of the vector with respect to some basis, in the case of functions it will be a string giving the binary expansions of the ordered pairs of elements in the domain and range of the function.

The algorithmic complexity of a string is obviously greater for random strings than ordered ones. Any repetition of digits or symmetry in a string will allow it to be compressed and so make the shortest programme that outputs the string shorter than the string itself. On the other hand, a truly random string cannot be compressed at all and so the shortest programme that outputs the string will have to contain the string itself. It is important that this measure is independent of the source of the string. For example, a string of apparently random numbers might in fact have been generated from the decimal expansion of π according to some rule, but this fact does not show up in the string itself.

C. Lempel-Ziv Complexity

A measure listed in many accounts of complexity is the Lempel-Ziv complexity. It is a compression scheme which approximates the algorithmic complexity of a given string. Hence, the Lempel-Ziv complexity is a measure of randomness and not of complexity in our sense. It is the algorithmic analogue to the Shannon entropy in that it uses any repetition in the data to achieve a higher compression. It does this, however, without having access to the statistics of the source. In practice, Lempel-Ziv coding is a very efficient data compression algorithm. Hence the Lempel-Ziv complexity is computable but it is deterministic in the following sense.

D. Deterministic versus statistical conceptions of complexity

One of the most intriguing ideas is that complexity lies between order and randomness. The latter two notions are here meant in the sense of algorithmic complexity theory. Systems of interest to complexity scientists are neither completely ordered nor completely random. What is needed then would seem to be a measure of complexity different from the algorithmic complexity and we make this argument more explicit below.

We first note that measures of complexity may be deterministic or statistical in the following sense. A deterministic measure of complexity will treat a random sequence of 0s and 1s as having high complexity. Although the Shannon entropy is a function over a probability distribution and the algorithmic complexity is a function of an instantiation, both are deterministic in this sense.

If we are interested in a measure that tells us when we are dealing with a complex physical system we clearly do not want a measure that is maximal for completely random data strings. Note also that the algorithmic complexity and the Shannon entropy of the data produced by a biased coin would be lower than that of a fair coin although there is no sense in which the former is more complex than the latter. Both functions are monotonically increasing with randomness whereas we want a function that is unimodal and peaked between randomness and complete order. Hence, neither Shannon entropy nor algorithmic complexity of a data string are an appropriate measure for the complexity of a physical system. We need a measure that is statistical in the sense defined above.[53] In the next section we review some important measures of complexity in the literature.

IV. MEASURES OF COMPLEXITY

There are various purported measures of complexity in our sense in the literature, some of which will be reviewed here. If we had an analytical framework for their classification we may then be able to decide what the appropriate measure or class of measures for complexity is on the basis of whether or not they meet appropriate criteria. We have already argued that an appropriate measure of complexity should not be maximal for either for random or for highly ordered systems, and hence that only what we call ‘statistical’ measures of complexity are plausible candidates. The following are quotes from scientists on further potentially necessary conditions for a good measure of complexity.

“Such a measure [...] should assign low complexity to systems in random states and in ordered

but regular states [...]. Complexity is then a measure of how hard it is to put something together.” [28][p. 189]

“A great many quantities have been proposed as measures of something like complexity. [...] A measure that corresponds much better to what is usually meant by complexity in ordinary conversation, as well as in scientific discourse, refers [...] to the length of a concise description of a set of the entity’s regularities.” [13][p. 1-2]

“Three questions that are frequently posed when attempting to quantify the complexity of the thing [...] under study are 1) How hard is it to describe? 2) How hard is it to create? 3) What is its degree of organization?” [27][p. 7]

Note that notions of complexity may be applied to three targets, namely, to the methods used to study certain systems, to data that are obtained from certain systems, or to the systems themselves. Some say that the complex sciences are simply those that use certain characteristic methods. In this view it makes little or no sense to speak of complex data sets or complex systems. Call this a pragmatic account of complexity. At the other extreme, some measures of complexity are supposed to be applied directly to systems to tell us whether they are complex or not. Call this a physical account of complexity. Among physical accounts of complexity are the theories of logical depth and thermodynamic depth discussed in the next section. These views allow for derivative notions of complex data (the data produced by complex physical systems), and complex methods (the methods appropriate to studying such systems). Finally, a third kind of account of complexity applies primarily to data and derivatively to physical systems and to the methods appropriate to their study.

The first quote above suggests that it might be possible to define and measure the complexity of a system by considering its history, and this is the approach taken by both Charles Bennett with his measure of logical depth [3], and by Lloyd and Pagels [29] with their notion of thermodynamic depth. Both accounts are attempts to rigorously define various intuitive notions of what it is for a system to be complex. And their authors agree on two fundamental facts: A physical system is complex or otherwise in virtue of its history; and complexity lies somewhere between complete order and complete randomness. In the next two sections we review these two prominent measures of complexity and others.

A. Logical depth

Charles Bennett’s approach begins with the intuition that some objects, such as, for example, the human body contain “internal evidence of a non-trivial causal history” [3, p. 227]. He formalizes the idea of an object that has such a history in the language of algorithmic information theory. Bennett’s key idea is that long causal processes are usually necessary in order to produce complex or ‘deep’ objects; deep objects are produced by short processes only with very low probability.

Suppose, following William Paley[33, p. 7-8], you find a pocket watch on a deserted beach, and consider all possible explanations for the watch’s presence. Some explanations will involve a watchmaker who, after training for many years finally manufactures the watch in question, plus some further story about how the watch came to be on the beach, such as that the chain snapped while the owner was walking. However, other explanations could be altogether more unusual, for instance, perhaps the moment you set foot on the beach the impact of your foot caused atoms in the sand grains to leap into the exact configuration of a pocket watch awaiting your discovery. Intuitively some of these explanations are more plausible than others. Bennet’s claim is that complex objects are those whose most plausible explanations describe long causal processes, and that hypothesising the objects to have had a more rapid origin forces one to adopt drastically less plausible stories. Here ‘explanation’ means a complete description of the object’s entire causal history.

Bennet’s challenge now is to say exactly how we grade various causal histories in terms of plausibility; why is the watchmaker hypothesis a better explanation than the springing-into-existence hypothesis? To this end Bennett turns to algorithmic information theory. Bennett proposes that the shortest programme for generating a string represents the a-priori most plausible description of its origin, while a print program on the other hand offers no explanation whatsoever, and is equivalent to saying ‘it just happened’ and so is effectively a null-hypothesis. If we view ad-hocness in an hypothesis as the use of arbitrary unsupported assertions that could be explained by a more detailed theory, then a hypothesis represented by a programme s -bits longer than the minimal programme is said to suffer from s -bits of ‘ad-hocness’ over the hypothesis represented by the minimal programme. Consider Bennett’s example; an observer in possession of the first million digits of π (highly compressible) is most likely to hypothesise some intelligent/computational/mechanical origin and reject alternative print program hypotheses that it simply appeared by chance since the print programme null-hypothesis involves a huge amount of ad-hocness over the minimal programme. The logical depth of a string depends on the running time of the programmes that

produce it. After reviewing several tentative definitions Bennett settles on the following one.

Logical Depth A string's *depth* denoted $D_s(x)$, will be defined as ... the least time required to compute it by an incompressible program.[54]

How successful is Logical depth as a measure of complexity? As a formalised account of an intuitive concept we should ask, firstly, how philosophically sound are the motivating intuitions? And second, to what extent does Bennett's formal theory agree with these intuitions? To recap the motivating intuitions are as follows.

1. Complex objects lie somewhere between complete order and complete disorder.
2. Complex or logically deep objects cannot be produced quickly so any adequate complexity measure should obey a *slow growth law*.
3. The history of a complex object is not only long, but non-trivial, that is, the most plausible explanation for the object's origin is one that entails a lengthy computation/causal process.

Recalling the characteristics of complex systems listed in Section 2.1, note that robustness may be formulated in computational language as the ability of a system to correct errors in its structure. In communication theory error correction is achieved by introducing some form of redundancy. This redundancy need not be explicit such as a copy of the string or its parts. It may be more subtle, for instance, exploiting parity checking [11] which is more computationally intensive but also more efficient (the message is shorter) than simple duplication. Bennett specifically mentions error correction:

“Irreversibility seems to facilitate complex behavior by giving noisy systems the generic ability to correct errors.”

The problem can be viewed like this: A living cell is arguably a paradigm complex object and does indeed have the ability to repair itself (correct errors), for instance a malfunctioning component may be broken down and released into the surrounding medium. Contrast the cell with a shallow object such as a gas in a box, a small perturbation of this gas is rapidly dispersed *with no limitations* to the many billions of degrees of freedom within the gas. The cell on the other hand has a one way direction for this dispersal, errors within the cell are transported out, and errors outside the cell are kept out (assuming, in both cases, the errors are sufficiently small).

Presumably error correction is important because objects capable of error correction are able to protect themselves from ‘environmental noise’ and therefore persist longer, thus meeting the long causal history requirement. However, it is unclear how error correction is reflected if at all in the definition of logical depth. Perhaps it is related to the stability issue discussed earlier, namely, if we introduce a small error in an otherwise deep object it should remain deep via the slow growth law. However this cuts both ways, shallow objects remain shallow in the presence of noise too.

The biggest problem with logical depth is that the function $D_s(x)$ is not computable because it is defined in terms of the algorithmic complexity which is provably non-computable. Returning to the philosophy of science metaphor the non-computability of $K(x)$ is related to the classic problem of underdetermination; for any set of data, not only are there many theories compatible with those data, but we have no way of telling which is a-priori the most plausible (the one with length = $K(x)$). If the complexity of an object is to be calculable, even if only in principle, then there is something wrong with Bennett’s definition. There is no reason to suppose that there could not be some true property of systems that measures their complexity even though we cannot compute it. However, we take it that a measure and associated characterization of complexity that can actually be used is to be preferred. We return to this issue in section 4.5. We now turn to another popular measure of complexity.

B. Thermodynamic depth

The notion of thermodynamic depth introduced by Lloyd and Pagels [29] shares much of its informal motivation with logical depth.

1. Complexity lies between order and disorder.
2. The physical notion of complexity should correspond to the notion of the computational complexity of a function f when the physical system in question is a computing device calculating the function f .
3. Complexity is not a property of a single physical state but rather a property of a process.

Items 1 and 3 bear some resemblance to the beginning of Bennett’s paper. Item 2 is somewhat surprising and is discussed further.

3 is motivated by the fact that Lloyd and Pagels do not consider multiple copies of a complex object to be significantly more complex than a single complex object, since, as they point out,

producing a single bull takes billions of years of evolution, but seven more bulls can be easily produced with a single bull and seven cows. The process from no bulls to seven bulls is only minutely harder than the process from no bulls to one bull. Putting aside concerns about the biological possibility of a single bull in isolation the idea, much like Bennett's, is that complexity has something to do with how difficult it is to produce something.

Like Bennett, Lloyd and Pagels consider the set of histories or *trajectories* that result in the object in question. A trajectory is an ordered set of macroscopic states $a_i, b_j \dots c_k$. If we are interested in the complexity of a state d , then we consider all trajectories which end in a state in d .

Their next step is to introduce three requirements that any reasonable measure of complexity defined over these trajectories should obey.

1. The measure must be a continuous function of the probabilities of the trajectories.
2. If all trajectories are equally likely then the measure should be monotonically increasing in n , the number of trajectories.
3. Additivity: the measure associated with going from state a to b to c should be equal to the sum of the measures for going from a to b and from b to c. In other words the complexity of building a car from scratch is the complexity of building the components from scratch and building the car from the components.

These are effectively the same demands made by Shannon in his pioneering paper on information theory [41] and it is well known that there is only one function that satisfies them (up to a constant) namely the entropy H we defined in section III A. The depth of a state is defined to be the Shannon entropy of the probability distribution over the set of trajectories that can lead to that state. If we happen to know that a state d was in fact a result of some trajectory $a_i, b_j \dots c_k \dots d$ then the depth of d is given by

$$\mathfrak{D}(d) = -k \ln p(a_i, b_j \dots c_k \dots d)$$

(where k is Boltzmann's constant)

How the probabilities p are to be interpreted is unclear but Lloyd and Pagels reject a subjective reading.

... although the complexity depends on the set of experiments that determine how a system reaches a given state, the measure as defined is not subjective: two different physicists given the same experimental data will assign the state the same complexity.[29, p. 190]

Because their discussion draws heavily on concepts from statistical mechanics, such as those of phase spaces, microscopic/macroscopic states and so on, it is reasonable to suppose the probabilities used in defining a system's depth are of the same sort as those found in statistical mechanics. But the correct interpretation of these probabilities is still controversial with some prominent authors, notably Jaynes [20, 21], arguing for a purely epistemic interpretation. If the biggest problems thermodynamic depth faces were those it inherits from statistical mechanics then Lloyd and Pagels could be relieved because, while statistical mechanics faces many philosophical problems like most (if not all) successful physical theories, the subject has hardly been brought down by them. However, bigger problems are highlighted by Crutchfield and Shalizi [7]. First we are not told how long the trajectories we consider should be. Complexity is claimed to be a property of a process but how do we identify when a process starts? To remedy this Crutchfield and Shalizi suggest looking at the rate at which depth increases, rather than its absolute value. The rate of increase of depth turns out to be the rate at which the system's entropy increases. Entropy rate, however, is a measure of randomness and does not capture complexity in the here desired way. This renders the thermodynamic depth itself unsuitable as a measure of complexity. Lloyd and Pagels specifically state that a reasonable complexity measure should not award high complexity to a purely random process but thermodynamic depth does exactly this. An N body ising spin system being cooled below its critical temperature assumes a frozen magnetic state 'up', 'down'... which has very many probable predecessors which results in a very large thermodynamic depth of this frozen state again showing that we must look elsewhere for a measure of complexity.

C. Effective complexity

Effective complexity, introduced by Gell-Mann [13, 15] (see also [14]), is a statistical measure of complexity based on Kolmogorov complexity. Gell-Mann turns Kolmogorov complexity into a statistical measure of complexity by asking for a shortest description, not of the entity itself, but of the ensemble in which the entity is embedded as a typical member. Here, 'typical' means that the negative logarithm of its probability is no larger than the entropy of the ensemble. The resulting definition of *effective complexity* of a string is the algorithmic information content of the ensemble

in which the string is embedded. The algorithmic information content of an ensemble is the length of the shortest programme required to specify the members of the ensemble together with their probabilities. Effective complexity is designed to measure the regularities of a string as opposed to the length of a concise description. The problem with this measure, as with logical depth, is that it is not computable, as Gell-Mann states himself: “There can exist no procedure for finding the set of all regularities of an entity.” [13, p. 2]. From a conceptual point of view, however, it is consistent with the above discussed notions of complexity. A random string is assigned zero effective complexity. Maybe more surprising, the digits of π have very low effective complexity. The Mandelbrot set, as well, has a very low effective complexity.

Gell-Mann considers effective complexity to be complimentary to logical depth. Whereas logical depth measures the time it takes to compute a string from a given prescription, the effective complexity measures the size of that prescription. A trade-off becomes apparent, since a very compressed description, such as a physical law – which has low effective complexity – might have a large logical depth since the computation of any particular solution can be computationally very costly.

D. Effective measure complexity

Introduced by Grassberger, the effective measure complexity, also known as excess entropy, shared information and others, is measuring the mutual information between two halves of a data string. In the case where the data string is a time sequence the effective measure complexity measures the mutual information between past and future. Measured in bits, it can be interpreted as the average number of bits a process transmits into the future. For a random string this is zero, as expected. The more structured or correlated the string is the higher the measure will be. Hence it is a measure of complexity in our sense and it is computable. The caveat of this measure is that it can only be applied to sequences with stationary probability distributions. It has been shown that it is a lower bound to the statistical complexity.

E. Sophistication

Sophistication, introduced by Atlan, is a non-quantitative measure of complexity and hence not truly a measure in even the loose sense. It is part of many lists of complexity measures because it captures qualitatively the complexity of a system. It is related to the excess entropy (effective

measure complexity) in that it also is the subextensive part of an entropy function. In the case of the excess entropy it is the Shannon entropy, for the sophistication it is the algorithmic information. Since the latter is non-computable, sophistication itself is non-computable.

F. Statistical complexity

Neither entirely ordered, nor entirely random processes qualify as ‘complex’ in the intuitive sense since both admit of concise descriptions. Ordered processes produce data that exhibit repetitive patterns describe the string ‘0101...0101 consisting of 1000 bits concisely as:

`‘01’ 500 times`

Whereas random ‘noise’ such as the outcomes of 1000 coin-toss trials, while lacking periodic structure, affords a concise statistical description:

`‘0’ with probability 0.5, ‘1’ otherwise`

The Kolmogorov Complexity[22–24, 26] and Thermodynamic Depth[7, 29] award random strings high complexity; the former by defining a random string as one whose minimal program is approximately the length of the string itself, the latter by measuring a systems rate of change of entropy. These measures assign high complexity to randomness because they ignore its statistical simplicity. The statistical complexity measure was developed to avoid this problem. Statistical complexity is the measure for structural complexity arising out of computational mechanics which is a framework developed by J. Crutchfield and co-workers [6, 8, 39]. Computational mechanics defines a process’s causal states and gives a procedure for finding them. A process is a bi-infinite sequence of random variables, which one can think of as a measurement sequence. A causal-state representation – a so-called ϵ -machine – is the minimal one consistent with prediction.[55] It is an application of Occam’s razor since it is assumed that the minimal representation is to be preferred, and showing the conditions under which it is optimal (namely statistically accurate) leads to the ϵ -machine. To a first approximation, the size of the ϵ -machine is a measure of structural complexity of the process from which the machine was reconstructed. The exact value for the statistical complexity of a process is derived as follows. From the measurement sequence – the process – probabilities for future observations conditioned on past observations are gathered and grouped into equivalence classes

of equal probabilities. Here, future is not necessarily to be taken literally. It could be a sequence in space just as well. All observations are attributed to one equivalence class and one only. These equivalence classes, or minimal sufficient statistics as they are called in the statistics literature, are the causal states mentioned above. Next, the probability of a future observation given a preceding sequence of past observations determines the probability of going from one equivalence class to the next while predicting the process. Hence, the equivalence classes, or causal states, represent our ability to predict the process's future behaviour. The final ϵ -machine consists of a finite number of causal states and stochastic transitions between them, that is called hidden Markov model by some, stochastic finite-state automaton by others. The statistical complexity, denoted C_μ , is the Shannon entropy over the stationary distribution over the causal states \mathcal{S} :

$$C_\mu = H(\mathcal{S}) \tag{1}$$

Note, that in this framework randomness is considered to be trivial and does not require any predictive capacity. Hence the ϵ -machine reconstructed from a random sequence would have one state only and zero statistical complexity. If the above outlined procedure, which is not unique, leads to an infinite number of causal states one goes up a computational hierarchy until one obtains a finite representation [6].

Computational approaches to complexity such as Bennet's Logical Depth[3] and Zurek's Physical Entropy[50] take the universal Turing machine[44] as their basic computational model. The universal Turing machine is the most powerful of the conventional computational models meaning that it computes the largest class of functions. Bennett and Zurek consider the computational resources (program size, space-time complexity etc.) needed to reproduce the given string verbatim. In contrast Crutchfield's approach is to consider the simplest computational model capable reproducing the statistical properties of the string. Analogously the task of quantifying a physical system's complexity amounts to finding the simplest computational model which captures the statistical properties of it. In this sense, the statistical complexity is not based on a Turing machine model but is more general and chooses the computational architecture which is appropriate to the process at hand.

Returning to the concept of an ϵ -machine, we now consider how it can be used to model the development of a system over time. The statistical complexity now measures the diversity of a system's behaviour over time. Observing a system over a long period of time allows repetitive behaviour, if it exists, to become apparent. We may think of the system as passing through a

set of generalised causal states, jumping from one to the other with a certain probability. An ϵ -machine represents each generalised causal state by a node in a graph with edges between nodes weighted according to the probability of the system making such a transition. The ϵ -machine with the smallest number of nodes needed to model a system's behaviour is known as the system's minimal ϵ -machine, which has been shown [39] to be optimal for predicting the system's future. Since our model is probabilistic to begin with we avoid the pitfall of assigning high complexity to random behaviour; a sequence of coin tosses may be modelled by a single state automaton with two equiprobable transitions.

Having identified a set of generalised causal states we may assign to each a probability according to the proportion of its time the system spends in each state, the Shannon entropy of these probabilities is the system's statistical complexity. For a system to have high statistical complexity it must have a large number of possible states each roughly equiprobable.

G. Computable versus non-computable measures

In this section it is argued that only computable measures are suitable as measures for the complexity of a physical system.

A good measure of complexity needs to be statistical as opposed to deterministic. This rules out thermodynamic depth and leaves us with effective complexity, logical depth, and statistical complexity. If we, in addition, impose computability as a requirement, which seems reasonable in so far as we are interested in using our measure of complexity to tell us which systems in the world are complex, then among the measures we have considered we are left with the statistical complexity as the only one that meets both requirements of being computable and statistical. In this case the primary application of the notion of complexity is to data sets. When we find a complex data set we may reasonably attribute the term 'complex system' to whatever system produced it.

We have not argued exhaustively by considering every putative measure of complexity in the literature, rather we have taken a representative sample of complexity measures and considered at least one that falls into each of the four categories of measure generated by our two binary distinctions among different kinds of measure. We do not attempt an exhaustive survey of such measures, since, that would be a Herculean task that would take up a whole long paper in itself. Our goal is to show that any adequate measure of complexity must have certain characteristics exemplified by the measure of statistical complexity due to Crutchfield.

H. Further remarks

A discussion of complexity measures in the context of dynamical systems can be found here [45]. The authors suggest a classification scheme into four categories similar to the scheme presented here. The scheme is devised such that it distinguishes between order and chaos. The distinction between those systems with a generating partition and those with a homogeneous partition is not applicable to anything but a dynamical system in the mathematical sense. The other property of structural or dynamical features is also not general. A discussion of complexity measures in the context of physical systems can be found here [2]. This book surveys a variety of different kinds of complex system and measures of complexity. Of the ones, mentioned there the grammatical complexity deserves mentioning. Opposed to the statistical measures reviewed so far grammatical complexity is concerned with higher level organisation while discarding finite-length correlations as not complex. Hence, it is worth to consider as a measure when the investigator is interested in higher-level order and infinite correlations. A survey of measures of complexity and their criticism can be found here [17]. The author concludes that no unified measure can be found. Nor should there be one since complexity in this view is intricately linked to meaning which is subjective and unmeasurable.

V. THE FEATURES OF COMPLEX SYSTEMS REVISITED

We have argued that there is something all complex systems have in common, namely that they are neither completely random nor completely ordered. Taking the statistical complexity as a good measure of complexity the interplay of order and randomness can be illuminated somewhat more. Indeed, a system with many levels of order and little if any randomness is also complex. The emphasis here is on the “many” levels of order. To quote Crutchfield and Young: “The idea is that a data set is complex if it is the composite of many symmetries.” [19, 227]. Randomness does not seem to be crucial here. What, then, is the reason for randomness to be stated as a necessary ingredient by so many authors? An answer to this is given by E. Morin in his treatise *Method* [32]. In his view the role of randomness is crucial to generate and sustain complexity. To summarise his view in our terms: Complexity could “be” without randomness, but it could not become nor sustain itself. The argument goes as follows. Assuming that initially there was no order. Elements can form structures and generate order only if they interact. The way they interact is by random movement, Brownian motion. Through chance encounters interactions take place and

correlations are built up. Hence, randomness acts as a source of interaction. Once correlations are built up a system may be complex. Because of the second law of thermodynamics, however, these correlations will decohere and the initial random state will be populated once more. Only, if the underlying mechanism of random encounters is maintained can decoherence be avoided and order maintained. Thus, there is no complex system which is static. They are all in a continuous process of maintaining their complexity. And hence randomness and order are intricately linked to complexity.

This leads us to the following tentative definition of complexity:

A complex system is an ensemble of many elements which are interacting in a disordered way, resulting in robust organisation and memory.

- Ensemble of many elements

Most definitions or descriptions of complexity mention ‘many elements’ as a characteristic. For interactions to happen and for pattern and coherence to develop, the elements have to be similar in nature. This is the prerequisite for the following condition, namely that of interaction. For systems to be able to *communicate* (in the broadest sense) with each other they have to be able to exchange energy or matter or information. Physical systems have to be particles comparable in size and weight, subject to the same physical laws. In biology, cells before they form multi-cellular organisms are indistinguishable / identical so they can maximally communicate (exchange genetic information) with each other. Non-physical systems, e.g. social structures have to be similar in character, behaviour, or rules obeyed. Other examples are the brain consisting of neurones, an ant colony consisting of ants, a financial market consisting of agents. While it is a necessary condition for a complex system that there are many similar parts of some kind it should be noted that not all the parts have to be similar and of the same kind. Weather systems are composed of many similar parts in so far as oxygen, nitrogen and water are all gases composed of molecules, yet of course these gases are importantly different in their properties and these differences give rise to important features of the climate. Moreover, weather systems also include geological features and radiation too. Also weather systems are composed of many similar parts in so far as there are many small volumes of the atmosphere that are similar and which interact to give rise to the weather. Every complex system has at least some parts of which there are many similar copies interacting. Of course, the idea of a large number of components is

vague, but we note that such vagueness is ubiquitous in the sense that it is vague whether, for example, a quantum system consists of a large number of particles so that it reduces to classical behaviour in accordance with the correspondence principle, or whether there are sufficient birds flying together to be regarded as a flock. [56]

- Interactions

The second condition is for the elements of the system to have the means of interacting. Interaction can be either exchange of energy, matter, or information. The mitigating mechanism can be forces, collision or communication. Without interaction a system merely forms a “soup” of particles which necessarily are independent and have no means of forming patterns, of establishing order. Here, we need to emphasise that interaction needs to be direct, not via a third party, a so-called common cause. Thus, we require not merely probabilistic dependence but direct dependence.

Locality of interaction is not a necessary condition, neither is it sufficient. Interactions can be channelled through specialised communication and transportation systems. Telephone lines allow for non-local interactions between agents of a financial market, nerve cells transport chemical signals over long distances.

It is important that the idea of interaction here is not necessarily one of dynamics. What is important is the idea of the dependence of the states of the elements on each other.

Non-linearity of interactions is often mentioned as necessary condition for a complex system. This claim can be easily refuted. Complex networks, which are a clearly complex in their structure and behaviour, are defined by matrices which are inherently linear operations. The fact that non-linearity is not necessary illustrates a commonly mentioned feature of complex systems. They can be perceived as complicated yet be defined by simple rules, where complicated means difficult to predict.

- Disorder

Disorder is a necessary condition for complexity simply because complex systems are precisely those whose order emerges from disorder rather than being built into them. Interaction is the basis for any correlations to build up and hence for order to arise from disorder.

- Robust order

The above three conditions are all necessary for a complex system to emerge but they are not sufficient because many similar elements interact in a disorderly way in a gas but they are not complex systems. However, a system consisting of many similar elements which are interacting in a disordered way has the potential of forming patterns or structures. An example are the Rayleigh-Bénard convection patterns. On an appropriate time scale the order is robust. This means that although the elements continue to interact in a disordered way the overall patterns and structures are preserved. A macroscopic level arises out of the microscopic interaction and it is stable (on that time or length scale). This kind of robust order is a further necessary condition for a complex system.

- Memory

“A system remembers through the persistence of internal structure.”[18] Memory is a straightforward corollary of robust order.

A. Statistical complexity revisited

Since statistical complexity is a measure applied to data, we offer the qualitative definition of a complex system:

A system is complex if it can generate data series with high *statistical complexity*.

Recall the features of complex systems discussed above:

1. Nonlinearity
2. Feedback
3. Emergence
4. Local organisation
5. Robustness
6. Hierarchical Organization
7. Numerosity

It is difficult to say how well the statistical complexity captures 1 and 2. Certainly some systems of high statistical complexity are nonlinear and exhibit feedback, for instance, see Crutchfield's investigation into the logistic map [19]. However, it is not yet clear that either are necessary or sufficient for high statistical complexity. Emergence, understood as an ontological thesis, is clearly consistent with high statistical complexity. But whereas emergentism precludes reductionism, a high statistical complexity is *prima facie* compatible with all kinds of reductions of high level phenomena reduce to lower level phenomena, as it could still be a useful measure. That said it is not clear what, if any, light statistical complexity might shed on emergence. There is an interesting suggestion to define emergence using the predictive efficiency of a process which is defined in terms of statistical complexity [40, p. 9].

As discussed earlier, local organisation and robustness are related, and the former appears necessary for the latter because order that is centrally produced may be easily disrupted by perturbing the point of control. Locally produced order, on the other hand, presents a picture of many distributed copies which serve as a back up to each other. Now a statistically complex system is one that exhibits a diverse range of behaviour over time i.e the ϵ -machine has many vertices. In contrast a system with no memory, say a coin being tossed repeatedly, can be modelled accurately with a one state machine with two equiprobable loop edges, 'heads' and 'tails', and thus has complexity of zero. Complex systems, unlike a coin toss must possess some memory, some record of their past, with an ϵ -machine we can infer something of its history by considering its current state and the probabilities of the paths leading up to it, this amount of memory is quantified by equation 1. Local organisation and robustness appear to be related by this idea of memory; both memory and robustness involve stability over time, and for this stability we need some form of back-up or redundancy with which to correct errors, this may be provided when we have a system whose order is not centrally controlled.

B. The hierarchical nature of complex systems

We now turn to the underlying architecture – by which we mean relative organisation of the of the elements – which leads to properties such as robustness or de-central control. As pointed out by H. Simon [42] a hierarchical organisation does lead to properties which are assigned to complex systems. A system is considered to have the form of a hierarchy when it is being composed of subsystems that, in turn, have their own subsystems, and so on. Many social system, such as business firms or universities, are hierarchical. A university, for example, consists of staff which are

grouped in departments which, in turn are grouped in faculties. The highest level of a university is the vice-chancellor. One can go up the hierarchy even further to networks of universities and so on. Not only social systems exhibit hierarchical organisation. Similarly, a biological organism is made of cells which, in turn, are made of nuclei, cell membranes, and mitochondria, etc. The parts of a cell are made of molecules, which are made of atoms, and so on. Clearly, natural and social hierarchical systems are in abundance. There is a subtle difference between social and physical systems in terms of the nature of the hierarchical levels to which we will come back to at the end of this Section.

A core problem in complex systems research is how a complex structure can have evolved out of a simple one. In particular, if the process of evolution (by which we do not mean biological evolution exclusively but any living or non-living process) is following a random walk through the space of possibilities the likelihood of anything non-random to evolve is negligible. This mystery of how complexity can evolve in a short time span can be solved with hierarchical evolution. Indeed, a “selective trial and error” can lead to complex forms on a geologically meaningful time scale. As Simon points out, in such a “selective” process each intermediate level forms a stable configuration which gets “selected” for further levels to build on top. In this way, the hierarchical nature of complex systems is explained by the efficiency of hierarchical building processes. Here, natural evolution and engineering tasks are very similar. Neither would be very successful if the final product was to be built from scratch every time (see the earlier discussion of the watch maker). This process of assembling stable intermediates to new stable configurations does not presume any particular nature of the subsystems. Hence the same reason for stable intermediates leading to complex systems applies to physical, biological, as well as social systems. This is why stable democracies cannot be created over night. It is apparent that hierarchical architecture and evolution of complex systems are intricately linked.

As we have argued, the robustness of a complex system arises from the hierarchical process which generates the complex system. In the same way, many other common properties of complex systems can be explained. The absence of a central control, for instance, springs automatically from its hierarchical structure. We illustrate this with an example. A cell is composed of molecules which are composed of atoms etc. Thus, an organism is composed of a multitude of identical particles. None of these has any particular role from the outset. The particles are equipped with the same interaction forces and dynamical degrees of freedom. The different levels of the hierarchy are made up of regroupings of lower levels, of a redefining of the system’s boundaries. Therefore, physical systems do not have a central control, since the highest level of the hierarchy is merely

the sum of the levels below.

Here we find a general difference between physical or biological systems on the one hand and social systems on the other that is not mentioned in Simon's discussion. In a social system an element at a higher level does not necessarily consist of elements at the level below. The CEO of a company, for example, does not 'consist' of her employees. She is, indeed, a novel central control element. This being said, many social systems do exhibit complex features. For example, the stock market is highly complex in the sense that there are many traders interacting, equipped with buying and selling capacity, and this gives rise to feedback, as in a market panic. The resulting dynamic is highly unpredictable though and it is often claimed that the market as a whole behaves like a random walk. Hence, whether or not a company is a complex system in the full sense of the term is an open question. Although the interaction between many companies make again for a good candidate.

The connection between social and physical hierarchies can be regained to some extent when we consider the additional element of representation. A head of a political party, for example, not only controls the party but also represents its members. Social hierarchies can combine to larger and larger systems through the mechanism of representation. An organisation like the United Nations could not exist if we didn't allow for representative governance. There is no such thing as "representation" in most physical systems and this makes the robust communication between millions of molecules even more impressive.

The hierarchical architecture of complexity is linked to the statistical complexity as a measure of a complexity. Crutchfield writes that "[...] a data set is complex if it is the composite of many symmetries." [5]. A hierarchical structure possesses exactly the architecture which can generate many symmetries. Such symmetries are, for example, the low- and high-frequencies which are associated with the upper and lower levels of the hierarchy, as described by Simon. Thus, the definition of a complex system as one which has high statistical complexity does overlap if not coincide with the definition of a hierarchy. It is an open question whether any non-hierarchical, non-engineered structure generates data with high statistical complexity.

It is fair to ask how the natural sciences were so successful before they began to take the hierarchical structure of nature into account. Clearly, nature has always been hierarchical. Subatomic particles form atoms, they in turn form molecules, which form organisms, which form living creatures, which form social groups. However, the natural sciences could generate meaningful and predictive theories for higher levels without knowing much about the lower levels. The reason for the ongoing success of the natural sciences was that often a higher level could be explained with

only a coarse picture of the level underneath it. Fluids, for example have been studied long before the existence of molecules was known. An atomic physicist does not need to know about subatomic particles to master her field. Yet understanding the levels above and below is becoming more and more important to compose a theory of one particular level.

Many of the phenomena studied now in the natural and social sciences are beyond the scope of any one discipline. For example, understanding human diseases requires knowledge of the physics of electromagnetism, the chemistry of molecular bonding, the biology of cellular organisms, and the psychology of the human mind. Hence, the importance of complexity indicates that our scientific knowledge has reached a point where we are able to put pieces together that we had previously been gathering separately.

VI. CONCLUSION

The right measure of complexity must be computable and not be maximal for randomness. Crutchfield's statistical complexity has these features and as such is illustrative of a good measure of complexity. Inferring the minimal ϵ -machine from a data sequence means identifying some form of pattern or structure, statistical or deterministic in the data. It is time to raise an important philosophical issue for complexity science, or, indeed, any scientific theory, namely the question of realism versus instrumentalism. What are the prospects for adopting a realist position toward the patterns represented by ϵ -machines as opposed to the view, say, that the patterns are merely useful tools for predicting the system's behaviour but not real features of the world?

In nature patterns are everywhere: the familiar arrow-head flight formation of geese; bees which assemble honeycomb into an hexagonal tessellation; the spiral formations of sunflowers and galaxies. Cicadas of the genus *magicicada* exhibit the interesting temporal pattern that their life-cycles are prime numbers of years to avoid synchrony with potential predators' lifecycles. The scientific study of naturally occurring patterns requires both a suitable means for formally representing patterns and a method of inferring patterns from data that picks out objective features of the world. The reconstruction of ϵ -machines meets the former challenge by representing patterns via the smallest computational model capable of statistically reproducing them. The latter challenge is an instance of the classic problem of natural kinds articulated by Plato:

... [T]o be able to cut up each kind according to its species along its natural joints, and to try not to splinter any part, as a bad butcher might do.[4, 542]

Patterns are doubly challenging for the scientific realist because not only are physical theories concerning patterns subject to the usual anti-realist arguments (namely the pessimistic meta-induction, the underdetermination argument and so on), but also all examples of patterns, such as the ones given above, are dependent on the existence of some underlying matter to exhibit them so there is always the objection that all that really exists is the matter not the pattern. A related but stronger motivation for antirealism about the patterns studied by the complexity sciences is reductionism. The reductionist begins with the plausible premise that the causal power of a high level phenomenon, say a brick, is nothing over and above the sum of its micro constituents' causal powers. Next she argues that only those objects whose causal interactions are necessary to produce the observed phenomena should be regarded as the real ones. The conclusion is that bricks and windows are superfluous causal terms in our explanation of observed phenomena. In other words, all phenomena arise from the low level knocking about of objects, and these are the only objects one needs to take to exist. An example of this strategy is eliminative materialism in the philosophy of mind, which seeks to replace folk-psychological discourse about minds, beliefs and emotions etc. with neuroscientific discourse about neural activity within the central nervous system. Contemporary metaphysics has also employed causal exclusion arguments against higher level phenomena, taking atomic collisions and suchlike as paradigm causal mediators and leading some authors [31] to dismiss talk of baseballs breaking windows, on the grounds that what *really* does the interacting are the brick's and the window's microstructures. Higher level phenomena such as bricks and windows are redundant causal terms and we are urged to view them as merely predictively useful abstractions, that is, we are urged to be instrumentalists about bricks and windows.

So how might we justify a realist attitude towards patterns? The most promising approach pioneered by Daniel Dennett in his paper "Real Patterns" [9] and endorsed in one form or another by David Wallace [46] and James Ladyman and Don Ross (2007) [25] is to appeal to the remarkable amount of computational effort saved at the expense of very little predictive power when adopting a theory of patterns. Consider Wallace's example: suppose we are interested in predicting the movements of a tiger when hunting its prey; one approach would be to model the entire tiger, prey, environment system on the sub-atomic level and study its dynamics according the laws of quantum mechanics, since quantum mechanics is by far the most accurately verified physical theory ever devised we should be confident our predictions will be highly accurate. This strategy is obviously completely impractical because there is no realistic hope of obtaining an accurate quantum mechanical description of the entire system let alone computing its time evolution. Ascending to a

higher level, we may view the tiger and its prey as biological systems and attempt to predict their behaviour through our understanding of biochemistry and neuroscience. This approach is vastly simpler than the previous approach but still way beyond our practical reach. Notice also that some accuracy will have been sacrificed. Highly improbable quantum fluctuations simply do not occur in the biological picture but this loss of accuracy is overwhelmingly insignificant in comparison to the reduction of the computational complexity of modelling the situation. Ascending to the even higher level of animal-psychology we have again compromised on predictive accuracy since the fine details of the individuals will not be modeled. However, as before this small loss in predictive power is more than compensated for by now having a tractable means for accurately predicting how the situation will play out. To quote Wallace's slogan:

A tiger is any pattern which behaves as a tiger.[46, 7]

This sounds like realism: tigers exist and they are anything that behaves like a tiger. However the discussion so far has been distinctly instrumentalist in tone reporting how ascending to higher levels serves to increase our practical predictive capacity.

To sum up:

1. We are pursuing a realist account of patterns in the physical world.
2. It is noted that some patterns allow for an enormous simplification of our model of reality while trading off relatively little predictive power.
3. Identifying patterns via their predictive utility is suggestive of an instrumentalist and anti-realist approach. How can a pattern's predictive utility be given ontological weight?

The beginnings of a theory to answer the latter question were given by Dennett [9] and refined by Ross [37] and Ladyman [25] resulting in an ontological account of patterns called *Rainforest Realism*. The main thesis is that since computation is a physical process there is a determinate matter of fact about whether a pattern is predictively useful, namely, if it is possible to build a computer to accurately simulate the phenomena in question by means of said pattern, and if doing so is much more computationally efficient than operating at a lower level and ignoring the pattern. Since computation is a physical process, it is the laws of physics which determine whether such and such a computation can occur and hence whether a given pattern is real. Crutchfield and Shalizi prove in section 5 of [39] that ϵ -machines are the unique, optimal predictors of the phenomena they are meant to model.

Our first theorem showed that the causal states are maximally prescient; our second, that they are the simplest way of representing the pattern of maximum strength; our third theorem, that they are unique in having this double optimality. Further results showed that ϵ -machines are the least stochastic way of capturing maximum-strength patterns and emphasized the need to employ the efficacious but hidden states of the process, rather than just its gross observables, such as sequence blocks.[39, 853]

The claim is that the patterns inferred and represented by ϵ -machines are the simplest, most accurate means for predicting the behaviour of the system they describe. According to rainforest realism it must be possible to construct a computer capable of simulating the phenomena under investigation by means of some pattern, in this case an ϵ -machine, in order for the pattern to be real. We know that in the vast majority of cases this will be possible. ϵ -machines are typically members of a simple class of computers such as finite state automata and we know from daily experience that we are adept at building computing devices capable of simulating these, namely, desktop computers. So long as we have enough memory, energy and time to run our computer the pattern described by the simulation is an objective feature of the simulated phenomena. The process of inferring ϵ -machines does indeed identify real patterns. Rainforest realism, despite its emphasis on the predictive utility of patterns, avoids instrumentalism through a criterion of pattern reality – the possibility or otherwise of such and such a computing device – whose truth is determined by objective facts about the physical world.

-
- [1] W. Brian Arthur. Complexity and the economy. *Science*, 284:107–109, April 1999.
 - [2] Remo Badii and Antonio Politi. *Complexity: Hierarchical Structures and Scaling in Physics*. Cambridge University Press, 1999.
 - [3] Charles. H. Bennett. Logical depth and physical complexity. In R Herken, editor, *The Universal Turing Machine, A Half-Century Survey*, pages 227–257. Oxford University Press, Oxford, 1988.
 - [4] John M. Cooper, editor. *Phaedrus in Plato Complete Works*. Hackett, 1997.
 - [5] J. P. Crutchfield and K. Young. Computation at the onset of chaos. *Entropy, Complexity and the Physics of Information, SFI Studies in the Sciences of Complexity*, VIII:223–269, 1990.
 - [6] James P. Crutchfield. The calculi of emergence: computation, dynamics and induction. *Physica D: Nonlinear Phenomena*, 75(1-3):11–54, 1994.
 - [7] James P. Crutchfield and Cosma Rohilla Shalizi. Thermodynamic depth of causal states: objective complexity via minimal representations. *Physical Review E*, 59(1):275–283, 1999.

- [8] James P. Crutchfield and Karl Young. Inferring statistical complexity. *Phys. Rev. Lett.*, 63(2):105–108, Jul 1989.
- [9] Daniel Dennett. Real patterns. *The Journal of Philosophy*, 88(1):27–51, 1991.
- [10] Editorial. No man is an island. *Nature Physics*, 5:1, 2009.
- [11] Richard Feynman. *Feynman Lectures on Computation*. Westview Press, 2000.
- [12] Richard Foote. Mathematics and complex systems. *Science*, 318:410–412, October 2007.
- [13] Murray Gell-Mann. What is complexity. *Complexity*, 1:1, 1995.
- [14] Murray Gell-Mann and Seth Lloyd. Information measures, effective complexity, and total information. *Complexity*, 2(1):44–52, 1996.
- [15] Murray Gell-Mann and Seth Lloyd. Effective complexity. In Murray Gell-Mann and Constantino Tsallis, editors, *Nonextensive Entropy – Interdisciplinary Applications*. The Santa Fe Institute, OUP USA, 2004.
- [16] Nigel Goldenfeld and Leo P. Kadanoff. Simple lessons from complexity. *Science*, 284:87–89, April 1999.
- [17] Peter Grassberger. Problems in quantifying self-generated complexity. *Helvetica Physica Acta*, 62:489–511, 1989.
- [18] John H. Holland. *Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control and Artificial Intelligence*. MIT Press, new edition edition, July 1992.
- [19] Karl Young James P. Crutchfield. Computation at the onset of chaos. In Wojciech H. Zurek, editor, *Complexity, Entropy and the Physics of Information*, volume 8, pages 223–269. The Santa Fe Institute, Westview Press, 1988.
- [20] E. T. Jaynes. Information theory and statistical mechanics. *The Physical Review*, 106(4):620–630, May 1957.
- [21] E. T. Jaynes. Information theory and statistical mechanics. ii. *The Physical Review*, 108(2):171–190, October 1957.
- [22] A.N. Kolmogorov. Three approaches to the quantitative definition of information. *Problems of Information Transmission*, 1:1–17, 1965.
- [23] A.N. Kolmogorov. Logical basis for information theory and probability theory. *Information Theory, IEEE Transactions on*, 14(5):662–664, 1968.
- [24] A.N. Kolmogorov. Combinatorial foundations of information theory and the calculus of probabilities. *Russian Mathematical Surveys*, 38(4):29–40, 1983.
- [25] James Ladyman, Don Ross, David Spurrett, and John Collier. *Everything Must Go: Metaphysics Naturalized*. Oxford University Press, 2007.
- [26] M. Li and P.M.B Vitányi. *An Introduction to Kolmogorov Complexity and Its Applications*. Graduate Texts in Computer Science. Springer Verlag, 2nd edition, 1997.
- [27] S. Lloyd. Measures of complexity: a nonexhaustive list. *Control Systems Magazine, IEEE*, 21:7–8, 2001.
- [28] Seth Lloyd and Heinz Pagels. Complexity as thermodynamic depth. *Annals of Physics*, 188:186–213,

November 1988.

- [29] Seth Lloyd and Heinz Pagels. Complexity as thermodynamic depth. *Annals of Physics*, 188(1):186–213, 1988.
- [30] Karl Mainzer. *Thinking in Complexity: The Complex Dynamics of Matter, Mind and Mankind*. Springer, 1994.
- [31] Trenton Merricks. *Objects and Persons*. Oxford University Press, 2001.
- [32] Edgar Morin and J. L. Roland Belanger. *Method: Towards a Study of Humankind : The Nature of Nature: 001*. Peter Lang Pub Inc, November 1992.
- [33] William Paley. *Natural Theology*. Oxford University Press, 2006.
- [34] Julia K. Parrish and Leah Edelstein-Keshet. Complexity, pattern, and evolutionary trade-offs in animal aggregation. *Science*, 284:99–101, April 1999.
- [35] P.W.Anderson. More is different: Broken symmetry and the nature of the hierarchical structure of science. *Science*, 177, 1972.
- [36] D. Rind. Complexity and climate. *Science*, 284:105–107, April 1999.
- [37] Don Ross. Rainforrest realism: A dennettian theory of existence. In Don Ross, editor, *Dennett's Philosophy: A Comprehensive Assessment*, chapter 8, pages 147–168. MIT Press, 2000.
- [38] R.S.MacKay. Nonlinearity in complexity science. 2008.
- [39] Cosma Rohilla Shalizi and James P. Crutchfield. Computational mechanics: Pattern and prediction, structure and simplicity. *Journal of Statistical Physics*, 104(3):817–879, 2001.
- [40] Cosma Rohilla Shalizi and Cristopher Moore. What is a macrostate? subjective observations and objective dynamics. *cond-mat/0303625*, 2003.
- [41] C. E. Shannon. A mathematical theory of communication. *Bell System Technical Journal*, 27:379–423;623–656, 1948.
- [42] Herbert A. Simon. The architecture of complexity. *Proceedings of the American Philosophical Society*, 106(6):467–482, 1962.
- [43] Christopher G. Timpson. The grammar of teleportation. *British Journal of Philosophy of Science*, 57(3):587–621, 2006.
- [44] Alan M. Turing. On computable numbers, with an application to the Entscheidungsproblem. *Proceedings of the London Mathematical Society*, 2(42):230–265, 1936.
- [45] R. Wackerbauer, A. Witt, H. Atmanspacher, J. Kurths, and H. Scheingraber. A comparative classification of complexity measures. *Chaos, Solitons & Fractals*, 4(1):133–173, 1994.
- [46] David Wallace. Everett and structure. *Studies In History and Philosophy of Modern Physics*, 34(1):87–105, 2003.
- [47] Gezhi Weng, Upinder S. Bhalla, and Ravi Iyengar. Complexity in biological signaling systems. *Science*, 284:92–96, April 1999.
- [48] B. T. Werner. Complexity in natural landform patterns. *Science*, 284:102–104, April 1999.
- [49] George M. Whitesides and Rustem F. Ismagilov. Complexity in chemistry. *Science*, 284:89–92, April

1999.

- [50] Wojciech H. Zurek. Algorithmic information content, church-turing thesis, physical entropy and maxwell's demon. In Wojciech H. Zurek, editor, *Complexity, Entropy and the Physics of Information*, volume 8, pages 73–89. Westview Press, 1990.
- [51] The relationship between macrostates and microstates is key to the complex sciences because very often what is interesting about the system is the way that a stable causal structure arises that can be described at a higher level than that of the properties of the parts (see Sec. IIE on emergence below).
- [52] One anonymous referee claimed that it is not possible to define chaos but on the contrary unlike complexity chaos can readily be defined as a system that exhibit so-called strong mixing. Moreover, recently, Charlotte Werndl has shown that there is a kind of unpredictability unique to chaos (2008).
- [53] Note that we are not here talking about whether the system that produces the data is deterministic or not. Of course, the Shannon entropy of a probability distribution is insensitive to whether that probability distribution was produced by deterministic or an indeterministic system. Our point is just that a good measure of complexity will not be maximal for random data strings.
- [54] Note that Bennett's definition is actually indexed to a 'significance level'. We omit mention of this since it is inessential to Bennett's basic idea and to our reason for rejecting his measure below.
- [55] ϵ refers to how finely we monitor the system's behaviour. In general it is a fine grained partition of the system's phase space.
- [56] We are grateful to an anonymous reviewer whose criticisms based on the example of the climate forced us to clarify these points.