Synthese
DOI 10.1007/s11229-010-9721-7

Determinism, predictability and open-ended evolution: lessons from computational emergence

Philippe Huneman

Received: 21 December 2009 / Accepted: 2 February 2010

Abstract Among many properties distinguishing emergence, such as novelty, irreducibility and unpredictability, computational accounts of emergence in terms of computational incompressibility aim first at making sense of such unpredictability. Those accounts prove to be more objective than usual accounts in terms of levels of mereology, which often face objections of being too epistemic. The present paper defends computational accounts against some objections, and develops what such notions bring to the usual idea of unpredictability. I distinguish the objective unpredictability, compatible with determinism and entailed by emergence, and various possibilities of predictability at emergent levels. This makes sense of practices common in complex systems studies that forge qualitative predictions on the basis of comparisons of simulations with multiple values of parameters. I consider robustness analysis as a way to ensure the ontological character of computational emergence. Finally, I focus on the property of novelty, as it is displayed by biological evolution, and ask whether computer simulations of evolution can produce the same kind of emergence as open-ended the evolution attested in Phanerozoic records.

Keywords

1 Introduction

Recently, an account of emergence in terms of computational compressibility in has been proposed by authors such as Bedau (1997, 2008), Holland (1998) and Crutchfield (1994, 2002) . This account, centered on computer science rather than on a view of a multilayered ontology for the physical world, had the advantage of being objective in

P. Huneman (✉)
IHPST (CNRS/Université Paris I Sorbonne), 13 rue du Four, 75006 Paris, France
e-mail: huneman@wanadoo.fr

the sense that it does not make the property of being emergent (or not) dependent upon the cognitive capacities or explanatory strategies of the human subject. Two related criticisms could, however, be raised against it: The first one concerns its involvement within the practice of computer simulations; the second one is the worry that it might make too many things emergent. In a word, one might concede that computational emergence makes sense of phenomena whose simulation involves emergent processes in the sense that they will not allow predictions, but cannot account for other characteristics usually related to emergence such as novelty, order or downward causation. In this paper, I will start by defending the account against some objections by showing how the account of unpredictability in models can entail accounts of the other dimensions of emergence. Then I will discuss what such an idea of emergence involves concerning predictability in general, and its relation with determinism. I will show that it allows different kinds of predictability, and therefore makes sense of the practice of qualitative predictions in the study of complex systems. The last section will consider the general issue of predictions in evolutionary theory, and therefore revisit the debates about open-ended evolution and classifying the kinds of evolution in biology and in artificial life (ALife).

2 Emergence: the concept and its instances

Among characters that entitle something—a property, a state, a process, or whatever—to be emergent are novelty, irreducibility, unpredictability and downward causation (e.g., Silberstein 2002; Chalmers 2006; Newman 1996). Authors frequently insist on "unpredictability" because the core idea of emergence involves that some phenomenon or property of a system, though not transcendent, cannot, in some form or another, be "derived" from the laws of nature and description of the "elementary" properties of this system. Specifications of this idea include characterizing what 'derivation' and 'elementary' mean; "prediction" is arguably a kind of derivation of a state of affairs, and "elementary properties" may often mean the properties of parts of a system. Numerous accounts of emergence have been centered on irreducibility, trying to capture the idea that some property of one level of a system is irreducible to properties of a lower level, or that properties of a whole system cannot be reduced to properties of parts (e.g., Wimsatt 2007; Wilson ~~forthcoming~~). Yet many of those accounts face the critique that they grasp something which obviously depends on the explanatory framework of the cognitive subjects, which for instance decide the relevant parsing of the system into levels. This gave rise to the worry that emergence is an epistemological property, and in no ways ontological (Silberstein 2002). The story of British emergentism, whose paradigm of emergent properties was chemical bonds, and which has been refuted when the progress of science explained this property away through quantum physics, is typical in this regard (McLaughlin 1992).

1997, H 2009

Recently, prospects of having a non-epistemic account of emergence arose in the wake of computer science. Genetic algorithms (Holland 1998), cellular automata (Bedau 2003; Crutchfield and Hanson 1993; Crutchfield 2002) and agent-based modelling (Epstein 1999) provide a field in which to discuss criteria of emergence that are rigorous and may have the rigidity of mathematical properties. Such accounts can

Journal: 11229-SYNT Article No.: 9721 MS Code: SYNTHB-4 TYPESET DISK LE CP Disp.:2010/2/23 Pages: 21 Layout: Small-X

be framed under the name of computational emergence (Humphreys and Huneman 2008). The basic idea is that a computational simulation process is emergent if and only if its result cannot be reached except by running the whole simulation; in other words, one cannot compress the simulation.[1] The interest of such a criterion is that emergence in this perspective is proved to be objective, in the sense that it does not rely on our epistemic capacities. In effect, the differences between computational classes do not depend upon the abilities of the mind exercising the computation; therefore they are objective and cannot be dissipated through the improvement of our mathematical skills (Huneman 2008b). Note that this account of emergence first draws upon the property of unpredictability, which is addressed by logical demonstrations of the computational difference between classes of automata (Buss et al. 1992).

However, one could argue that such a concept of emergence cannot be the whole story, since it concerns only simulations, and does not touch upon actual causes in the world (Corning 2002).[2] However, the point of providing a computational view of emergence is precisely about figuring out a *concept* of emergence that cannot be identified, in last analysis, to an epistemic weakness. Another question is *whether* this concept is *instantiated* in the world, and this entails the issue of the relationships between simulation, or computational models, and processes in nature. But still, it might be that this very concept of emergence entails something about causes (as I argued in Huneman 2008a),[3] whether or not it is instantiated in the world. This means that, if something in the world falls under this concept, it will display some specific causal properties. (Sect. 4 below addresses the issue of how we know whether some real system or process actually instantiates the concept).

Another worry is that even if unpredictability were accounted for in this way, it would not be equal to emergence, since many unpredictable features are not at all interesting. Among those features, for instance, we could find random sequences, and such a fact is at odds with our intuitive understanding of emergence. I tried to meet this objection in various ways (Huneman 2008a,b), and will summarize this argument here.

First, consider cellular automata (CA) like the Game of Life. Some of these display computational incompressibility, so let's take an incompressible instance of a Game of Life, and keep it running. The whole point of incompressibility is that, while there is a deterministic rule so that one can go from a set of cells $C(i-j, i+j)^n$ to the state of cell Ci^{n+1}, there is no shortcut to go from $C(1, M)^n$ to $C(1, M)^{n+1}$ (M being

[1] Proving that some simulation cannot be compressed is a difficult issue, and indeed Chaitin's results in complexity theory state that this problem is unsolvable. But this concerns our grasp on compressibility, which does not affect the question of whether incompressibility provides a satisfying criterion of emergence. If this is the case, then the trouble of establishing incompressibility would simply mean that we can have no certitude that a process is indeed emergent, only very probable judgments.

[2] The objection was addressed to a Holland-style account of emergence: "Consider Holland's chess analogy. Rules or laws have no causal efficacy; they do not in fact "generate anything". They serve merely to describe regularities and consistent relationships in nature. These patterns may be very illuminating and important, but the underlying *causal agencies* must be separately specified (though often they are not)" (Corning 2002, p. 26).

[3] The argument is basically that within a given simulation counterfactual dependencies hold between sets of states of cells at various steps, which therefore means that causal relationships (*sensu* Lewis 1973) are at stake within it.

Journal: 11229-SYNT Article No.: 9721 MS Code: SYNTHE-4 TYPESET DISK LE CP Disp.:2010/2/23 Pages: 21 Layout: Small-X

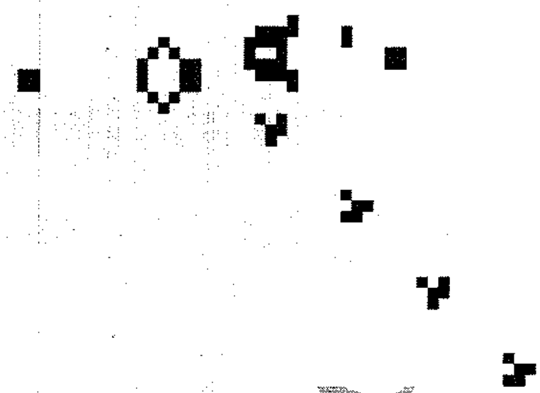

Fig. 1 Game of Life, "gospel glider gun" throwing gliders. The moving sets of cells called gliders takes successive positions along the diagonal; the position of a glider at step n counterfactually depends on its position at step n-m, at least for a period of the simulation

the number of cells) or from n to the whole state of $C(1, M)^n$. However, among the cellular automata, it might well be that some display gliders, glider guns or all this fauna of digital creatures familiar to computer scientists. This means that, though still incompressible, there is another way to describe the Game of Life which relies on counterfactual dependencies between sets of ~~sets~~ of cells at various steps.[4] A glider can indeed be characterized as a set of ~~sets~~ of cells; gliders have typical behaviors, so that if at step n the glider is at one point, analysis of the cellular automaton may entail that it will be somewhere else a while after (after N steps, let's say). One can then claim that, on the background of the rules of the automaton, the glider would not be where it is at step $n + N$ if it had not been where it was at step n (Fig. 1). This is a counterfactual dependency, and it illustrates the fact that even in cellular automata showing computational incompressibility, some regularities of this type can be produced.

states
states

The class of automata that, in addition to being computationally incompressible processes ("weakly emergent" *sensu* Bedau), display such counterfactual dependences between sets of ~~sets~~ of cells, match our expectations concerning the concept of emergence (namely, both unpredictability and a kind of "order")—and I thereby call them "robustly emergent." The global behavior of such an automaton, far from being random, displays global regularities that were not implemented in its first designing or description. This meets our intuitive idea that emergent features should exhibit a kind of order not reducible to properties of the parts.

states

Bedau (2008), defining weak emergence, insists that it is both dependent (on the properties of the elementary building blocks) and autonomous. The irreducibility of entities involved in novel counterfactual dependency exactly fits his definition. Here, "irreducible" does not mean that it is something ontologically different from the basis of the entities (in this case, the cells of the CA). In fact the gliders, and all the global-level entities emerging within the automaton used to track counterfactual dependencies within it, are not superadded. Because they are produced by the building blocks, they are the deterministic effects of the building blocks' movement (see Epstein 1999 for the same demonstration about agent-based models). However, they are irreducible in the sense that deriving the exact position and movement of an entity like a glider

[4] A more precise characterization, in terms of counterfactual dependencies between internal properties of the CA defined by disjunctions of sets of states of cells, is given in Huneman (2008a).

Journal: 11229-SYNT Article No.: 9721 MS Code: SYNTHB-4 TYPESET DISK LE CP Disp.:2010/2/23 Pages: 21 Layout: Small-X

from the rules and position of the initial cells is precisely incompressible; there is no shortcut (because the whole emergence of gliders was incompressible). In this sense emergence does not oppose reductionism; ~~causal~~ *processes* take place in the H the building blocks or the cells, according to their rules, and they completely determine the behavior of the automaton. However, *a more explanatory set* of counterfactual dependencies—so causal relationships, in a weak sense of causation—is produced through the emergent, incomputable process.[5] Thereby, notwithstanding the computational incompressibility proper to the simulation, we can still look for macrostates that could be predictors of what follows, and then parameters that would define a class of those macrostates – in the sense of maximized "efficiency of prediction" defined by Shalizi (Shalizi and Moore 2003; Shalizi and Crutchfield 2008) in the context of his conception of macrostates. So the claim here is that robust computational emergence includes some *specific causal patterns*, and is not only formal; Sect. 4 will examine under which conditions those causal patterns exhibited by computer simulations match real causation in nature.

I will now turn to the consequences of such a complete concept of emergence ital (computational irreducibility and counterfactual dependencies between global-level entities) for the idea of predictability.

3 Emergence, predictability and determinism

One interestesting point about this view of robust emergence is that it disentangles the two concepts of determinism and predictability, in a formal manner. Cellular automata are purely deterministic; yet, some processes might be incompressible, which means that they yield unpredictable results.

If we cash out the idea of predictability, it intuitively means the following. Imagine a system, even a deterministic one, with a given initial state (i), i.e., an initial value of the descriptive state variables or initial position in the state space, and a small range (d) of values around those initial values. This system allows predictions if the final values it reaches, starting from all different initial values in $[i - d; i + d]$, are in a range $f(d)$ which is not too much larger than the range (d) of initial values. If not, it means that the margin of error (represented here by d) of our measurements of those initial values will not ensure that the final result yielded by calculating the final state is in a same or analogous margin of error, so there will be no possible prediction. The baker's transformation (Fig. 2) provides a famous example of a system which is clearly deterministic, but for which two arbitrarily close points in the initial state will have at some step n a distance which will be very large, but they will still be able to get very close at some remote step $n + m$. Any margin of error around a point A initially will /a render prediction of the final state of A impossible. Interestingly, although predicting H the system is related to our epistemic capacities, this condition for being predictable or not is a

[5] One must note that this account of emergence in terms of a new explanatory level that is not computationally (or analytically) involved in the explanatory practices for the building blocks level comes close to the way Batterman (2002) characterizes emergence in terms of gaps between explanatory regimes (which is the way he solves the riddle of downward causation).

Journal: 11229-SYNT Article No.: 9721 MS Code: SYNTHB-4 TYPESET DISK LE CP Disp.:2010/2/23 Pages: 21 Layout: Small-X

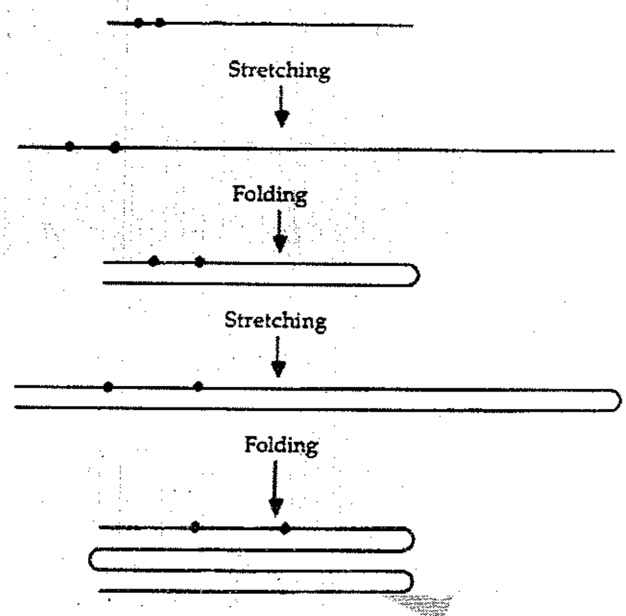

Fig. 2 The baker's transformation—one of the simplest cases of unpredictability in deterministic systems. At each step the line is stretched to twice its length and then folded. Two points randomly chosen, no matter ~~is~~ close they are, will at an arbitrary step in the future be arbitrary far from each other. This renders any prediction as to the fate of one initial point impossible (given that an interval around it, as small as it can be, still yields future states that lie in an arbitrarily extended interval)

How

property of the system, so it is *objective*. Famously it is expressed by the notion of sensitivity to initial conditions.

In the case of computational emergence, a logical result shows that when in some set of logical automata the initial state is changed, there is no polynomial function of the predicted result that would yield the correct prediction for the subsequent new final state (Buss et al. 1999). On the contrary, it is at least an exponential function of this result, which means that no computation can correct the prediction. So those deterministic systems like CA that display emergent behavior are not predictable, although they are deterministic. In stochastic computer simulations like some genetic algorithms, the same criterion of emergence can be applied; this means that the property of unpredictability is orthogonal to the difference between determinism and stochasticity, while it is also an objective property of systems.

Now, if we remember that computational emergence still allows discriminating a class of robustly emergent processes, what does it mean for predictability? The issue is all the more difficult since, though genetic algorithms may sometimes be emergent in this sense, they are somehow predictable because they are often designed to build a predicted result. In this case, the distinction is straightforward: Though the global result is predictable, the pathway toward it is not, since the way the program reaches its optimum is precisely incompressible.

What *in general* corresponds to this notion of pathway unpredictability? Considering the counterfactual dependencies (such as between the gliders) actually yields some predictions in emergent, hence unpredictable, processes. So *prima facie*, if you want to track down the trajectory of one cell in a CA satisfying the incompressibility criterion (for more details on such criteria see Hovda 2008), it will quickly become computationally intractable and thus impossible. The very definition of unpredictability is met here: Changing slightly the initial state of a cell's neighbor from i to i' will be such that no computable function of the final state f_i will give you the modified final state $f_{i'}$ of this cell. This view can easily been generalized: In a robustly emergent process,

Journal: 11229-SYNT Article No.: 9721 MS Code: SYNTHB-4 TYPESET DISK LE CP Disp.:2010/2/23 Pages: 21 Layout: Small-X

the trajectory of one point, or one micro-element of the system, is not predictable. However, as I show here, some predictions—here called macro-predictions—can be reached at the level of global entities whose counterfactual dependencies emerge in the course of the process. Israeli and Goldenfeld (2004) made the point that all the 2-D CA (rules 1–256) could be coarse-grained into other CA. You just have to take sets of cells x_j as one cell X_j, and project the rule that changes x_i according to the state of their neighborhood onto a rule changing X_j according to the state of its neighbours $X_{j'}$. Many computationally reducible CA appear to be reducible at a coarse-grained level. For example, the CA with rule 146 has a coarse-grained version with rule 128, which is of class II, so computationally reducible.

One can roughly distinguish between agent-based models (ABM), or in general computer simulations, which embody equations of a theory explaining a phenomenon (a), and ABM or simulations that just take as rules some plausible regularity, observed at a large scale (b). A famous example of the latter category is of course Reynolds's simulation of the movement of flocks of birds or schools of fish (1987). The generative rules are hypothesized on the basis of what is observed (birds, or "boids" in the simulation, have three simple rules, like fly in the direction of the mean direction of the others, do not go too far from the immediate neighbors, etc.). Another example of the latter category (b) for elaborating macro-predictions comes from Epstein (2002), in his "generative social science" project of designing agent-based models for social phenomena. He created ~~agent-based models~~ ABM for rebellion and social uprisings, whose agents are individuals, and parameters are the level of legitimacy of the government and the level of oppression. He varies the parameters, and the behavior of the agents obeys just a few rules, like the intuitive rule that agents will not easily rebel if nobody else does, but likely will if many of their neighbors are uprising. The results are striking to various degrees, depending on the levels of the parameters and above all the kind of change those parameters undergo; all simulations display a few characteristic kinds of patterns. For example, when the legitimacy of governments drops, what is crucial for the likelihood of uprising is the speed of this fall, not the width. So ABM here allows one to predict likelihood and patterns of uprising on the basis of the values and changes of two global parameters; the ranges of values and types of behaviors are indeed predictors of the global uprising behavior in the set of agents. Another simulation in the field of social sciences shows the same possibility of macro-predictions of emerging behaviors (see Tassier's (2004) study on fads). The outcomes of simulations, here, might display some emerging general behaviors of agents—namely, fads. Then, according to the ranges of variables, we will witness an extension of this pattern of behaviors or, more interestingly, some life cycles for those collective behaviors; those patterns adequately represent the empirically attested phenomena proper to fashion. The resulting patterns can be classified according to the following categories: "If agents imitate each other in order to express an affiliation, there are at least three phenomena that may develop: First, all agents in the population may coordinate on one alternative. Second, another type of stable behavior may develop; agents may separate themselves into stable groups with each group acting in one particular way, or buying a particular good. Third, stable or unstable cycles may develop. These cycles are examples of what I am calling fads or fashion cycles" (Tassier 2004, p. 51). The simulation uses models of networks with connected agents,

Journal: 11229-SYNT Article No.: 9721 MS Code: SYNTHB-4 TYPESET DISK LE CP Disp.:2010/2/23 Pages: 21 Layout: Small-X

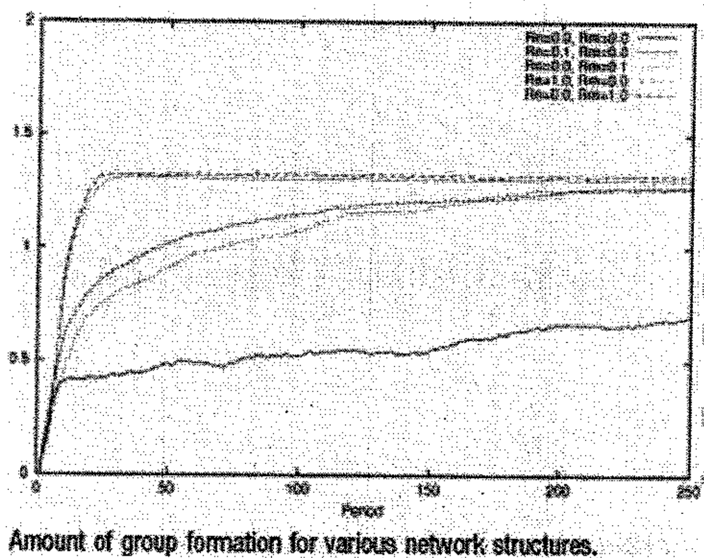

Fig. 3 Group formation according to probabilities Rm and Rn in Tassier's models of fads (2004). Credits Tassier, Wiley and Son

like in Watts (2002), Watts and Strogatz (1998), and Barabasi and Albert (1999). One of the important parameters for the simulation is the "preference parameters for having agents of similar idiosyncratic type and high attraction" in the group to which the focal agent belongs; the other parameters are the probability of having a connection with some individual not connected to the neighbors, and the probability of having a quick encounter with individuals of other sub-networks "such as people one notices when walking down a street that the agent does not know or other random meetings that may occur" (p. 53). Group formations and cycles of fads (defined by the creation of those groups) can be predicted through the simulation from the values of those parameters (see Fig. 3).

The simulations exemplified by the above studies achieve the following: Although trajectories of some agents are not predictable (and this is entailed by computational emergence), one can run several rounds of simulations and check whether ranges of values of the state parameters (legitimacy, level of oppression, etc.) will always yield the same pattern of behavior.[6] The "pattern" here is exactly what I alluded to with the example of the gliders: There are dependencies between sets of states of agents, even if those sets are in each simulation realized variously by the agents (e.g., uprisings are not always led by the same agents, in the same time, with the same length, in the same places, etc.). However, the counterfactual dependencies between range of values, uprisings, their frequency and their succession, are attested by the computer simulations. The fact that given the initial state of the model—positions of the agents, values of state variables—one cannot predict the fate of an agent does not preclude a prediction of general patterns of uprising, based of the counterfactual dependencies between those patterns (here noted m-regularities), on the one hand, and the range of values of the initial parameters on the other hand. In those examples, what is at

[6] On the notion of patterns, see Humphreys (2008), in the framework of cellular automata research.

Journal: 11229-SYNT Article No.: 9721 MS Code: SYNTHB-4 TYPESET DISK LE CP Disp.:2010/2/23 Pages: 21 Layout: Small-X

stake is the distinction of several levels of prediction, in processes that are "robustly emergent" (in the sense of computational emergence as defined above).

More precisely: Robust emergence implies the possibility of predictions based on the counterfactual dependencies between global level entities (noted here m-predictions). What appears in the present examples is a sort of second-degree macro prediction (here M-prediction), which holds between sets of parameters and sets of global-level entities (which can define some patterns). Now, what is considered is a set of simulations, and the M-predictions relate on the one hand the initial values of the initial states (given the rules of the simulation), and on the other hand those regularities emerging in each simulation. In M-predictions, we are concerned with what is generic in multiple runs of the same simulation with different initial distributions, and different values for the parameters (e.g., the levels of oppression and the trust towards government, in Epstein's models of uprising). Of course, those M-predictions somehow generalize the m-predictions that are possible in a given simulation on the basis of those emergent m-regularities (in the sense of counterfactual dependencies instantiated by the ones that tie the gliders together, or which link several successive global steps of an uprising process in Epstein's model).

18

Let us explicate this distinction between levels of prediction. In several simulations (various initial distributions, and possibly various values of initial parameters, as in the case of agent-based modeling), different patterns (with different sizes and positions) emerge. The m-regularities proper to the patterns in a token simulation allow m-predictions internal to this token simulation. Their behaviors (in the sense of m-regularities) are likely to be clustered into *more general classes*. (e.g., in runs of a Game of Life with various initial distributions we will have "gliders with glider guns," "gliders with no glider guns," etc.). Those general classes in turn can then be related to the *ranges* of the variables (about the initial distribution, or about the parameters for the rules) that ~~determined~~ them. This range is determined by all the values of variables for initial distribution (and/or all the values of parameters for the rules of the simulation) which generate m-regularities likely to be clustered in the same class (e.g., "late uprising," "uprising resolved," "no uprising," etc., in Epstein's models of social rebellion).[7]

condition
A specific

Computational emergence entails that changing the value of initial states of cells, or agents, precludes any prediction of the final states of the agents on the basis of the earlier knowledge of the initial states—micro-unpredictability; but this does not logically entail that concerning the m-regularities those changes will still display such sensitivity to initial conditions. So the unpredictability ontologically proper to those systems—in terms of unpredictability of the trajectories of agents—is however superseded, first by

[7] One can refer to the notion of attractors in order to understand what is at stake here; basically attractors are those configurations allowing all varieties of a same system (differing by the values of their parameters) to revolve around them in the phase space and display a similar sort of qualitative pattern of behavior. Exploring *many* simulations therefore allows one to identify the counterfactual dependencies between ranges of values of the parameters and possible attractors, so that the shifts of attractors ("bifurcations") can be made manifest and related to the changes of values of parameters upon which they depend. The relations between attractors and ranges of parameters defining classes of systems exactly correspond to M-predictions in our case. (Newman 1996 develops an attempt to understand emergence in the light of the notion of strange attractors, usual in chaos theory. I sketch here the reverse conceptual pathway, although both theorizations are related and compatible.)

Journal: 11229-SYNT Article No.: 9721 MS Code: SYNTHB-4 TYPESET DISK LE CP Disp.:2010/2/23 Pages: 21 Layout: Small-X

the m-predictions yielded by the m-regularities manifested by global-level entities in one simulation, and second by the M-predictions yielded by the counterfactual dependencies attested between sets of m-regularities and ranges of parameters in classes of varied simulations. For example, the emergence of local norms is a m-regularity involving evolving sets of agents, and is regularly correlated to ranges of values of the parameters of the simulations in Burke and colleagues' (2006) study of the dynamics of social norms.

This explains why the study of complex systems, whose state equations are often analytically intractable, can be done through the design of various simulations and their comparisons. Such comparisons indeed, notwithstanding the unpredictability proper to those systems, will allow predictions in the sense of connections between ranges of values of the parameters or variables, and qualitative patterns of behaviors of the systems. Although such study is experimental, i.e., one has to run many simulations and try to explore wide ranges of values of parameters, the proper unpredictability of the system as emergent does not ruin it, in the sense that it would render absolutely meaningless the inference towards general classes of behaviors.[8]

4 Robustness analysis, relevant causal factors and the instantiations of the formal computational concept of emergence

The last worry about computational emergence was that the concept fits only simulations but is not instantiated in the real world, or that we cannot know when it is indeed instantiated. More precisely, one could argue: This natural process is modeled by a computer simulation that displays emergent processes, and such emergent characteristics are absolutely *objective*; fair enough. Computational emergence is also a *causal* concept (Huneman 2008b; Bedau 2008), so when a computer simulation falls under it, such model tells us something about the causal structure of the modeled system. However, who knows whether this simulation is an accurate one? Couldn't the process be modeled by another, non-emergent simulation, which for the moment we cannot grasp? An answer can be provided by consideration of the methodology of complex systems study through qualitative M-predictions, as exemplified above.

In such studies, identifying the qualitative patterns of behavior is done by varying the values of the variables in the simulations, and then running them. Yet one can also vary the *sets of parameters* used in a simulation. This would probably refine the predictions (for example, suppose that in Epstein's simulation we add a parameter like degree of education). However, it might be that some parameters do *not* change the patterns of behavior.

Robustness analysis is the method of comparing several models with varied values of the variables or parameters used, for a given phenomenon. There are two kinds of robustness analysis, one weaker and one stronger, according to whether you vary the values of the parameters, or modify the number and nature of the parameters

[8] Of course, varying the values in order to find how those changes affect the trajectories of one given agent will be pointless, precisely because of this unpredictability. in an ABM

Journal: 11229-SYNT Article No.: 9721 MS Code: SYNTHB-4 TYPESET DISK LE CP Disp.:2010/2/23 Pages: 21 Layout: Small-X

themselves (Weisberg 2006a,b).[9] When the model yields the same predictions and the same patterns of behavior even if we change some parameters, then it is said to be "robust."[10] In this case, even if we were mistaken in identifying the ~~relevant~~ parameters, our model would still be accurate, and somehow fitting the "causal structure"[11] of the target systems, since nothing in its behavior would change when we substituted other parameters ~~(the relevant ones)~~ for the ones actually used (Levins 1966; Wimsatt 2007; Weisberg 2006a).[12]

An example of this case is provided by the neutral theory of ecology (Hubbell 2001). In modeling relationships of vicariance, distribution and abundance between species in a community (species of the same trophic level, e.g., trees), the theory poses the unrealistic hypothesis that rates of death and birth are the same for all individuals of all species; this is highly implausible and entails a negation of natural selection. Actually the model repeatedly correctly predicts distributions, for example of species of trees in tropical rainforests; abundance distributions of species of trees in the Barrio Alto rainforest have been accurately modeled in this framework (Hubbell 2001). These models are often sophisticated simulations, and it is plausible that in many of them the outcome is *computationally emergent*; moreover such processes display *robust* emergence, because in the end we have patterns appearing and following each other in a regular manner, which makes room for the m-regularities I talked about in Sect. 3. Now, the outcomes of such models often intriguingly match the actual distribution (Chave 2004; Holyoak and Loreau 2006).

Taking natural selection as a parameter defines what ecologists call "niche effects." If we incorporate "niche effects" in those models which successfully represent actual distributions, the fact is that the outcomes will not change much (Bell et al. 2006, p. 1383; Chave 2004, p. 246). So the model of species abundance distribution in the tropical rainforest is robust in the sense that adding niche parameters does not change the outcome; moreover, numerous outcomes of the simulations designed according to neutral theory (given various initial states) match the patterns displayed by real systems. From this viewpoint it is plausible that the neutral dynamics, for which the only parameters are the general parameters of the population (namely the size, the age of reproduction, etc.), and which mainly occurs because of dispersal limitation,[13] is causally responsible for the emergent patterns of distribution (Chave 2004). The species

9 In what follows I will focus mostly on strong robustness analysis.

10 Above, "robust" emergence means computationally emergent process with robust counterfactual dependencies between macro patterns; here "robustness" concerns a set of models. Those meanings are not identical: We can speak of "robust models of robustly emergent processes" when a model satisfying a robustness analysis displays robustly computationally emergent processes. In order to keep with previous uses in both modeling literature and emergence literature, I kept this unfortunate dual use of a single word. t

11 The connection between a robust model and its consistency with the "causal structure" of the system is argued for in Weisberg (2006a,b).

12 The same reasoning also holds of course for robustness analysis concerning the values of the variables (weak robustness analysis).

13 Connectivity of space and number of species yields dispersal limitation, which in turn influences the relative strength of intra- and inter-specific competition, and finally leads to equalizing fitnesses, both inter- and intra-specific (Holyoak and Loreau 2006, p. 1374). This mechanism explains the adequacy of the assumptions of the neutral model. partly

Journal: 11229-SYNT Article No.: 9721 MS Code: SYNTHB-4 TYPESET DISK LE CP Disp.:2010/2/23 Pages: 21 Layout: Small-X

robustness of such neutral models dismisses niche effects as putative causal factors for this outcome.

Analyzing the examples involves two important consequences. First, the initial objection about the purely formal character of emergence can be dismissed. Suppose that Epstein's model of civil uprisings is robust. It means that the system modeled—here collectives of people—is captured in its causal structure by the model. So when there is emergence in this model (of course, all the parameter values may not generate emergence in the above sense of robust emergence), the *reality* itself can be said *to display emergent processes* or emergent properties, since this emergence has been made manifest in a robust model *of its causal structure*. So robustness analysis, here, allows one to rebut an objection based on the fact that someday we might have a better knowledge of the system. Computer simulations of a process which display robust emergence and which have passed the test of robustness analysis can rightly allow us to infer that the real system has emergent features, in a sense which is in no way epistemic, and even if we later learn of other parameters responsible for the behavior of the system which we could insert in a new model. This new model will indeed be similarly emergent, because of the robustness of the family of models considered.

The second insight concerns the relevant causal factors. A robust model does not change its global outcome, which maps onto the actual outcome of the system,[14] when we vary variables or parameters. I argue here that the parameters not causally responsible for the emergent features are the ones that are not making a difference to the robust model's behavior. Suppose that indeed adding or withdrawing the parameter "education" does not change the qualitative behaviors in Epstein's model of uprising (behaviors are the connections between ranges of values and m-regularities, as stated above). This means that nothing in the m-regularities depends upon such parameters. So whereas those parameters involve many effects in the behaviors of the individual "agents," they must not be held as causes of the qualitative patterns of outcome behaviors considered here, namely the trends for uprising. Thus, while robustness analyses of models enables one to infer from emergence in simulations to emergence in natural processes, it also allows one to identify the parameters that will be causally relevant in the production of emergent features in the processes (or at least exclude the ones that are not responsible).

More precisely, the robustness of the model means that it is not sensitive to adding or withdrawing some parameters. Why is this so? An inference to best explanation gives here the answer. What could possibly explain the insensitivity to other parameters? In fact, in reality there are always a very large number of parameters for a given system, but parameters are connected such that the effects of some of their values are compensated for by others, so that in the end they are inefficient regarding the outcome. Hence, the insensitivity to change of parameters, on the side of the model, together with the fact that outcomes of the model map onto real situations, would mean that those connections are captured through the model, which means in turn that it is somehow connected to reality. In this sense, the best explanation for the parameter insensitivity of robust modeling is some connection with the "causal structure" of

[14] There are many ways to define this mapping, but I will not consider this issue here. Suffice to say that the M-predictions from the model are confirmed in many actual instances of the system.

Journal: 11229-SYNT Article No.: 9721 MS Code: SYNTHB-4 TYPESET DISK LE CP Disp.:2010/2/23 Pages: 21 Layout: Small-X

reality.[15] Therefore, the consequence, in the case of a robust model displaying emergent processes, is that this model parses the parameters causally responsible for emergent features, and the ones that are not.

Generally speaking, the philosophical issue here is reading causal factors off scientific models. When models are equations, it is assumed that the variables in the equations are the relevant causal factors in reality.[16] Analogously, in the case of computer simulations designed when equations are out of reach or intractable, reading the causal factors off the model consists in identifying which parameters are constantly present in a robust modeling of a system. When emergence is displayed in such models, we have therefore a grasp on the factors responsible for emergent features.

The last consequence involves the controversial issue of "downward causation." Traditionally emergence also goes with novelty and downward causation. The latter is the most controversial, and in this context may be explicated by the relationship between the explanatory force of the global emergent level of m-regularities, and the level of building blocks, which enable only local explanations. Indeed, the m-regularities, once stated, will somehow constrain the behavior of those building blocks. If I know that an uprising is likely to happen here (because the parameters are in the range of values upon which uprising depends), then I know that most of the agents in these circumstances will often enter into rebellion, even if the uprising is not a cause but is generated by the acts of the agents. Recognizing this constraint by global-level m-regularities neither deprives the agents or cells or building blocks of their unique causal efficacy, nor add other sets of causes and forces acting somehow from above the agents (see also the critique of strong emergence in ABM by Epstein 1999). So the value of the parameters is a relevant causal factor (of the uprising), as well as the acts of the agents. Yet the classical arguments (e.g., Kim 1999) against downward causation—in terms of overdetermination—do not hold here, because the relationship between those two categories of causes is not exactly one of levels: The values of the parameters are not *caused* by the acts of the agents. So we have indeed a case of "downward causation"—except that precisely because the "levels" relation does not hold, the term may be inappropriate. (see also Humphreys 1997)

As for novelty, the analysis provided here may not be directly illuminating. In order to develop this last traditional connotation of emergence in the framework of computational emergence, I consider in the last section a case study, novelty in evolution or the problem called open-ended evolution.

5 Novelty: a case study about artificial life and open-ended evolution

Evolution of life on Earth shows continual adaptability and innovation. Adaptation is traditionally explained by natural selection (Brandon 1996; Williams 1966), and so all genetic algorithms based on natural selection display adaptation. Innovations

[15] Notice that the argument proposed here corresponds to the "inference to best explanation" for scientific realism.

[16] Metaphysically speaking, an argument for that could derive from Quine's dictum on reality ("to be is to be the value of a linked variable") and the widely-accepted metaphysical thesis that real entities necessarily have a causal effect.

Journal: 11229-SYNT Article No.: 9721 MS Code: SYNTHE-4 TYPESET DISK LE CP Disp.:2010/2/23 Pages: 21 Layout: Small-X

concern novelties such as the wings of insects, the gills of fish or language: They are a radical qualitative difference in morphology or behaviour, and trigger new phylogenetic patterns and adaptive radiation (Müller 2002). Some biologists argue that natural selection alone cannot account for innovation, but only, through cumulative selection, for continuous adaptation. Processes like developmental scaffolding (e.g., Müller and Newman 2005; Carroll 2005; Arthur 2001) have been proposed to supplement natural selection for mutation and recombination here—the bulk of the critique being that normal mechanisms of variation such as mutation and recombination cannot allow a selective process that would yield such innovations. The alternative view of evolution trying to emphasize the role of development, and known as evolutionary developmental biology or evo-devo, is widely concerned by such issues (Amundson 2005). Other suggestions have been made for processes responsible for innovations, for example symbiosis—and some computer algorithmic modelling has helped to understand the potential for novelty proper to those processes (Watson 2005).

(Lenaft, 2000)

In this sense, artificial life[17] could model not only adaptations (in genetic algorithms), but also processes of innovations. Some argue that life on Earth displays "open-ended evolution" in the sense that innovation is pervasive and also that no limit can be set to this propensity for novelty. Evolutionary theory can yield predictions, but unpredictable events are always ~~opening~~ new kinds of possible adaptations—for instance, evolution of wings is explained by natural selection, but before insects appeared no one could have predicted that day animals would also be living in the sky, as Kauffmann likes to emphasize. Some term this propensity "creativity" (Taylor 2001), others "open-ended evolution." Anyway, the problem remains: Is artificial life, which already gives rise to adaptation and some innovations, able to display open-ended evolution, so that nothing uniquely proper to life on Earth is involved in this feature of evolution? Or are there several large classes of evolution, all possibly instantiated by some kinds of computer simulations, among which one class displays open-ended evolution and includes, as a token, life on Earth, along with other evolution designed by some appropriate ALife design? In this last case, it means that we could pinpoint the causal factors relevant to open-ended evolution.

creating

one

s

Such a question of course is entangled with the issue of contingency in evolution. If evolution is contingent, then it is unpredictable—meaning that some events extrinsic to the course of natural selection (such as the fall of an asteroid) triggered major events within evolution. It is not absolute unpredictability—for example, given the appropriate hypothesis, e.g., the fall of the asteroid, we could predict some outcomes; but it is unpredictability within a given model. Gould says that a model of biological microevolution at the times of the Cambrian would not predict evolution afterwards because the most relevant cause (the asteroid) is a hypothesis extrinsic to such a model. On the other hand if evolution is necessary, it means that even if we "replay the tape of life" (Gould 1989; Fontana and Buss 1994), as in the famous Gouldian metaphor, we will still find the same kinds of features: mobile entities, detectors, light detectors, motion tracking devices, etc. Basically, the models of evolution elaborated in the classical modern synthesis theory of evolution (macroevolution extrapolated from

[17] See Langton (1989) for an overview, and Lange (1996) for a philosophical assessment.

Journal: 11229-SYNT Article No.: 9721 MS Code: SYNTHB-4 TYPESET DISK LE CP Disp.:2010/2/23 Pages: 21 Layout: Small-X

micro evolution) would be robust against all extrinsic possible circumstances (i.e., on average natural selection and drift account for what happens). So in the end we would not have mammals and nervous systems, but we would have creatures able to track changes in their environment in order to prey and resist parasites. Dennett (1995) calls this "good tricks." Running many simulations in artificial life is meaningful under this hypothesis, because it will display some causal factors relevant to evolution and the main qualitative patterns of evolution, i.e., the "good tricks" with their patterns of causal efficiency. The "contingency" issue might be a question of fine-grained vs. coarse-grained views of evolution. ~~The~~ predictions, if we think of Kauffman's wings example, are always possible when some novelty happens. So evolution seems predictable at a local scale (at least in its gross features), but not at the global scale, because those novelties are themselves contingent. ~~The~~ two debates (necessity or contingency in evolution, uniqueness of open-ended evolution in life) are not identical, ~~but~~ are quite related. Open-endedness means the potential for novelties (such as the wing) that are contingent in a very large-scale view of evolution. Many classifications have been proposed in the evolutionary computer sciences literature, and these final paragraphs sketch a survey of them and suggest a synthetic classification based on the present analysis of the status of algorithmic models and discontinuous evolution. The consequence of such classification is that it could help us to understand what is proper to life in our world—in opposition to a would-be evolution by natural selection in other possible worlds.

Eventually, it
Coarse grained
Although the
they

Those classifications are built, either in a sort of "semantic" perspective (considering what it means to say that "X is novel"), or from an epistemological point of view (considering the various cases of robust emergence in the various scientific contexts).

- *Semantically*, from the computational point of view, Cariani and Ray (1992) distinguished syntactic, semantic and pragmatic emergence; Crutchfield (2002) distinguished emergence in intuitive sense, emergence as patterns newly appearing, and "intrinsic emergence;" Pattee (1989) seems equally inspired when he indicates "new measurement" as the third and truest kind of emergence.
- *Epistemologically*, Bedau et al. (1998) distinguished three kinds of emergence: Class II is "bounded emergence"—Holland's (1995) GA Echo—as opposed to class I, no emergence, in an Echo simulation *with no selection* (what they call "Echo neutral shadow"), and class III is "unbounded emergence," manifest in the phanerozoic fossil records (i.e., the history of life). "Bounded" for Bedau and Packard means that the range of adaptations exhibited is somehow finite, which is not the case in class III. Contrary to the example of the contingent emergence of wings and the subsequent colonization of avian space, in Echo, it makes no sense to speak of another environment likely to be colonized: "Digital organisms" will be restricted to adapting to computer environments; hence the range of adaptations is bounded.
- Channon's classification is also tripartite, elaborating on his Geb simulation: (Channon and Damper 1998, 2000) (1) artificial selection in the SAGA simulation, (2) natural selection of program codes in Ray's Tierra, which seems a now-limited evolution, and (3) less limited

Journal: 11229-SYNT Article No.: 9721 MS Code: SYNTHB-4 TYPESET ☑ DISK LE CP Disp.:2010/2/23 Pages: 21 Layout: Small-X

evolution through Channon's "natural selection"[18] in Geb simulation. Hence the major question about novelty and open-endedness: Is the Channon class 3 in the latter analysis the same as Bedau and Packard's class III (fossil records, i.e., life)?

I would suggest another classification, based on the considerations developed here.

A. Imagine a simulation with no selection acting like Kauffman's "order for free" in Boolean networks. This category allows us to embrace those cases of emergence that are obviously not biological adaptations, namely phase transitions. However, it can allow for some discontinuities in the history of life that do not need any selection as explanation. Those situations are likely to concern the lowest stages of life, for example, the first emergence of replicators (Fry 1995). Selection is likely to act, and act powerfully, at highest levels, in ways that strongly counter those kinds of evolution with no selection. / (1993)

B. Computer simulations can display features of natural selection climbing a peak on an adaptive landscape. In this class fall classical genetic algorithms with fixed fitness function, or some of the Bedau–Packard class II (Echo does not have a fixed fitness function), or Channon's "artificial evolution." This is the realm of gradual evolution by cumulative selection; adaptations are continuously produced,[19] but there will not be many discontinuities, and then few novelties and no open-ended evolution.

C. Non-gradual evolution; here emergence, innovations and novelties are not conceivable in terms of cumulative selection, like in the case of those novelties considered by evo-devo (Carroll 2005). Some maladapted processes ("maladapted" in the sense of not-fitness-increasing, i.e., no simple "hill climbing") entail new kinds of adaptation, as the decreases of mean fitness implied by the novelties named "evolutionary transitions" by Maynard-Smith and Szathmary (1995) or Michod (1999) (i.e., cohesive groups of entities becoming new units of fitness (cells, multicellular organisms, colonies)). Interesting cases in this category are the processes of "compositional evolution," as modeled by Watson (2005) in modified genetic algorithms that incorporate features proper to sexuality and symbiosis in real life. In this class we therefore find Channon's class 3, but also Pattee and Crutchfield's class III. Robust emergence (in the sense of Sect. 1) clearly characterizes those features; the pending question is to find subclasses of models for which a range of parameter values would robustly entail not only m-regularities between emergent sets of cells or agents, but continuous production of new m-regularities, or even of new entities supporting novel m-regularities. For the moment (Bedau et al. 1998) no such thing has been found, although member of the ALife community are working on it (Taylor 2001). /s

Those categories leave an open question: Is there anything particular with Bedau and Packard's class III, namely *the history of life*? Is it included, in fact, in my class C

[18] Channon opposes *artificial selection* and *natural selection* as genetic algorithms with a fixed fitness function vs. genetic algorithms with evolving fitness functions. / (Channon and Damper 2000)

[19] The problem for populations is to avoid being stacked on a local optimum; hence the shifting balance theory—yet some more recent concepts such as neutral tunnels of fitness (Gavrilets 2003) account for solutions that do not need drift.

Journal: 11229-SYNT Article No.: 9721 MS Code: SYNTHB-4 ☐ TYPESET ☑ DISK ☐ LE ☐ CP Disp.:2010/2/23 Pages: 21 Layout: Small-X

here, which would mean that life—in the vernacular sense—has nothing more than all the features of those processes that are likely to be simulated. Or is there something more about it?

Bedau (1999) argues that for the moment, open-ended unbounded adaptation as well as increase in complexity which we find in the phanerozoic fossil record is not reached as a result of the various simulations ~~we~~ currently ~~have~~ available—be they from class II in his classification or class C of mine. This is just a current fact, and nothing precludes that some more sophisticated future simulations will yield patterns of complexity and unbounded adaptation that will match the fossil record's evolutionary pattern (Bedau et al. 1998). The point is that the processes we are acquainted with, in the sense that we can now implement them in computer simulations, are not enough to recreate the whole pattern of adaptive radiation or rates of novelties met in the phylogenic trees. On the other hand we can reproduce increases of complexity that match the increase of complexity found in the tree of life (Adami et al. 2000), estimating complexity with some informational-based measures of information (Ofria et al. 2000; Adami 2002). This implies that the Darwinian features (variability, inheritance, fitness) of the entities used in our computer simulations are enough to produce such trends towards ~~novelty~~; ~~inversely, it means that *computational criteria of emergence seem sufficient to account for the field of life and nature in general*~~ complexity

Yet, there remains the fact that open-ended adaptive processes, with continuous emergence of innovations and new kinds of adaptations, characteristic of the Phanerozoic record, are not yet reproduced by our simulations. This means that the default hypothesis of investigations into the processes of long-term biotic evolution is that transparent ~~pure possible~~ processes—transparent in the sense that they are likely to be implemented on computers—are not alone responsible for the open-ended adaptive trend in evolution. The major benefit of this approach is not a definitive statement of *ignorabimus* but, on the contrary, a firm basis for an empirical approach to such an issue: Find a way to falsify the default hypothesis, find evidence for a clear and testable alternative hypothesis, or progressively defeat these alternatives. Such a program is not at all complete, but is indeed feasible. To this extent, the specificity of novelty in biological evolution—open-ended evolution, as we say—is still in question; the computational concept of emergence has therefore the virtue of allowing us to formulate the empirical question about the uniqueness of features of novelty in life.

6 Conclusion

Emergence is a word whose use exponentially increased in science and in discourses about science since hard reductionist programs in physics and biology seemed to reach some limits two decades ago. However, this use is often quite loose, and most of the time irreducibility and novelty invoked by its tenants are in fact likely to be dissolved in a reductionist framework, as happened to earlier emergentism of the 1920s. Computational incompressibility, however, provides a concept of emergence likely to resist such objection, and to be not merely epistemological—which also includes "robust emergence" as computationally emergent processes that display "m-regularities." On this basis one can recapture the other dimensions of what is understood under the name

Journal: 11229-SYNT Article No.: 9721 MS Code: SYNTHB-4 TYPESET DISK LE CP Disp.:2010/2/23 Pages: 21 Layout: Small-X

of emergence, namely irreducibility, downward causation (as relations of explanatory constraint) and novelty. Furthermore, robustness analysis of several simulations displaying computational emergence entitles one to ascribe the property of emergence to a target system and specify the causally relevant parameters for the emerging features, so that the concept of emergence can exceed the context of digital processes and simulations and concern the ~~objective~~ world. real

Reciprocally, the very idea of predictability is enriched by such an approach. Some systems do not allow predictability, albeit they are deterministic—and such a property of predictability is ontologically proper to the system, although it concerns an epistemic feature, namely prediction. Computationally emergent systems are not predictable as to the fate of their component elements. Yet robust emergence—in the sense of the arising of sequences displaying counterfactual dependencies (m-regularities) between sets of building blocks—allows one to elaborate partial predictions at this macrolevel (m-predictions) as well as M-predictions based on the global relations between sets of outcomes of simulations and ranges of the values of parameters of simulations. The study of complex systems extensively draws on this possibility of forging qualitative macro predictions, and identifying bifurcations between qualitative behaviours—sometimes understood as attractors—on the basis of the ranges of values that such bifurcations discriminate. Along those lines, the dimension of novelty, proper to the vernacular concept of emergence, can be circumscribed. However, biological evolution displays not only novelties but also increasing trends towards novelties (open-endedness). The rates of producing novelties are themselves proper to some systems, and investigating the conditions for open-ended evolution means investigating the dependency of those rates upon the various parameters proper to biological evolution. It is now an empirical question to decide whether the highest rates of production of novelties, which define open-ended evolution, are within the reach of computer simulations or involve some other features proper to biological evolution. present However, the point is that the general concept of computational emergence is instantiated by lots of systems, and can be further specified into more refined classes—first, "robust emergence," and then, degrees of emergence producing more or less high rates of novelty. The present analysis therefore provides a multifaceted objective concept of emergence of which several nuances allow one to ask questions about the varieties of novelty in natural systems in general, and the specificity of novelty in evolutionary biology in particular.

References

Adami, C. (2002). What is complexity. *Bio Essays, 24*, 1085–1094.

Adami, C., Ofria, C., & Collier, T. (2000). Evolution of biological complexity. *PNAS, 97*(9), 51–61.

Amundson, R. (2005). *The changing role of the embryo in evolutionary theory*. Cambridge: Cambridge University Press.

Arthur, W. (2001). *The origins of animal body plan*. Oxford: Oxford University Press.

Barabasi, A.-L., & Albert, R. (1999). Emergence of scaling in random networks. *Science, 286*, 509–512.

Batterman, R. (2002). *The devil in the details*. Oxford: Oxford University Press.

Bedau, M. (1997). Weak emergence. In J. Tomberlin (Ed.), *Philosophical perspectives: Mind, causation, and world* (Vol. 11, pp. 375–399). Oxford: Blackwell Publishers.

~~Bedau, M. (1998). Four puzzles about life. *Artificial Life, 4*, 125–140.~~ HS

Journal: 11229-SYNT Article No.: 9721 MS Code: SYNTHB-4 TYPESET DISK LE CP Disp.:2010/2/23 Pages: 21 Layout: Small-X

Bedau, M. (1999). Can unrealistic computer models illuminate theoretical biology. In A. Wu (Ed.), *Proceedings of the 1999 genetic and evolutionary computation conference workshop program* (pp. 20–23). Florida.

Bedau, M. (2003). Downward causation and the autonomy of weak emergence. *Principia, Revista Internacional de Epistemologia, 6*, 5–50.

Bedau, M. (2008). Is weak emergence just in the mind? *Minds & Machines, 18*, 443–459.

Bedau, M., Snyder, & Packard, N., et al. (1998). A classification of long-term evolutionary dynamics. In C. Adami (Ed.), *Artificial life VI* (pp. 189–198). Cambridge, MA: MIT Press.

Bell, G., Lechowicz, M., & Waterway, M. (2006). The comparative evidence relating to functional and neutral interpretations of biological communities. *Ecology, 87*(6), 1378–1386.

Brandon, R. (1996). *Adaptation and environment*. Princeton: Princeton University Press.

Burke, M., Furnier, G., & Prasad, K. (2006). The emergence of local norms in networks. *Complexity, 11*(5), 65–83.

Buss, S., Papadimitriou, C., & Tsisiklis, J. (1992). On the predictability of coupled automata: An allegory about chaos. *Complex Systems, 5*, 525–539.

Carroll, S. (2005). *Endless forms most beautiful*. New York: Norton.

Chalmers, D. (2006). Strong and weak emergence. In P. Clayton & P. Davies (Eds.), *The reemergence of emergence*. (pp. 244–256). Oxford: Oxford University Press.

Channon, A., & Damper, R. (1998). Perpetuating evolutionary emergence. In R. Pfeifer, B. Blumberg, J. A. Meyer, & S. Wilson (Eds.), *From animals to animats*. (pp. 534–539). Cambridge, MA: MIT Press.

Channon, A., & Damper, R. I. (2000). Towards the evolutionary emergence of complex increasingly advantageous behaviours. *International Journal of Systems Science, 31*(7), 843–860.

Chave, J. (2004). Neutral theory and community ecology. *Ecology Letters, 7*, 241–253.

~~Chu, J., & Adami, C. (1999). A simple explanation for taxon abundance patterns. *Proceedings of the National Academy of Sciences of the United States of America, 96*(26), 15017–15019.~~

Corning, P. (2002). The re-emergence of "emergence": A venerable concept in search of a theory. *Complexity, 7*(6), 18–30.

Cracraft, J. (2000). The origin of evolutionary novelties: Pattern and process at different hierarchical levels. In M. Nitecki (Ed.), *Evolutionary innovations*. (pp. 21–43). Chicago: University of Chicago Press.

Crutchfield, J. (1994). The Calculi of emergence. *Physica D*, special issue on the Proceedings of the Oji International Seminar. *Complex Systems—from Complex Dynamics to Artificial Reality*.

Crutchfield, J. (2002). When evolution is revolution: Origins of innovation. In J. Crutchfield & P. Schuster (Eds.), *Evolutionary dynamics: Exploring the interplay between accident, selection, neutrality and function*. (pp. 101–133). Oxford: Oxford University Press.

~~Crutchfield, J., & Hanson, J. (1993). Turbulent pattern bases for cellular automata. *Physica D, 69*, 279–301.~~

Crutchfield, J., Mitchell, M., & Hraber, D. (1993). Revisiting the edge of chaos: Evolving cellular automata to perform computations. *Complex Systems, 7*, 89–130.

Dennett, D. (1995). *Darwin's dangerous idea*. New-York: Simon & Shuster.

Epstein, J. (1999). Agent based computational model and generative social science. In J. Epstein (Ed.), *Generative social science*. (pp. 1–46). Princeton: Princeton University Press.

Epstein, J. (2002). Modeling civil violence. *PNAS, 99*(3), 7243–7250.

Fontana, W., & Buss, L. (1994). What would be conserved if "the tape were played twice? *PNAS, 91*, 757–761.

Fry, I. (1995). Are the different hypotheses on the emergence of life as different as they seem? *Biology and Philosophy, 10*, 389–417.

Gavrilets, S. (2003). Evolution and speciation in a hyperspace: The roles of neutrality, selection, mutation and random drift. In P. Schuster & J. Crutchfield (Eds.), *Evolutionary dynamics: Exploring the interplay of selection, accident, neutrality and function*. (pp. 135–162). Oxford: Oxford University Press.

Gould, S. J. (1989). *Wonderful life: The Burgess Shale and the nature of history*. New York: W. W. Norton.

~~Hillis, D. (1998). *The pattern on the stone: The simple ideas that make computers work*. London: Basic Books.~~

Holland, J. (1995). *Hidden order: How adaptation builds complexity*. New-York: Addison-Wesley.

Holland, J. (1998). *Emergence: From chaos to order*. New York: Basic Books.

Journal: 11229-SYNT Article No.: 9721 MS Code: SYNTHB-4 TYPESET DISK LE CP Disp.:2010/2/23 Pages: 21 Layout: Small-X

Hovda, P. (2008). Quantifying weak emergence. *Minds & Machines, 18*, 461–473.

Hubbell, S. P. (2001). *The unified neutral theory of biodiversity and biogeography*. Princeton: Princeton University Press.

Humphreys, P. (1997). How properties emerge. *Philosophy of Science, 64*, 53–70.

~~Humphreys, P. (2004). *Extending ourselves*. Oxford: Oxford University Press.~~ H ℘

Humphreys, P. (2008). Synchronic and diachronic emergence. *Minds & Machines, 18*, 431–442.

Humphreys, P., & Huneman, P. (Eds.). (2008). Special issue on dynamic emergence and computation. *Minds and Machines, 18*.

Huneman, P. (2008a). Emergence and adaptation. *Minds and Machines, Special Issue on Dynamic Emergence, 18*(4), 493–520.

Huneman, P. (2008b). Combinatorial vs. computational views of emergence: Emergence made ontological? *Philosophy of Science, 75*, 595–607.

Israeli, N., & Goldenfeld, N. (2004). On computational irreducibility and the predictability of complex physical systems. *Physics Review Letters, 92*, 074105.

Kauffmann, S. (1993). *The origins of order*. Oxford: Oxford University Press.

Kim, J. (1999). Making sense of emergence. *Philosophical Studies, 95*, 3–36.

Lange, M. (1996). Life, "artificial life" and scientific explanations. *Philosophy of Science, 63*, 225–244.

Langton, C. (1989). Artificial life. In C. Langton (Ed.), *Artificial life: SFI studies in the sciences of complexity, Proc.Vol. VI.*. Redwood City, CA: Addison-Wesley.

Levins, R. (1966). The strategy of model building in population biology. In E. Sober (Ed.), *Conceptual issues in evolutionary biology.* (1st ed., pp. 18–27). Cambridge, MA: MIT Press.

Maynard-Smith, J., & Szathmary, E. (1995). *The major evolutionary transitions*. San Francisco: Freeman.

McLaughlin, B. (1992). The rise and fall of British emergentism. In A. Beckermann, H. Flohr, & J. Kim (Eds.), *Emergence or reduction?* (pp. 44–93). Berlin: de Gruyter.

Michod, R. (1999). *Darwinian dynamics*. Oxford: Oxford University Press.

Müller, G. (2002). Novelty and key innovations. In M. Pagel (Ed.), *Encyclopedia of evolution.* (pp. 827–830). Oxford: Oxford University Press.

Muller, G., & Newman, S. (2005). The innovation triad: An evodevo agenda. *Journal of Experimental Zoology (Mol Dev Evol), 304B*, 487–503.

Newman, D. (1996). Emergence and strange attractors. *Philosophy of Science, 63*(2), 245–261.

Pattee, H. (1989). Simulations, realizations and theories of life. In C. Langton, et al. (Ed.), *Artificial life I* (pp. 63–78). Redwood City, CA: Addison Wesley.)

Ray, T. (1992). An approach to the synthesis of life. In C. Langton (Ed.), *Artificial life II.* (pp. 371–408). Redwood City, CA: Addison Wesley.

Reynolds, C. (1987). Flocks, herds and schools: A distributed behavioural model. *Computer Graphics, 21*(4), 25–34.

Shalizi, C., & Crutchfield, J. (2008) *Computational mechanics: Pattern and prediction, structure and simplicity*. Santa Fe Institute Working Paper 99-07-044

Shalizi C., & Moore C. (2003). What is a macrostate? Subjective observations and objective dynamics. http://philsci-archive.pitt.edu/archive/00001119/

Silberstein, M. (2002). Reduction, emergence and explanation. In M. Silberstein & P. Machamer (Eds.), *Blackwell companion to the philosophy of science.* (pp. 80–107). Oxford: Blackwell.

Tassier, T. (2004). A model of fads, fashions and group formations. *Complexity, 9*(5), 51–61.

Taylor, T. (2001). Creativity in evolution: Individuals, interactions and environments. In P. J. Bentley & D. W. Corne (Eds.), *Creative evolutionary systems.* (pp. 74–108). San Francisco: Morgan Kaufman.

Watson, R. (2005). *Compositional evolution: The impact of sex, symbiosis and modularity on the gradualist framework of evolution*. Cambridge: MIT Press.

Watts, D. J. (2002). A simple model of global cascades on random networks. *Proceedings of the National Academy of Sciences, 99*, 5766–5771.

Watts, D. J., & Strogatz, S. H. (1998). Collective dynamics of small-world networks. *Nature, 393*, 440–442.

Weisberg, M. (2006a). Robustness analysis. *Philosophy of Science, 73*, 730–742.

Weisberg, M. (2006b). Forty years of "the strategy": Levins on model building and idealization. *Biology and Philosophy, 21*(5), 623–645.

Williams, G. (1966). *Adaptation and natural selection*. Princeton: Princeton University Press.

Wilson, J. (~~forthcoming~~). Non reductive physicalism and degrees of freedom. *British Journal for Philosophy of Science*. 2009

Wimsatt, W. (1997). Aggregation: Reductive heuristics for finding emergence. *Philosophy of Science, 64*(4), 372–384.

Wimsatt, W. (2007). Emergence as non-aggregativity and the biases of reductionisms. In W. Wimsatt (Ed.), *Piecewise approximations to reality*. Cambridge: Harvard University Press.

Journal: **11229-SYNT** Article No.: **9721** MS Code: **SYNTHB-4** ☐ TYPESET ☑ DISK ☐ LE ☐ CP Disp.:**2010/2/23** Pages: **21** Layout: **Small-X**