Emergence Made Ontological? Computational versus Combinatorial Approaches

Philippe Huneman[†‡]

I challenge the usual approach of defining emergence in terms of properties of wholes "emerging" upon properties of parts. This approach indeed fails to meet the requirement of nontriviality, since it renders a bunch of ordinary properties emergent; however, by defining emergence as the incompressibility of a simulation process, we have an objective meaning of emergence because the difference between the processes satisfying the incompressibility criterion and the other processes does not depend on our cognitive abilities. Finally, this definition fulfills the nontriviality and the scientific-adequacy requirements better than the combinatorial approach, emergence here being a predicate of processes rather than of properties.

1. Introduction. Defining emergence in philosophy of science has to fulfill two requisites: first, the concept has to match a great deal of the scientific practice (i.e., the "scientific adequacy requirement") and cover some of the uses of the term; second, the concept has to be such that not too many things will fall under it (i.e., the "nontriviality requirement"). The latter is obvious since, if "emergent" actually means something, it is regarding issues about novelty, so the concept should not be trivially instantiated.

In the literature, emergence is often addressed through the question of emergence of *properties* (O'Connor 1994; Newman 1996; Crane 2001;

†To contact the author, please write to: Institut d'Histoire et de Philosophie des Sciences et des Techniques (CNRS/Université Paris I Sorbonne), 13 rue du Four 75006, Paris; e-mail: huneman@wanadoo.fr.

‡I warmly thank Paul Humphreys and John Sumons for their precious comments on an earlier version of this article.

Philosophy of Science, 75 (December 2008) pp. 595–607. 0031-8248/2008/7505-0048$10.00
Copyright 2008 by the Philosophy of Science Association. All rights reserved.

https://doi.org/10.1086/596777 Published online by Cambridge University Press

Silberstein 2002; Chalmers 2006).[1] Another trend in approaches is the focus on *whole-parts relationships*. Emergence is often considered to be the problem of understanding properties of the wholes that are irreducible to properties of the parts, which I call combinatorial emergence. In this article I challenge the property approach and the dominant combinatorial understanding of emergence by focusing on the computational definition of emergence. I argue that this approach (focusing on what is an emergent process, rather than on the emergence of properties) fulfills the initial requisites better than the combinatorial/property approach.

2. The Whole-Parts Approach of Emergence and Its Flaws. Traffic jams (Nagel and Rasmussen 1994), fads (Tassier 2004), temperature, and chromosomes at the time of meiosis exhibit a behavior that is not understandable by adding up the considerations of the behavior of their parts. Hence, one considers them emergent behaviors: this emergence is thought to be something proper to the whole and irreducible to the parts. Philosophers such as Bechtel and Richardson (1992), O'Connor (1994), and Silberstein (2002) addressed the issue of emergence through this scheme of wholes and parts. Dessalles and Phan (2005) see emergence as a drop in complexity, and Wilson (forthcoming) sees it as a decrease in the degrees of freedom—contrary to the simple product of the properties of the parts (where there would be additivity of degrees of complexity/degrees of freedom; see also Atay and Jost 2004, 18).

But take Schelling's model of segregation (1969): in this model, according to their color, agents will arrange themselves into some homogenous clusters. The behavior "join the cluster" is surely not given in the rules of behaviors of the agents, but the clusters are not exactly *composed* of the agents, since they remain themselves even if some agents are added and some "die" (Gilbert 2002). So, since the parts are transient relative to the whole, a simple view of emergence as irreducibility of properties of the whole to properties of the parts is misguided.

Wimsatt (1997, 2007) defined emergence as the failure of aggregativity; the point is thereby to provide some criteria of aggregativity—this definition addresses the emergence problem the other way around. Wimsatt's criteria of failure of aggregativity are a more sophisticated formalization of what is the case when we say that we cannot reduce properties of the parts to properties of the whole. These criteria are: invariance through

1. Among other options is the emergence of laws. To the extent that laws are conceived of as theoretical constructs accounting for regularities observed, this makes emergence epistemological as opposed to ontological. But law itself is such a problem that it is not a good strategy to address emergence through a notion quite as hard as emergence to construe (Klee 1984).

https://doi.org/10.1086/596777 Published online by Cambridge University Press

intersubstitution of parts; qualitative similarity under addition/withdrawal of parts; invariance regarding decomposition/reaggregation of parts; no cooperative/inhibitory interactions. They are general criteria of invariance; hence they take into account the case of alternating and changing parts within a whole. However, it appears that, except for mass, almost nothing is really aggregative, that is, nothing satisfies all the criteria, for example, of invariance regarding permutation of parts, and so on. This surely is a problem if one wants to capture the meaning of emergence by this perspective—emergence should surely apply to fewer properties than "everything but the mass," so it would then need a supplementary criterion that is not provided in this analysis. Hence, we are just left with the idea that emergence comes in degrees (see also Bechtel and Richardson 1992). However, in this view the meaning of emergence is quite superfluous; that is, we could talk merely of degrees of aggregativity, but it would make sense only if we already have this criterion that determines how emergence is more than some lack of aggregativity—but our combinatorial analysis will not provide it.

Actually, emergence is supposed to cover several features: unpredictability, novelty, and irreducibility (Klee [1984], O'Connor [1994], Humphreys [1997], Crane [2001], Silberstein [2002], Seager [2005], and Chalmers [2006] largely agree on those features). Irreducibility construed as irreducibility of the properties of wholes regarding properties of parts appears now quite trivial and is too frequent to yield something as "emergence" (see also Bar Yam 2004). Concerning novelty, since properties of the wholes are quite always novel regarding the properties of the parts—think of color, or mass, or volume—the problem is to choose *which* novelty should count as emergent. We are left here with no objective criterion. Novel means most of the time something that has yet to be named (Epstein [1999] 2007). So this inevitably leads to a widely shared conclusion: if emergence means something, it is restricted to epistemological emergence, that is, relative to a set of theories and of cognitive abilities—apart perhaps from the exceptional case of qualia (O'Connor 1994; Silberstein and McGeever 1995; Crane 2001; Seager 2005; Chalmers 2006).

3. The Incompressibility Criterion and Emergent Processes. In the framework of computer simulations some people have been able to define what Mark Bedau (1997, 375) called "weak emergence." According to this criterion, a state of a computation process is weakly emergent iff there is no shorthand to get to it except by running the simulation (this defines an incompressibility *criterion of emergence*; see Huneman and Humphreys 2008).

Such an approach avoids the subjectivity problem proper to the novelty issue in the previous approach. Hence we would have a major clue about

https://doi.org/10.1086/596777 Published online by Cambridge University Press

an emergence that would be, if not ontological, at least *objective* in the same way that *conceptual* truths in mathematics are objective, independent of our cognitive abilities or epistemic choices. But one could object that our incompressibility criterion is only *provisory*, since we cannot assert that in a distant future, with enlarged computational abilities, we will still be unable to find analytical shortcuts to get to a final state faster than by simulation. Yet, there is some evidence that this objection fails.

I will extract arguments for the objectivity of computational criteria from Buss, Papadimitriou, and Tsisiklis (1992). The basic idea is to construe a set of automata whose values change according to a global rule ***R***. Each automaton transforms the value of its cells according to an input, 0 or 1. The application of the global rule ***R*** is itself dependent on the numbers of each value (q_1, q_2, . . .) in the set of automata at step n; the input function that will determine the input to all automata at step $n + 1$ is determined by the global rule. Hence the system is perfectly deterministic.

> Input function: If $Z_n = 0$, then $F(n + 1) = g_0(F(n))$, and if $Z_n = 1$, then $F(n + 1) = g_1(F(n))$, where functions g_0 and g_1 take their values in $\{q_1 \ldots q_j \ldots q_n\}$.
>
> Global rule ***R***: Z_i takes its values in $\{0, 1\}$. $Z_i = M(N_i(q_1) \ldots N_i(q_j) \ldots N_i(q_n))$, where $N_i(q_j)$ is the number of times value q_j is taken at step i (see table 1).

Some global rules are *constant free*, meaning that they can be enunciated with no reference to some of the actual values q_i . . . of the constants, and others are not. "If there is as much q_i as q_j, for each value of i and j, $Z = 0$; else $Z = 1$" would be an example of a constant-free rule. Buss et al. (1992) show that if the global rule is non–constant free, then the problem of predicting the state of the system at time T is PSPACE complete; hence the problem cannot be solved in polynomial time (since we

TABLE 1

Step 0	$F_1(0)$	$F_2(0)$		$F_i(0)$				$F_m(0)$
Step 1	$F_1(1)$	$F_2(1)$		$F_i(1)$				$F_m(1)$
Step k	$F_1(k)$	$F_2(k)$		$F_i(k)$				$F_m(k)$
Step n	$F_1(n)$	$F_2(n)$		$F_i(n)$				$F_m(n)$

https://doi.org/10.1086/596777 Published online by Cambridge University Press

assume that no $P = NP$[2] and NP problems are included in PSPACE problems, so that being PSPACE complete entails being such that all NP problems can be translated into that problem, which makes it at least harder than NP complete). The detail of the proof relies on the fact that constant-free global rules are preserved through permutation of the q_i, which makes a major difference concerning the pattern of computation of the prediction.

This result perfectly illustrates the fact that some computational devices are objectively incompressible. As the authors write: "If the prediction problem is PSPACE complete, this would mean essentially that the system is not easily predictable, and that there seems to be no better prediction method other than simulation" (Buss et al. 1992, 526). Even with infinite cognitive abilities, there would still be a genuine difference between PSPACE complete prediction problems and others, so the computational definition of emergence is *objective*.

Weak emergence defined as inaccessibility except by simulation is thereby *not something trivial*, since, in this framework, all global rules that are constant free are computable in polynomial time, so we have a clear delimitation between some weakly emergent cases and some other cases.

4. Patterns, Order, and Unpredictability. Now, speaking of computational emergence, *what* precisely is emergent? Humphreys (2004, 2008) talks about *emerging patterns* in cellular automata; the robust pattern, not computable except by simulation, is then an emergent property in a cellular automaton (CA). But for Dennett (1991), a pattern is an array of traits that can be easily distinguished from others and recognized; hence it should have a high level of redundancy. A pattern (D-pattern) in this sense is easily deducible from one of its parts (contrary to Humphreys [H-patterns]). So defined, D-patterns seem rather mind-dependent, since Dennett addresses patterns through their recognition, while recognition is linked to our cognitive abilities. Nevertheless, we can specify those redundancy patterns by some purely internal properties, like redundancy, symmetry, or any case of invariance through homeomorphisms: this renders D-patterns mind-independent.

It seems that there will be no way to understand through the computational approach some of the instances of "emergent D-patterns" identified by combinatorial approaches. Basically, if someone draws a D-

2. Even if this assumption turns out to be false, the fact that NP is strictly included in PSPACE implies that being PSPACE complete makes the prediction problem for non-constant-free rules incommensurably harder for any cognitive ability than the one for constant-free rules that is solvable in polynomial time.

https://doi.org/10.1086/596777 Published online by Cambridge University Press

pattern on an $n \times n$ grid, there will be some symmetries, and so forth, so it will need less than n^2 pieces of information to be described. Thereby, the whole is not the aggregation of the parts. But in precisely this case the final pattern seems to contradict the incompressibility criterion, because the symmetries provide shorthand to figure out the total pattern.

Therefore, there are two poles included in the vernacular idea of emergent patterns: order (mostly instantiated by D-patterns) and unpredictability. Segregation clusters, traffic jams, and fads (Nagel and Rasmussen 1994; Gilbert 2002; and Tassier 2004) reveal some order instead of the pure heterogeneity that we should have expected. They exhibit a clear drop in complexity (e.g., a segregation cluster can be defined in n zones rather than $p >>> n$ points). This "order" aspect easily raises the objection that it is epistemological (Dessalles and Phan 2005): emergent orders always depend on our epistemic abilities since they are relative to an "expected" complexity.

On the other hand, the "unpredictability" aspect is captured by the incompressibility criterion. Yet there is a similarity between Chaitin's definition of randomness of a random sequence of numbers and the incompressibility criterion of weak emergence, but if this criterion includes randomness, how could weak emergence be reconciled with the "order" aspect? I must now explicate such a similarity.

In a random sequence, the nth digit and its precursors do not give any information regarding the $n + 1$th digit, x, and this is of course the intuitive idea of randomness; that is, there is no reason for x to be something rather than something else. Chaitin's idea of the lack of an algorithm that would save us from this step-by-step examination of the sequence formally captures this naïve understanding. Hence x is unpredictable relative to the previous sequence: so if you want to know x you have to read all the $n + 1$ first digits. On the other hand, the incompressibility criterion of emergence formulates the unpredictability of a given state from the knowledge of the rule and initial state—but of course, not of step $n + 1$ relative to step n, since this is perfectly deterministic and determined. Briefly said, Kolmogorov/Chaitin randomness requires that there not be a program whose length is significantly shorter than the length of the sequence of states. But there can be a short program—the rules of the CA—that produces an incompressible sequence of states.

The order and the unpredictability aspects seemed to contradict one another, precisely because we implicitly oppose order and randomness here. However, given the aforementioned crucial difference between randomness incompressibility criterion and emergence incompressibility criterion, there is no genuine opposition between a computational criterion of emergence, and order—so that some usual meanings of emergence that connote "order" are not in principle excluded by the computational view.

https://doi.org/10.1086/596777 Published online by Cambridge University Press

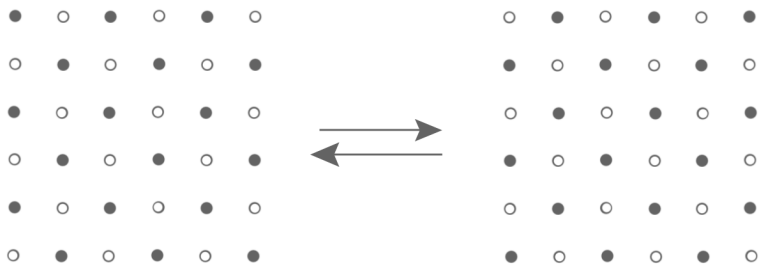


Figure 1. Blinking pattern b.

The problem is then to make sense of those connotations within computational emergence.

5. Triviality and the Randomness Exclusion Issue. The two approaches—emergence through order in a combinatorial perspective, and computational emergence—radically diverge on some simple cases. For instance, take the CA defined by $C_0 = 1, 0, 1, 0,$ etc., 0, and the rule: "if at least one neighbor is 1, turn 0 and reciprocally." The outcome will clearly be the blinking pattern b shown in Figure 1.

Here, we have a kind of pattern that is not included in each of the cells, but that can only be given through their relations. No descriptions of the parts (black, white, etc.) include the figure of the whole, which clearly exhibits a failure of aggregativity. But there is no incompressibility: to know the state of the CA at step n, I just need to check whether n is odd or even.

The computational notion of emergence does not precisely exclude this pattern as such, but only the CA (rule + initial state) so described. Yet, in some cases we could consider that this precise pattern emerges; for example, in Burke, Furnier, and Prasad's (2006) study about local norms, we are shown that with simple rules of imitation and fitness enhancement, and some specific initial patterns of distribution, in an agent-based model local norms will emerge. A blinking pattern b′ eventually appears as a final stable state in several configurations when some parameters concerning the initial distribution of signals and the choosing dispositions of the agents are settled.

Thus, according to the computational view it is never the pattern as itself that is emergent (since patterns b and b′ are identical), but *emergence is a feature of the whole agent-based simulation process*—otherwise there would be no difference between blinking patterns b and b′. (This makes

https://doi.org/10.1086/596777 Published online by Cambridge University Press

the computational view immune to the triviality proper to the combinatorial view.) As Humphreys (2008) points out, emergent patterns are always token patterns (and not types) since other tokens of the same types, for example, a photo of our blinking pattern b′, would not be emergent. What clearly individualizes the pattern token as a token is the CA *process* in which it is a part.

Such a consideration allows this approach to meet our initial scientific-adequation requisite. Actually, if X incompressibly results from an initial state X_0 in a CA, nothing prevents it to be a random number ("randomness exclusion issue"). In effect: suppose that $X = \Sigma_{(i=0,\ldots)} X_i\, 10^{-i}$ is a random number. Since class IV CA are universal Turing machines (see Wolfram 1984), there must exist a CA noted H with a given rule R and an initial state X_0 whose final state is X. So, given that there are so many CAs whose outcome is random, we can assume there will be plenty of "emergent states" according to the incompressibility criterion that are random sequences. Basically, there is no way to warrant that random patterns are precluded as such to be counted as emergent patterns, since nothing in the computational definition of emergence distinguishes between random and nonrandom patterns. Moreover, fishy properties like P "being in state X, Y, or Z," where X, Y, and Z are random numbers, are trivially emergent properties of H.

If our definition allows all those random numbers to count as emergent, we will lose its connection with the usual meaning of "emergence," which seems too high a cost to pay. In this case we should keep calling those patterns "incompressible" and not mess with emergence any more.

But since *emergence is a predicate of processes rather than of properties*, what is primarily said to be emergent is the process, and only secondarily some properties (states of an emergent process) will be said to be emergent. So fishy properties (like P) that are not individualized outcomes of a process cannot satisfy the predicate "emergent." And since we individualize the emergent items through the processes, if there exist descriptions allowing us to distinguish *various classes of processes*, this would involve a *principled distinction*—in the set of processes fulfilling the incompressibility criterion—between the noninteresting random patterns and other ones that better match the usual meaning of emergence.

This is available through the *computational mechanics* description language initiated by Crutchfield and Hanson (Crutchfield and Hanson 1993; Hordijk, Crutchfield, and Mitchell 1996; Hanson and Crutchfield 1997; Crutchfield and Shalizi 2001). The basic idea is that we can filter out any CA in such a way that "domains," "particles," and "frontiers" will naturally appear. Processes in the CA can then be described as interactions between particles, openings, and collisions of domains, and so forth. As some authors write:

https://doi.org/10.1086/596777 Published online by Cambridge University Press

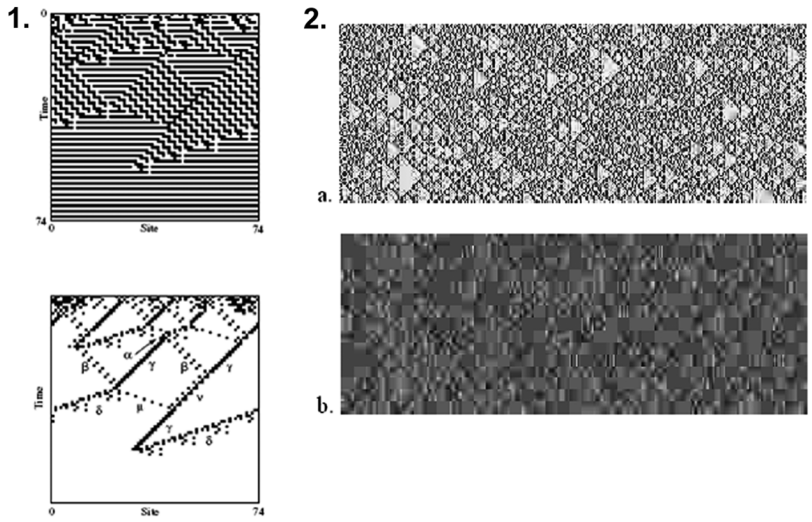


Figure 2. (*1*) After filtering, the types of particles (shown by Greek letters) and their interactions appear in a CA, whose final pattern seems to result from those "mechanics" (from Crutchfield and Hanson 1993). (*2*) A CA (rule 22, random initial configuration) (*a*) without and (*b*) with automatic "local sensitivity" filtering—then displaying no real domains and thereby appears chaotic (from Shalizi et al. 2006).

> The early empirical categorization of space-time patterns into four "classes" [reference to Wolfram]—loosely based on an analogy with those found in continuous state dynamical systems—has resisted numerous attempts at formalization. In many CA, it is immediately evident that the system self-organizes into some type of emergent pattern. In other CA, the structure or even existence of an emergent pattern is less clear. The question that naturally arises therefore is how to characterize the spatiotemporal patterns that emerge during the CA's evolution. If such a characterization is possible, it can be used as a basis for numerical and analytical tools that discover, analyze, and filter patterns that emerge in CA. (Hanson and Crutchfield 1997, 170)

In any case, we can define some sets of particles and domains: those sets will be proper to a CA, and defined by their generative structure (Figure 2, *1*). "Computational mechanics attempts to discover and characterize the *typical patterns* occurring in a given CA," wrote Crutchfield and Hanson (Hanson and Crutchfield 1997, 171). Delimitating the domain patterns that several CAs use as "building blocks," and then expressing them as finite automata, amounts to a possible typology of CAs. Hence

https://doi.org/10.1086/596777 Published online by Cambridge University Press

TABLE 2

Class: Computational Incompressibility			
Emergence of Properties:			
Kind	A. Fishy properties (disjunctive, etc.)	B. Emergent processes	
Species		B.1. Random patterns	B.2. "Interesting" patterns (specific subclasses characterized by proper particles-domains geography, measurable in complexity)

we will get *types* of elementary processes, which will allow us to define classes of emergent processes. Thus, as soon as there is a periodic "domain," we are ensured that it will display some regularities, and this prevents pure randomness; the range of periodic "building blocks" and their types of entanglement in CAs can thereby be the basis of a characterization and typology of incompressible processes that allows us to reliably *exclude* a large class of purely random patterns (not necessarily all) from what is likely to be emergent according to the mere incompressibility criterion. This amounts to integrate the usual "order" connotation of emergence into computational emergence. Accordingly, Shalizi et al. (2006) designed automatic filters that allow us to characterize processes with no domains and walls, and hence purely chaotic, and whose outcomes are random (e.g., in rule 22 the local sensitivity filter makes no pattern salient; see Figure 2). Computational mechanics hence provides a set of tools—filters—that help us to answer the randomness exclusion problem (Figure 2).

In the right column of Table 2, labeled B.2, we get an *objective*, *nonepistemological* meaning of emergence in CA (at least) that *fulfills* the nontriviality requirement and that is able to *confront correctly* the usual meaning of "emergence" (see also Huneman 2008 for biology).

6. Agent-Based Models and the Triviality Issue. Let's now widen the scope of computational emergence by considering simulations other than CAs. Epstein concludes his ([1999] 2007) investigation with a critique of the popular use of "emergence" in the discourse of people doing social simulations. His argument holds against a naïve version of the part-whole approach of emergence. Concerning *social agent–based modeling*, the dilemma he displays goes against any formulation of weak emergence. He formulates a comparison with the bee/beehive structure. "Typical of classical emergentism would be the claim: *No description of the individual bee can ever explain the emergent phenomenon of the hive*. . . . The mischievous piece of the formulation is the phrase 'description of the individual bee.' . . . Does 'the bee's' description not include its rules for interacting with

https://doi.org/10.1086/596777 Published online by Cambridge University Press

other bees? . . . My 'rules of social interaction' are, in part, what make me me. And, likewise, the bee's interaction rules are what make it a bee—and not a lump. When (as a designer of agent objects) you get these rules right—when you get 'the individual bee' right—you get the hive, too. . . . Unless the theoretical (model) bees generate the hive when you put a bunch of them together, you haven't described 'the bee' adequately. Thus, contrary to the opening emergentist claim, *it is precisely the adequate description of 'the individual bee' that explains the hive*" (Epstein [1999] 2007, 36–37). So, reciprocally, if you say that the beehive "emerges," you must say that *any outcome of a complete description of a set of entities* is emergent.

Basically, Epstein argues against emergentists that either

A. *There is no emergence* since the outcomes of any agent-based model—although they are obviously not aggregative—are always proceeding from the behavior of the agents, so there is no novelty in their so-called emergent pattern compared to the ensemble of all the behaviors and their consequences (regarding to what this final state could count as novel since there is no other collective outcome expected?); or
B. *Everything is emergent*, since "to emerge from the initial rules and configurations" in an agent-based model means to be "generated," which amounts to proceeding from the activation of the rules of all the agents, and this is the case of any final state in any agent-based model.

The trouble here is what "deduction" means: to "deduce" the final state of a set of agents, according to Epstein, is only to let all those agents act according to their rules. But "deducibility" in this sense is taken for granted by everybody, since the very definition of a computer simulation implies that all the agents behave according to their rules, so the final state can always be said to be deduced from the initial conditions. Yet, defending an idea of computational emergence that should contrast with deducibility, one will highlight another sense of deduction (noted *deduction**), namely, the deduction of one final global state of the system from previous global states: the incompressibility criterion means precisely that there is no possibility of such a deduction*. Using the work of Buss et al. (1992), we just saw (in Section 3) that such cases in general exist.

So all states are trivially deducible, and therefore, in agreement with Epstein's claim A, that there is no emergence *if* emergence means the opposite of deducibility. But emergence computationally understood excludes *only* deducibility*. In essence, Epstein argues: either emergent = generated = deducible, or emergent is opposed to deducible; because everything is generated in agent-based models, emergence is either trivial

https://doi.org/10.1086/596777 Published online by Cambridge University Press

or impossible. But since "emergent" is indeed only opposed to "deducible*," which is a subset of "deducible," the alternative fails. It is possible that some phenomena might be deduced = generated, but not deducible*, and are thereby emergent; indeed it is the case of many agent-based models encountered up to now—fashion cycles (Tassier 2004), traffic jams (Rasmussen and Nagel 1994), and local norms (Burke et al. 2006). Thus, emergence in the sense of the incompressibility criterion can be applied in agent-based modeling.

Reciprocally, *if "emergent" were a trivial category, all that in agent-based models is deducible would be not-deducible**; but this is false. Suppose, for instance, that we have an initial blinking distribution, and a rule that states "do the opposite of what your immediate horizontal neighbors do (and in case they have between them opposite behaviors, don't move)." It is clearly the previous case of blinking pattern b (Figure 1)—and so there exists a shorthand to compute the state of the system at n, which makes this state deducible*. Thereby, emergence in ABM—deducibility with no deducibility*—is not trivial.

7. Conclusion. The computational notion of emergence, embedded in the incompressibility criterion and naturally focusing on processes rather than properties, meets the two requirements of nontriviality and of scientific adequacy better than combinatorial views. In both fields of agent-based modeling and cellular automata, one can formally describe a class of emergent processes that is limited and broadly corresponds to our linguistic inclination to talk of emergence in those cases.

REFERENCES

Atay, F., and J. Jost (2004), "On the Emergence of Complex Systems on the Basis of the Coordination of Complex Behaviours of Their Elements", *Complexity* 10 (1): 17–22.

Bar Yam, Y. (2004), "A Mathematical Theory of Strong Emergence Using Multiscale Variety", *Complexity* 9 (6): 15–24.

Bechtel, W., and R. Richardson (1992), "Emergent Phenomena and Complex Systems", in A. Beckermann, H. Flohr, and J. Kim (eds.), *Emergence or Reduction?* Berlin: de Gruyter, 257–287.

Bedau, M. (1997), "Weak Emergence", in James Tomberlin (ed.), *Philosophical Perspectives: Mind, Causation, and World*, vol. 11. Oxford: Blackwell, 375–399.

Burke, M., G. Furnier, and K. Prasad (2006), "The Emergence of Local Norms in Networks", *Complexity* 11 (5): 65–83.

Buss, S., C. Papadimitriou, and J. Tsisiklis (1992), "On the Predictability of Coupled Automata: An Allegory about Chaos", *Complex Systems* 5: 525–539.

Chalmers, D. (2006), "Strong and Weak Emergence", in P. Clayton and P. Davies (eds.), *The Re-emergence of Emergence*. Oxford: Oxford University Press, 244–256.

Crane, T. (2001), "The Significance of Emergence", in C. Gillett and B. Loewer (eds.), *Physicalism and Its Discontents*. Cambridge: Cambridge University Press, 207–224.

Crutchfield, J., and J. Hanson (1993), "Turbulent Pattern Bases for Cellular Automata", *Physica D* 69: 279–301.

https://doi.org/10.1086/596777 Published online by Cambridge University Press

Crutchfield, J., and C. Shalizi (2001), "Pattern Discovery and Computational Mechanics", arXiv:cs/0001027v1.
Dennett, D. (1991), "Real Patterns", *Journal of Philosophy* 88 (1): 27–51.
Dessalles, J. L., and D. Phan (2005), "Emergence in Multi-agent Systems: Cognitive Hierarchy, Detection, and Complexity Reduction", in P. Mathieu, B. Beaufils, and O. Brandouy (eds.), *Artificial Economics*. Lecture Notes in Economics and Mathematical Systems 564. Berlin and New York: Springer, 147–159.
Epstein, J. ([1999] 2007), "Agent-Based Computational Models and Generative Social Science", in *Generative Social Science: Studies in Agent-Based Computational Modeling*. Princeton, NJ: Princeton University Press, 4–46.
Gilbert, N. (2002), "Varieties of Emergence". Paper presented at the Social Agents: Ecology, Exchange, and Evolution Conference, Chicago (http://www.soc.surrey.ac.uk/staff/ngilbert/ngpub/paper148_NG.pdf).
Hanson, J., and J. Crutchfield (1997), "Computational Mechanics of Cellular Automata: An Example", *Physica D* 103: 169–189.
Hordijk, W., J. Crutchfield, and M. Mitchell (1996), "Embedded Particle Computation in Evolved Cellular Automata", in T. Toffoli, M. Biafore, and J. Leao (eds.), *PhysComp'96*. Cambridge, MA: New England Complex Systems Institute, 153–158.
Humphreys, P. (1997), "How Properties Emerge", *Philosophy of Science* 64: 53–70.
——— (2004), *Extending Ourselves*. New York: Oxford University Press.
——— (2008), "Synchronic and Diachronic Emergence", in Huneman and Humphreys 2008, 431–442.
Huneman, P. (2008), "Emergence and Adaptation", in Huneman and Humphreys 2008, 493–520.
Huneman, P., and P. Humphreys, eds. (2008), "Special Section on Dynamic Emergence and Computation", special issue, *Minds and Machines*, vol. 4.
Klee, R. (1984), "Microdetermnisms and Concepts of Emergence", *Philosophy of Science* 51: 44–63.
Nagel, K., and K. Rasmussen (1994), "Traffic at the Edge of Chaos", in R. Brooks (ed.), *Artificial Life IV*. Cambridge, MA: MIT Press.
Newman, D. (1996), "Emergence and Strange Attractors", *Philosophy of Science* 63: 245–261
O'Connor, T. (1994), "Emergent Properties", *American Philosophical Quarterly* 31: 91–104.
Schelling, T. (1969), "Models of Segregation", *American Economic Review* 59 (2): 488–493.
Seager, W. (2005), "Emergence and Efficacy", in C. Erneling and D. Johnson (eds.), *The Mind as a Scientific Object between Brain and Culture*. Oxford: Oxford University Press, 176–192.
Shalizi, C., R. Haslinger, J. B. Rouquier, C. Klinkner, and C. Moore (2006), "Automatic Filters for the Detection of Coherent Structures in Spatiotemporal Systems," ArXiv CG/0508001.
Silberstein, M. (2002), "Reduction, Emergence and Explanation", in M. Silberstein and P. Machamer (eds.), *Blackwell Guide to the Philosophy of Science*. Oxford: Blackwell, 80–107.
Silberstein, M., and J. McGeever (1995), "The Search for Ontological Emergence", *Philosophical Quarterly* 49: 182–200.
Tassier, T. (2004), "A Model of Fads, Fashions and Group Formations", *Complexity* 9 (5): 51–61
Wilson, J. (forthcoming), "Non-reductive Physicalism and Degrees of Freedom", *British Journal for Philosophy of Science*.
Wimsatt, W. (1997), "Aggregation: Reductive Heuristics for Finding Emergence", *Philosophy of Science* 64 (Proceedings): S372–S384.
——— (2007), "Emergence as Non-aggregativity and the Biases of Reductionisms", in *Re-engineering Philosophy for Limited Beings: Piecewise Approximations to Reality*. Cambridge, MA: Harvard University Press, 274.
Wolfram, S. (1984), "Universality and Complexity in Cellular Automata", *Physica D* 10: 1–35.

https://doi.org/10.1086/596777 Published online by Cambridge University Press