Computational and Conceptual Emergence

Paul Humphreys[†‡]

A twofold taxonomy for emergence is presented into which a variety of contemporary accounts of emergence fit. The first taxonomy consists of inferential, conceptual, and ontological emergence; the second of diachronic and synchronic emergence. The adequacy of weak emergence, a computational form of inferential emergence, is then examined and its relationship to conceptual emergence and ontological emergence is detailed.

1. Introduction. It is useful to consider two complementary taxonomies for emergence. The *relational taxonomy* is based on relations between the emergent entity and the entities with respect to which it is emergent. There are three divisions in this taxonomy: the *inferential* approach, the *conceptual* approach, and the *ontological* approach.

1.1. The Relational Taxonomy. The inferential approach holds that an entity, such as a state or a property instance, is emergent with respect to a domain D if and only if it is impossible, on the basis of a complete theory of D, to effectively predict that entity or to effectively compute a state corresponding to that feature. The terms 'complete theory', 'impossible', and 'effectively' serve as parameters in this approach that, when specified, produce a particular inferential account. To avoid an inferential account being too easily satisfied, it should be accompanied by arguments that the domain D has some special features, such as universality (as when D is the physical domain) or closure (as when D is the domain of a well-established science). Domains specific to a computational model can also be used. The primary relation involved holds between the emergent feature

†To contact the author, please write to: Corcoran Department of Philosophy, 120 Cocke Hall, University of Virginia, Charlottesville, VA 22904-4780; e-mail: pwh2a@virginia.edu.

‡Conversations and correspondence with Mark Bedau, Philippe Huneman, and Cyrille Imbert have much improved this paper.

Philosophy of Science, 75 (December 2008) pp. 584–594. 0031-8248/2008/7505-0047$10.00
Copyright 2008 by Paul W. Humphreys. All rights reserved.

https://doi.org/10.1086/596776 Published online by Cambridge University Press

and the predictive basis. When the complete theory involved is one of limit science, this may count as a reason to speak of the underivability as absolute, but logically, the attribution of emergence is still relational. A classic account of this type can be found in C. D. Broad's 1925 book *Mind and Its Place in Nature*,[1] where the impossibility condition is taken to be the impossibility of a formal derivation. Broad's account was formulated before the development of completeness theorems and formal theories of computability, and it is, forgivably, unspecific in certain respects. More precise accounts of the relevant derivational relations underlie certain computationally based treatments of emergence, one of which, weak emergence, is discussed below.

The conceptual approach to emergence often results from a situation in which some form of inferential emergence occurs. It maintains that an entity, such as a state or a property, is conceptually emergent with respect to theoretical framework F if and only if a conceptual or descriptive apparatus that is not in F must be developed in order to effectively represent that entity. This apparatus may include new predicates, a new law, or an entirely new theory, and the introduction of the new concepts will sometimes involve moving to a new scientific domain. The phrase 'effectively represent' serves as a parameter for these approaches and may involve efficient prediction or description. For instance, the need to introduce the term 'liquid' in order to effectively describe the behavior of a large collection of water molecules that has undergone a transition from the solid phase to the liquid phase is an example of conceptual emergence.[2] The primary relational aspect involved in conceptual emergence is one of conceptual irreducibility between the new and the old theoretical frameworks. Both the inferential approach and the conceptual approach apply to representations of the phenomena rather than directly to the phenomena themselves.

In the last part of the relational taxonomy, the ontological approach considers emergent entities to be genuinely novel features of the world itself, where an entity is ontologically emergent with respect to domain D if and only if that entity is ontologically irreducible to entities in domain D. Candidates for ontologically emergent phenomena include entangled

1. "The only peculiarity of [a transordinal law] is that we must wait till we meet with an actual instance of an object of the higher order before we can discover such a law; and that we cannot possibly deduce it beforehand from any combination of laws which we have discovered by observing aggregates of a lower order" (Broad 1925, 79).

2. One concise statement of conceptual emergence is: "This, then, is the fundamental philosophical insight of twentieth century science: everything we observe emerges from a more primitive substrate, in the precise meaning of the term 'emergent', which is to say obedient to the laws of the more primitive level, but not conceptually consequent from that level" (Anderson 1995, 2020; quoted in Castellani 2002, 265, note 6).

https://doi.org/10.1086/596776 Published online by Cambridge University Press

states in quantum mechanical systems and, in some approaches, properties that supervene on, but are not reducible to, properties in *D*. The former examples have the distinctive feature that the states of the individual systems before interaction do not determine the joint state after the interaction, whereas the joint state does determine the individual states. This feature makes compositional accounts of the joint system implausible. The primary relata in the ontological approach are the emergent feature and the features from which it emerges, but the relation can be one of functional irreducibility (Kim 1999), supervenience (McLaughlin 1997), fusion (Humphreys 1997b), or some other relation.

The three approaches are not mutually exclusive, although between them they do cover every contemporary account of emergence of which I am aware, except for mystical views that consider emergent phenomena to be immune from understanding. The fact that elements falling into each of these categories are relational is logically important, for rather than the criteria for emergence being applicable to an entity *E* in isolation, such as consciousness, one must consider a domain *D* or a framework *F* with respect to which *E* is emergent, in the sense that it is the relevant relations between *E* and *D* or *F* that determine the emergent status of *E*. The fact that some of these relations are negative relations, such as irreducibility, does not affect this general point.

I have not included in this taxonomy the approach that considers emergent features to be unexplainable because a striking characteristic of most contemporary scientific candidates for emergent phenomena is that we do have credible explanations of why those candidates have their emergent features. They are thus essentially different from the examples of emergence cited in earlier eras, such as life, and they are also different from some contemporary candidates such as consciousness. No 'natural piety' is needed to accept modern, scientifically based, examples of emergence.

1.2. The Temporal Taxonomy. The second taxonomy, the *temporal taxonomy*, appeals to the differences between diachronic and synchronic emergence and is easier to articulate. Diachronic emergence primarily, but not exclusively, involves the emergence of novel phenomena from preceding phenomena as a result of a temporally extended process; synchronic emergence involves the simultaneous coexistence of novel 'higher-level' objects or properties with objects or properties existing at some 'lower level'. This division cuts across the tripartite division of the relational taxonomy; each one of those three divisions has diachronic and synchronic variants. For example, synchronic conceptual emergence can occur in theoretical contexts within which, according to a familiar tradition, Nagel's inhomogeneous reducibility (Nagel 1961, Chapter 11) cannot be carried out, whereas, when a material makes the phase transition to ferro-

https://doi.org/10.1086/596776 Published online by Cambridge University Press

TABLE 1.

	Inferential Unpredictability	Conceptual	Ontological
Synchronic	Broad 1925	Anderson 1972	van Cleve 1990
Diachronic	Bedau 1997	Rueger 2000	Humphreys 1997b

magnetism, the brief but temporally extended relaxation into the global state given by **M** involves a diachronic process, the end state of which requires the introduction of a new concept, the long-range order, to fully describe. Supervenience accounts of emergence (e.g., van Cleve 1990; McLaughlin 1997) are usually synchronic; a diachronic account of conceptual emergence is given in Rueger (2000). An example of each dual category is given in Table 1.

2. Weak Emergence. A contemporary example of a computationally based inferential approach is the concept of weak emergence that has been developed by Mark Bedau (Bedau 1997, 2003): "Assume that P is a nominally emergent property possessed by some locally reducible system S. Then P is weakly emergent if and only if P is derivable from all of S's micro facts but only by simulation" (Bedau 2003).[3] Weak emergence is a diachronic, inferential approach to emergence, and it is most easily understood with the aid of some examples. Consider a two-dimensional cellular automaton. The cells on an infinite discrete grid[4] are labeled by their coordinates (*i, j*), and each possesses a discrete *N*-valued property (for concreteness, think of these states as colors), the value of which constitutes the state of the cell. The state of every cell is simultaneously updated at each time step by a deterministic rule that is a function of the current state of the cell and the states of its immediate neighbors.[5] Over a wide range of initial states of the cellular automaton, appropriate updating rules can produce randomly distributed arrays of differently colored cells, stable patterns that persist across time, and dynamic patterns that evolve over time.

Weak emergence fits into the inferential tradition because if a state of

3. An earlier account of emergence based on the unavoidability of simulation is Darley (1994).

4. Finite grids can emulate infinite grids by using periodic (wraparound) boundary conditions.

5. Using all eight immediate neighbors gives the Moore neighborhood. Von Neumann neighborhoods, in which the four diagonal neighbors do not influence the active cell, are not essentially different. The difference between sequential cellular automata in which the states of the cells are updated asynchronically and regular cellular automata in which cell states are updated simultaneously becomes relevant in Section 3.3.

https://doi.org/10.1086/596776 Published online by Cambridge University Press

a system is weakly emergent, then the existence of that state can be discovered for the first time only through a step-by-step simulation of the system, so that the system must go through all the intermediate developmental states before the final state is displayed. In other words, the process that leads up to the state is computationally incompressible. In yet other words, unlike the prediction of future solar eclipses for which the computational difficulty of prediction is almost independent of how far into the future the eclipse will take place, predictions of future states of computationally incompressible systems must run through each of the intermediate time steps between the initial state and the predicted state. Letting the computational model work out its own development is thus the only effective way to discover how the system's states evolve. Adapting the approach in Rasmussen and Barrett (1995), the computational incompressibility aspect of weak emergence can be represented in the following way. Consider the kinematics of a system S that is modeled as a temporally discrete process. An updating function U takes the current state of the process and transforms it into a future state. There is then an essential difference between these two cases:

$$S(t + T) = U[S(t)], \tag{1}$$

for arbitrary T, and

$$S(t + T) = U_T[U_{T-1}[\cdots[U_1[S(t)]]\cdots]], \tag{2}$$

where U_i is the updating operator at time step i. In the usual case in which U_i is fixed, we have $S(t + T) = U^T[S(t)]$.

3. Weak Emergence, Conceptual Novelty, and Some Reasons Why Weak Emergence Is Different from Conceptual Emergence. The definition of weak emergence succinctly captures a great deal of what is considered emergent within the fields of dynamical systems theory and complexity theory, but one aspect of it is controversial. As Cyrille Imbert (2006) has noted, Bedau's definition places no restriction on the kind of end state that can weakly emerge from a simulation except that it be nominally emergent, where a nominally emergent property is one that can be possessed at the macro level but not at the micro level. Employing a number of illuminating examples, Imbert argues that this is a defect in the definition of weak emergence and goes on to argue that the additional criteria for genuinely emergent states generated by computationally incompressible processes must be contextual and based on the immediately preceding states in the generating process. Intuitions and arguments can pull us in different directions on this issue, and rather than argue for or against the adequacy of weak emergence as originally presented, I want to use the

https://doi.org/10.1086/596776 Published online by Cambridge University Press

contrast between Bedau's and Imbert's positions to emphasize the logical independence of weak emergence from conceptual emergence. The issue of whether to impose restrictions on the kinds of outputs generated by weakly emergent systems then turns out to involve considerations that are based on significantly different conceptions of emergence.[6] Some preliminary arguments are needed before we address the core issues involved.

3.1. Weak Emergence Is about Instances. In Humphreys (2008), I provide an argument that diachronic emergence concerns property instances and not types. Here is a different argument for that conclusion, formulated specifically for weak emergence. Consider a process within which weak emergence takes place over the first N time steps, the Nth state being a token of a property type P that has not previously been instantiated in the process. The rules governing the development of the process then regenerate that state at later stages in a way that is predictable in a highly compressed way; one simple case occurs when the system remains in state P after time step N, and this steady state is predictable from the rules once state P has been reached. In such a system, the property instance at the Nth state is weakly emergent, but no further instances of P count as weakly emergent because the derivation of the $2N$th state from the Nth state is trivial. Contrast this with a system having a different set of rules that generates a state of type P after N steps and in which another token of property type P is produced by a computationally incompressible process at the $2N$th step. For simplicity, we can suppose that the system reverts to its initial state at step $N + 1$ and then cycles through the same sequence of states so that no states of type P occur between the Nth and $2N$th time steps. In this second case, both instances of P are weakly emergent. This shows, from a specifically diachronic perspective, that, in addition to information about the state type itself, one needs information about the process leading up to a token state in order to determine whether that token is emergent. That is, we can have two instances of the same property type, one instance of which is weakly emergent and the other

6. Bedau did address this issue in his 1997 article. As he wrote, "One might object that weak emergence is *too* weak to be called 'emergent', either because it applies so widely or arbitrarily that it does not demark an interesting class of phenomena, or because it applies to certain phenomena that are not emergent. For example, indefinitely many, ad hoc [Game of] Life macrostates are (for all we know) underivable without simulation . . . . But this breadth of instances, including those that are arbitrary or uninteresting to 'emergence theorists', is not a problem or flaw in weak emergence" (Bedau 1997, 395). Bedau has confirmed in conversation that this is still his position.

https://doi.org/10.1086/596776 Published online by Cambridge University Press

of which is not.[7] This is in sharp contrast to conceptual emergence. It is property types that are conceptually emergent because definitional reducibility and irreducibility apply to predicates in isolation from their instances. (Specifically, singular features such as tropes or haecceities seem to be best treated as primitives rather than as emergent, but it may well be that emergent objects, taken here to be concrete entities denoted by uniquely referring terms, require some form of conceptual emergence for a proper representation.) For present purposes, the important point that follows from the primacy of instances is that we can have two instances of a given type of state within a process and both instances are weakly emergent.

3.2. Novelty. Now consider a criterion that is often imposed on emergent features, that they be novel. The sense of novelty that is involved in inferential accounts of emergence is different from that used in conceptual emergence, and it is intertwined with the unpredictability requirement. First, as Imbert has emphasized, what counts as novel in the case of weak emergence has to be decided with respect to property instances occurring within the process that leads to E's appearance. To require that a property instance be absolutely new is far too strict, for then only the first ever instance of a property would count as emergent. Next, conceptual emergence always requires supplementation or replacement of the original theoretical framework to represent the emergent state, and this sometimes takes us out of the original scientific domain entirely. Weak emergence imposes no such condition, and a system that produces recurring instances of the same property whose instances are produced by a computationally incompressible process is nevertheless exhibiting repeatable weakly emergent states. The novelty consists in the impossibility of having noninductive advance knowledge of those states within the specified theoretical domain. In terms of the Rasmussen and Barrett representation, there is no single updating operator U' that can replace U^T.

3.3. Random to Random. We are now in a position to assess one argument in favor of imposing a further structural condition on the definition of weak emergence. Suppose that the system transitions via a com-

7. This is what one should expect from any element in the relational taxonomy. Talk of a property or property instance being emergent tends to suggest that emergence is an intrinsic feature of a property or its instances. Yet as soon as we keep in mind that on both the synchronic and the diachronic approaches an emergent entity must emerge *from* something, we see that this reflects the fact that emergence is unavoidably a relational property and that unless the appropriate relation holds between the property under consideration and its origins, tokens of the former can occur without their being emergent.

https://doi.org/10.1086/596776 Published online by Cambridge University Press

putationally incompressible process from an initial random configuration to a final random configuration. If all the intermediate states are nonrandom, then the argument of Section 3.1 establishes that the final state is weakly emergent. But what if all the intermediate states are also random? Although random patterns are token distinct, qua random, they are type identical, and so no new property type has been produced when one random array is transformed into another; this is something that undoubtedly underlies the hesitation in allowing that any kind of emergence has occurred in this situation. Furthermore, it looks like none of the successive random states are weakly emergent because the 'derivation by simulation' involves only one step. But this will not be true in all cases. Synchronic updating of cellular automata (and of agent-based models as well) rarely involves a single transparent derivational step; what occurs in the single time step is not only the computational updating of the state of each cell but frequently the production of a global state that is reachable only by simulation and not by traditional closed-form prediction. Finally, although these intermediate states might seem to violate the novelty criterion for emergence, we have seen that this violation is apparent only because weak emergence can allow repeated instances of the same property as long as the computational connections between the instances are of the appropriately complex kind, as they usually will be here. Hence a process within which all states are random can be a candidate for producing weakly emergent states.

3.4. Ordered to Random. Now suppose that the system makes the transition from an initially structured state to a final random state. An appeal to the motivations behind considering self-organizing systems as producers of emergent states might in this case incline us toward insisting that this should not count as a case of emergence. The basic idea behind self-organization is that large-scale structure spontaneously emerges in a dynamic system of interacting constituents (i.e., the structure emerges without the need for external interventions) solely in virtue of local interactions between those constituents, and this structure is not accidental in the sense that were the system to be restored to its initial state or to a qualitatively equivalent state and the microdynamics rerun, a similar large-scale structure would, with a high degree of probability, reappear.[8] It is the apparent ability to run counter to the general trend produced by the Second Law of Thermodynamics that makes the transition from disorder to order so interesting. Order emerging from disorder strikes us as important because

8. Self-organization is almost always a statistical property of a system. In many models, unusual initial conditions will prevent the usual structures from emerging.

https://doi.org/10.1086/596776 Published online by Cambridge University Press

it is unusual; disorder arising from order as uninteresting because it is so common.

Whether or not one finds this argument a compelling reason to insist that random end states generated from structured states should not count as emergent, and thus that some additional constraint should be imposed on the end state of such processes, it is clear that the argument relies on considerations that are completely separate from those involved in the computational basis of weak emergence.

3.5. Weak Emergence and Ontological Emergence. Finally, what of the connections between weak emergence and ontological emergence? Weak emergence does not entail ontological emergence for a straightforward reason, a reason that makes weak emergence different from older versions of the inferential approach. Discussions of weak emergence are often based on examples using cellular automata, and I followed that tradition above. One example that is often used, the Game of Life, has the great advantage of simplicity, but emergent phenomena in the Game of Life, such as gliders, are pure model phenomena; they do not, and are not intended to, represent features of systems external to the model. This reflects the special sense in which the term 'simulation' is used in the definition of weak emergence: not only is it strongly skewed toward agent-based and related simulations (for a discussion of other types, see Humphreys 2004, Chapters 3 and 4), but there is no requirement that the computational process is a simulation of a real system. The weak emergence approach thus bifurcates into two subgenres. The first considers the possibility of predicting phenomena only within the model; the other is concerned with the possibility of predicting states that have an intended scientific interpretation. One should not restrict ontological emergence to cases that actually occur in the world because it is a legitimate function for theories of ontological emergence to tell us how emergent phenomena could occur in nonrepresentational systems. However, when a simulation is used as a model of another system, our knowledge of that second system may serve as a basis for imposing a further restriction that only states of the simulation that correspond to states of the real physical, biological, or social system can be candidates for weak emergence. It is thus easier to be a purist about weak emergence as a stand-alone criterion for emergence within the first subgenre than it is within the second.

4. The Sparseness Issue. We can draw one further philosophical conclusion from our analysis. What I call the sparse perspective on emergence is probably the dominant perspective in the philosophical literature on emergence. The sparse perspective holds that emergence, if it exists at all, is a rare phenomenon. Indeed, the sparse perspective often seems to be

https://doi.org/10.1086/596776 Published online by Cambridge University Press

used as an informal touchstone for accounts of emergence: if the account makes emergence too common, that is a reason to think the account is defective. A particularly clear version of this sparse view is held by Jaegwon Kim. Having laid out the first plausible account of reduction since Ernest Nagel's version was undermined by multiple realizability arguments, Kim asserted that

> emergentism may yet be an empty doctrine. For there may not be any emergent properties, all properties being physical properties or else functionalizable and therefore reducible to physical properties. . . . So are there emergent properties? Many scientists have argued that certain "self-organizing" phenomena of organic, living systems are emergent. But it is not clear that these are emergent in our sense of nonfunctionalizability. And, . . . , the classic emergentists were mostly wrong in putting forward examples of chemical and biological properties as emergent. It seems to me that if anything is going to be emergent, the phenomenal properties of consciousness, or "qualia," are the most promising candidates. (1999, 18)

The sparse position on emergence seems plausible, I believe, because it tends to focus almost exclusively on synchronic emergence. If weak emergence and other forms of diachronic, computational emergence are legitimate kinds of emergence, then examples of purely diachronic weak emergence will be quite common and the sparse position correspondingly false.

5. Conclusion. I have argued that the three types of accounts in our relational taxonomy of emergence are very different from one another. This then raises the question of whether there is a higher-level framework for emergence that can unify the three types. In Humphreys 1997a, I identified six features that are commonly required of emergent phenomena—novelty, qualitative difference, nominal emergence, applicability of different laws, essential interaction, and holism—and suggested that a cluster analysis based on those features might serve as a bridge between different approaches. There are, no doubt, other, better, frameworks within which all six categories of our two taxonomies fit, but the challenge now is to identify and construct it.

REFERENCES

Anderson, P. W. (1972), "More Is Different", *Science* 177: 393–396.

——— (1995), "Historical Overview of the Twentieth Century Physics", in L. M. Brown, A. Pais, and B. Pippard (eds.), *Twentieth Century Physics*. New York: American Institute of Physics Press, 2017–2032.

Bedau, Mark (1997), "Weak Emergence", in J. Tomberlin (ed.), *Philosophical Perspectives: Mind, Causation, and World*, vol. 11. Malden, MA: Blackwell, 375–399.

https://doi.org/10.1086/596776 Published online by Cambridge University Press

——— (2003), "Downward Causation and the Autonomy of Weak Emergence", *Principia Revista Internacional de Epistemologica* 6: 5–50.
Broad, C. D. (1925), *The Mind and Its Place in Nature*. London: Routledge and Kegan Paul.
Castellani, Elena (2002), "Reductionism, Emergence, and Effective Field Theories", *Studies in History and Philosophy of Modern Physics* 33: 251–267.
Darley, Vince (1994), "Emergent Phenomena and Complexity", in R. Brooks and P. Maes (eds.), *Artificial Life IV: Proceedings of the Fourth International Workshop on the Synthesis and Simulation of Living Systems*. Cambridge, MA: MIT Press, 411–416.
Humphreys, Paul (1997a), "Emergence, Not Supervenience", *Philosophy of Science* 64: S337–S345.
——— (1997b), "How Properties Emerge", *Philosophy of Science* 64: 1–17.
——— (2004), *Extending Ourselves: Computational Science, Empiricism, and Scientific Method*. New York: Oxford University Press.
——— (2008), "Synchronic and Diachronic Emergence", *Minds and Machines* 18: 431–442.
Imbert, Cyrille (2006), "Why Diachronically Emergent Properties Must Also Be Salient", in Carlos Gershenson, Diederik Aerts, and Bruce Edmonds (eds.), *Philosophy and Complexity: Essays on Epistemology, Evolution, and Emergence*. Singapore: World Scientific.
Kim, Jaegwon (1999), "Making Sense of Emergence", *Philosophical Studies* 95: 3–36.
McLaughlin, Brian (1997), "Emergence and Supervenience," *Intellectica* 25: 25–43.
Nagel, Ernest (1961), *The Structure of Science*. New York: Harcourt.
Rasmussen, Steen, and Chris Barrett (1995), "Elements of a Theory of Simulation", in *European Conference on Artificial Life*, Lecture Notes in Computer Science. Berlin: Springer-Verlag.
Rueger, Alexander (2000), "Physical Emergence, Diachronic and Synchronic", *Synthese* 124: 297–322.
van Cleve, James (1990), "Mind-Dust or Magic? Panpsychism versus Emergentism", *Philosophical Perspectives* 4: 215–226.

https://doi.org/10.1086/596776 Published online by Cambridge University Press