

The Hierarchic Logic of Emergence: Untangling the Interdependence of Evolution and Self-Organization.¹

Terrence W. Deacon, Department of Anthropology, University of California, Berkeley

Chapter 14 in: Bruce Weber & David Depew (eds.) *Evolution and Learning: The Baldwin Effect Reconsidered*. MIT Press, 2003

1 Introduction

". . . something more from nothing but."²

Biological evolution is characterized by a collection of highly convoluted processes that produce a remarkably complex kind of combinatorial novelty. The contention of this paper is that biological evolution and evolutionary processes in general are a subset of processes drawn from a much larger set of novelty-producing processes that also includes self-assembly and self-organizing processes. Not only does it appear that these are related concepts, I think it is also clear that they are interdependent in complex and subtle ways that have yet to be fully delineated, especially in the processes of life and mind. It is also suspected by many writers that a synthesis that successfully integrates the logic of these various kinds of creative processes will do more than significantly advance our understanding of how life came about and how thoughts and experiences are generated. It could possibly also provide new insights into the very nature of physical causality. But there are some broad theoretical issues that stand in the way of this outcome. These issues derive from a set of unresolved problems about the nature of physical novelty itself, and how we conceive of its origination in terms of current theories of causality. These most enigmatic physical phenomena all have something to do with creative or originative processes in nature, and for this reason seem inevitably to come in conflict with our otherwise quite successful reductionistic account of most other aspects of the world.

A more general term often used to describe this larger class of spontaneous, and only weakly predictable, order-generating processes is "emergence."³ This is a promising abstract explanatory concept, but one that is at risk of becoming overused and too vague for any technical purposes, precisely because of its generality and only partially specified meanings. The purpose of this essay is to take the concept apart and to attempt to discern what (if any) features about physical causal processes it accurately reflects, so that it contributes to the empirical investigation of biological and mental processes, and not just to philosophizing about them.

The concept of emergence probably has gained its worst reputation when it has been used in a primarily negative sense, i.e. to point to something missing in reductionistic explanations. In explicitly anti-reductionistic criticisms of standard accounts of such phenomena as life and mind, it has come to be a code word identified with a complex systems theoretic perspective. In this use, the concept of emergence is a place-marker intended to indicate points where standard reductionistic accounts fail or seem incompletely to explain apparent discontinuities in properties exhibited at different levels of physical scale. This negative usage has unfortunately led many more orthodox thinkers to suspect

that there is no underlying phenomenon to be described, only a vague suspicion due to incomplete analysis. On the other hand, in examples where it has been more precisely described (e.g., the emergence of liquidity or surface tensions from the interactions of water molecules), it is seen as adding nothing of empirical significance to standard physical reductionistic accounts. And finally, where it is used to describe more complex phenomena (e.g., emergence of life or mind), the details and logic are sufficiently obscured by incomplete scientific investigation to be of much use. Incautious uses allow critics to rightfully claim that it mostly serves only as a sort of philosophically motivated promissory note.

Nevertheless, the term emergence has become a kind of signal for research paradigms sensitive to systemic factors. There is growing awareness among biologists, physicists, and computation scientists studying diverse kinds of complex phenomena that many of these share a curious general feature in common. Complex dynamical ensembles can spontaneously assume ordered patterns of behavior that are not prefigured in the properties of their component elements or in their interaction patterns. Moreover, unprecedented global forms can develop along parallel lines to reach similar patterns of behavior despite arising from components of radically different constitution, interacting according to quite diverse physical principles. Thus, retrospectively at least, the "cause" of such unprecedented forms appears to be attributable to properties that can only be described at the level of the whole ensemble. Systems that behave this way are often described as "self-organized," since the most relevant antecedent condition is at the same level of description and scale as the consequent ordered behavior.

The appearance of this general form of systemic causal dynamic in diverse domains motivates the suspicion that phenomena like this may be related in some underlying ways to one another, independently of their material composition and interactions. There is also a sense, however, that living phenomena exhibit a mode of unpredictability in their self-organization that is something more than mere self-organization. It is also becoming clear that the kinds of processes we call "evolutionary" are potentially quite diverse and of varying levels of complexity. In this volume alone, diverse views about the relative importance of the Darwinian account of natural selection are expressed, with a few authors significantly departing from the orthodoxy. Others (including this author) suggest that, in cases of complex organisms able permanently to alter their environment or, in cases where higher-order evolutionary phenomena (like language and culture) contribute influences to biological evolution, it becomes difficult to distinguish where evolution leaves off and self-organization begins.

The term "emergence" connotes the image of something coming out of hiding, coming into view for the first time, something without precedent and perhaps a bit surprising. But is there some new physical quality emerging in the physical processes we are considering or merely some change that allows us to see something that previously was in a cryptic form and is now unmasked? In other words, we need to be clear about whether we are investigating a limitation of scientific description or something about actual physical processes in the world. There will always be gaps in our understanding. Phenomena that we cannot

predict often merely reflect limitations of theory or of modeling and computing power. It should come as no surprise that we encounter difficulties producing commensurate descriptions that are adequate at very different levels of scale. And something that defies continuity of description with respect to current tools across levels of scale is an artifact of science, a descriptive discontinuity, not a causal one.

Many scientists suspect that most things described as emergent phenomena are actually matters of descriptive inadequacy across levels of scale. To the extent that new kinds of descriptive methodologies are forced upon us by the scale-dependent features of our analytical tools, the novelty we encounter can be described as epistemological emergent. The question is whether there is something we might, in contrast, call ontological emergence. We need to distinguish between recognizing a descriptive discontinuity and the generation of true physical novelty. This distinction has been made by many others, and is at the heart of a long-running philosophical debate. I enter this discussion with the hypothesis that true physical novelty can be generated by many natural processes including evolutionary and mental processes, but also with the suspicion that we tend to look for the wrong sort of novelty. So I take my challenge to be the task of trying to be as explicit and concrete as I can be about what "emerges" in these cases. This also means trying to make explicit the connection between emergent processes and those we think of as plain old garden-variety mechanical processes.

A good place to start is to ask "what" emerges? The answer is not some "thing" but rather something like a form, or pattern, or function. The concept of emergence seems to apply to phenomena in which relational properties tend to dominate over constituent properties in determining aggregate features. To some extent the modern conception of physical "cause" tends to combine and confuse what might be called constituent or material features of things and events with what might be called topological or relational features.⁴ This contributes to confusion about emergence because it is with respect to the configurations and topologies, not the specific properties of constituents that we trace processes of emergence. Global topological properties are often shared by many systems irrespective of any particular physical laws or any particular constituent properties. Correspondingly, it is the spontaneous, unprecedented production of new relational properties that constitutes emergence, not the production of new kinds of substance or physical law.

2 Synchronic and Diachronic Senses of Emergence

The conception of novelty implicit in all senses of emergence can implicitly or explicitly involve a "before and after" comparison, even when the major focus is on phenomena involving differences of scale. Emergent phenomena are always defined with respect to time in some regard, but the constitutive role played by temporal change differs with different senses of emergence. As a further preliminary, we need to distinguish forms of emergence in which development in time leads to a kind of self-similarity across time, allowing us to treat it as synchronic, and self-similarity that is not in the outcome, but instead is in the developmental process from moment to moment. This requires us to consider it as irreducibly diachronic.

Synchronic uses of the concept of emergence are usually restricted to properties of matter that become evident with ascent in scale. Consider water: H₂O molecules in large aggregates in the right conditions exhibit a set of dynamical properties we refer to as liquidity. These higher-order properties depend only on the specific properties of the constituent molecules determining what might be called parameters of liquidity (e.g., viscosity, vaporization point, etc.), but in some sense not for the generic property of liquidity itself. At appropriate temperatures and pressures and time domains (e.g., glass is a "slow," i.e., viscous liquid), many different kinds of molecular constituents can exhibit the same general behavioral properties. But liquid properties are only one of a few phases of matter (solid, gas, and plasma being others) and transitions from one to another can occur with changes of boundary conditions affecting the aggregate. Specific details of the boundary conditions are determined by the molecular constituents, but, again not the existence of these phases themselves. A general characteristic of all types of molecular constituents is that they can exhibit these phases under the right conditions. So we can consider these phasic behavioral regimes to be generic properties irrespective of the material. Whereas the specific conditions at which phase transitions occur depend on the material properties of the constituents, the form of these phases depends on something more, and yet not more. Phase "states" are a sort of oxymoron in a microscopic sense since they are dynamic processes. But they are processes with a special kind of temporal topology. They are "attractors," in the lingo of complexity theory: highly self-similar ensembles of configurations of the global dynamics that tend to lead to one another "neighboring" configurations with very high probability. And these attractors are separated from each other by unstable intermediate transients.

The existence of multiple regular higher-order attractors means that they can come into existence and fall out of existence over time. One might be tempted to treat the coming into being of a new phase as emergent in a different, diachronic sense than when self-similar dynamics (in effect) cycle. But this would ignore the fact that in some sense the new configuration of the material remaining in the same phase is also being diachronically "generated" anew, and so forth, in each ensuing microsecond. This might suggest that it is an artificial distinction to say that, for example, the emergence of liquidity is synchronic whereas the evolution of life on earth is diachronic. The first unfolds in time as well. But I want to argue that in this comparison an important distinction about this unfolding is missed.

Evolution is a historical phenomenon in ways that liquidity isn't. I would even want to say that evolution is, in some sense, "more" emergent than liquidity. Evolution involves time in a convoluted way. It is not self-similar in any simple way. A raindrop analyzed two hundred million years ago and another analyzed five minutes ago would be vastly more similar than two individual vertebrates even from the same lineage. There is self-similarity to life across time but there is also a similar trend of difference. Life and mind are, similarly, more temporally convoluted than liquidity, although in a very vague sense they each are emergent. Attending to this kind of difference in mode of development when comparing different kinds of emergent phenomena is thus important. I believe that failure to do so has been the most consistent reason

this concept has more often obfuscated than clarified the causal logic of the most vexing scientific questions of the past century. Its long use as an explanatory term despite little advance in the sophistication of its scientific use appears to attest to this. (See for example its early use in Lewes's "Problems of Life and Mind.")

Francis Crick, co-discoverer of the structure of DNA, calls the realization that thoughts, feelings, passions, and so forth, are the results of chemical processes in the brain an "Astonishing Hypothesis." One has the suspicion, however, that the astonishing part of this is to be found in what is not stated in Crick's reductionistic claim. Mind can't be "nothing but" this, and yet isn't exactly "something in addition" to these chemical processes either. But here is where an indiscriminate use of the concept of emergence can be unhelpful. Many who feel uncomfortable with the "nothing but" characterization express this by arguing along the following lines: "Consciousness is an emergent property of the cellular and molecular processes within a brain in the same sense as surface tension is an emergent property of the interactions of water molecules."⁵ By invoking this analogy, an implicit claim about a presumably well-understood relationship is meant to inform our understanding of the general logic that can be applied to a much less well-understood (indeed deeply mysterious) relationship. Unfortunately, it is precisely the huge number of unstated details on both sides of the analogy that turns this into an exercise in question-begging.

What motivates the analogy is that both surface tension and consciousness are in some sense compositional. Both, too, involve a contrast between levels of description for which quite different descriptive tools are required. But that fact is pretty self-evident. The question is whether these are analogous in any other way. By hypothesis, water molecule collisions (obeying lower-order component properties) stand in relation to surface tension as neuro-chemical reactions stand in relation to conscious experience: i.e. both exhibit higher-order aggregate properties. Without question, activities supported by whole living brains must at the same time be supported by the component processes of these brains, since the brain is also a composite or aggregate dynamic entity. Beyond that of mere aggregate form, however, we are left with little confidence that the kind of causal relationships implicitly being compared allows for any further extension of the analogy. It is a bit like the comparison of the production of liquidity to the evolution of living forms.

It becomes obvious that we are actually comparing different kinds or modes or levels of emergence, however, when we consider comparing different properties. Notice, for example, that we would be on pretty solid ground comparing the viscosity or the opacity of brains to the relative viscosity or opacity of a liquid, and we would feel confident treating both cases as similar forms of emergent properties from their molecular interactional properties. It is apparent that we implicitly recognize that there can be different "species" of emergence. And if there are, then cross-category comparisons of emergent relationships, such as this, may not be as informative as we might hope.

Whatever we mean by "emergence," there can be no doubt that mental phenomena are emergent from the subordinate neuro-chemical interactions occurring in a brain in a more complex way than liquid phenomena are emergent

from water molecule interactions. The fact that we can identify similar kinds of properties in brains to those of liquids (e.g., viscosity) also suggests that the brain-mind kind of emergence might at least be described as of a higher order. The same can be said of life as an emergent phenomenon. Living metabolism is not emergent from its constituent chemical processes in the same way that a flame's dynamical features are emergent from its component chemical processes, even though a metabolic process shares some features in common with burning. The critical features that make metabolism emergent, and sets it apart from, and at least a level above, those features that make flames emergent, can be traced to the fact that in some sense flamelike emergence is one component in lifelike emergence.

At the very least, then, we need to articulate a more elaborate and hierarchically organized taxonomy of emergence relationships that unpacks the temporal complexity of the issue a bit more clearly. This can at least aid us in avoiding these sorts of too-simplistic comparisons. Though I hope to finesse philosophical debate about the reality of irreducible emergence- ontological emergence-I hope to provide explicit positive definitions and categories of emergence to counter the current vagueness in its usage, so that irrespective of whether we believe there is something like "true" emergence, we will be more clear about what we are claiming in either case. Though by describing as precisely as possible a causal logic distinguishing specific examples we risk sacrificing the strong sense of emergence implied by proponents of antireductionism, I think we do not necessarily lose the constructive sense it contributes to explanations of natural creativity that is behind its most important uses. At the very least, an explicit analytic categorization of types of emergence can help to distinguish among theoretical claims that are commonly confused in discussions of evolutionary and mental phenomena.

Although I will offer only a qualitative descriptive account of what I mean by "levels of emergence," I believe it can be made sufficiently precise and unambiguous to be easily rendered in mathematical terms using current theoretical tools.

3 Compound Interest

What is it about certain phenomena we label as emergent that prompts us to consider them novel and unprecedented? It is not the energy that courses through them nor the material they are composed of that matters. It is rather their topology or configuration, that is, how the matter and energy is organized, and how it behaves. Emergence is about the topology of causality. What I believe to be critical for explaining the emergent properties of some system is an account of the origins and development of the configurational or topological regularities of both the constituent interactions and also of whole aggregates. This account must also address the fact that these topological relationships between levels of phenomena can be multiply realized by different physical systems, and are thus neither only nor merely expressions of substrate properties.

A related point is made by researchers simulating complex systemic dynamics in computers and by computational theorists in the cognitive sciences. In many of these model systems, the large scale configurations that are produced more

often reflect configurational properties of the whole ensemble (e.g., number of elements and their interactional connectivity) rather than properties derived only from interactional dynamics of constituents. The point is that similar systemic properties can be exhibited by systems made up of unrelated elements with very different component dynamics. This is a specific case of a much more common phenomenon that is axiomatic in the computer sciences. Radically different machine architectures can "run" the "same" program or algorithm (e.g., a high-level description of some entire class of input-output relations) even if the physical mechanism involved is different in different cases of its running. There is, of course, a mundane, commonsense version of this in the case of arithmetic done with the aid of an abacus, a slide-rule, or with paper and pencil manipulations of scribbles. All of these embodiments can assist us in "doing" the same arithmetic operation in some sense. By analogy, cognitive scientists have argued that the "functional" properties that define a given cognitive operation are like the logical architecture of a computer program. A given thought may also be the same in certain critical respects irrespective of the device that "embodies" it.

Philosophically, this general form of argument is known as functionalism, and although I think there are problems with this view of mind, I nevertheless think the general logic is relevant to the problem of emergence. Claiming the existence and relative autonomy of higher-order ensemble properties of physical systems and arguing that they are constituted by some sort of global influence over the properties and dynamics of system constituents is a claim for a kind of physical-causal analogue to functionalism. In functionalist accounts, only certain substrate details are essential; others may vary without functional consequence. But function is defined extrinsically, at least in computer science and cognitive theories, because function is a semiotic distinction, not a physical one. What is the physical analogue? It is, as I suggested above, something like topology: the form, configuration, or distribution of component features. In these terms, my argument will not be grounded on any more radical assumption than the claim that the same topology can be differently instantiated. This is noncontroversial. But two additional points need to be made about this assumption, and about the analogy I am drawing.

First, topology is not just a descriptive feature of a physical system, it is a constitutive fact about the spatial-temporal relationships among component elements and interactions with intrinsic causal consequences. Second, it is not a single relationship, but a space of relationships, or perhaps more accurately the space of spaces of relationships. There are various levels of abstraction we might use to describe a system configuration. Each might include more or ignore certain details, but the description is irrelevant, it is the way that these topological features contribute to future topological features and influence those in adjacent systems that matters. This propagation of effect is selective and relational and always subject to higher-order topological factors or boundary conditions. What is gained in a physical theory of causality by noting that, under certain reduced descriptions, multiple physical systems with divergent details are isomorphic in some respects? It matters, of course, only if it is not just description, but a matter of dynamical importance. It matters if certain aspects of topology at some level are more relevant for an account of

topology at that or some other level or stage of development. It matters, then, if topology at one level of scale or locus affects topology at some other level or locus, irrespective of the physical substrate. Topology is a global or ensemble relationship. So it may also include whole-to-part influences that are the most cited characteristic of emergent phenomena. The "top-down" character is both the key defining character and yet the most criticized aspect of the concept of emergence. Reframing it in terms of topological influences and constraints is a first step toward removing the confusion of "causality" claims associated with it.

What needs to be explained, then, is not a new form of causality, but how some systems come to be dominated by their higher-order topological properties so that these appear to "drag along" component constituent dynamics, even though, at the same time, these higher-order regularities are also constituted by lower-order interactions. In other words, an explanation of how topologies come to make a difference is required. I believe that the secret to explaining this lies in what can be called amplification processes. I believe that we can understand emergent phenomena as all being variant forms of what might be called topological reinforcement or amplification in pattern formation.

Albert Einstein once quipped that "compound interest" is the most powerful force in the universe.⁶ Perhaps he was thinking of nonlinear field dynamics, like the compounding effects of gravitation, or possibly he was just sardonically reflecting on his house payments. In any case, it is a useful metaphor for explaining many kinds of amplification effects. I will argue that a kind of compound interest of topologies is the basis for all forms of emergence. We are justified in calling something "emergent," I will suggest, if it is the result of a recurrent amplification of configuration or topology. This recurrent architecture is itself a topological concept, so in some sense emergence is a special case of topological transformation of topologies.

Wherever it occurs, amplification is accomplished by a kind of repeated superimposition of similar forms. Amplification can be achieved by mathematical recursion in a computation, by recycling of a signal that reinforces itself and cancels the background in electronic sound amplification or other signal processing circuits, or by repetitively sampling the same biased set of phenomena in a statistical analysis. In each of these cases, the formal or configurational regularities of the medium are what serve as the basis for the amplification. The "stuff" that constitutes the medium in which amplification is exhibited is largely irrelevant, as is the electricity in an audio amplifier. (Audio amplification can for example be achieved using a liquid medium flowing through pipes or light beaming through light guides.) Amplification can be a merely physical process or an informational process—the latter usually depends on the former—as in statistics and certain computations. In all cases, it is a form that is amplified; repetition either multiplies the number of its appearances in some physical medium or else gets it embodied in a progressively larger fraction of the physical medium (as in sound amplification). Thus amplification is a kind of compound interest of physical form.

The relevance of "compound interest logic" in the analysis of emergence is that it helps explain how certain minor or even incidental aspects of a complex phenomenon can come to be the source of its dominant features. In this analysis,

then, emergence will be treated as a form of amplification of certain topological features within a physical system. Specifically, I will argue that amplification occurs because of iterated superimposition of events sharing the same form occurring across levels of scale. Using the logic of compound interest in physically-embodied forms as my guide, I will argue that we can distinguish three hierarchically related classes of emergent phenomena on the basis of their constraint architectures. Specifically, I will suggest that the most useful architectonic feature is whether there is recurrent or circular redundancy of the influence of constraints or biases exhibiting the same form across levels of scale. In the history of the development of theories of complexity and emergence, writers have regularly recognized the important role of "circular causality" in some form or other.⁷ Problems of recursion and self-reference have also regularly been the basis for emergent concepts in mathematics: from the calculus, to the concept of imaginary numbers, to the inspiration for Gödel's incompleteness proof. To date, however, I know of no effort to formalize this intuition with respect to a theory of ontological emergence, or to apply it as a general analytic tool for describing the common features of all forms of emergent phenomena.

4 Circular Definition

Not all cases of circular causality may qualify as full instances of emergence, at least in the sense that it is most often invoked in physical systems. A simple and well-known example of "circular causality" is embodied in a thermostatic control system. By connecting a heating and/or cooling device to a temperature-sensitive switch, the coupled devices can be configured to respond to temperature in a manner that will eventually invert the state of the switch and with it the state of the device that heats or cools the environment. In other words, there is a self-undermining pattern of cause and effect—so-called negative feedback—that tends to produce behavioral oscillation around some set-point. If this causal linkage is reversed, of course, so that deviation away from the set-point activates mechanisms that cause the environmental temperature to deviate yet further, a very different and unstable behavior results—so-called positive feedback. This runaway effect is checked only by outside constraints. Even simple deterministic engineering devices, in which a number of such feedback control devices are coupled, can produce highly complex, quasi-periodic behaviors—or even "deterministic chaos," in which long-term prediction is essentially impossible. So feedback effects of this sort can be quite complicated, and even deeply unpredictable, in complex dynamical systems. They produce higher-order quasi-regularities whose patterning can only be understood in terms of the whole causal system, and are only produced when the whole ensemble of components is assembled.

Emergent effects in most contexts inherit their surprising characteristics from some version of this logic. This basic feedback servomechanism design has a self-undermining architecture in terms of the transformation of energy and the "signals" this energy conveys. This self-undermining logic is a bit like the famous liars' paradox— "This statement is false." In a servomechanism, internal self-contradiction of this kind serves a sort of self-correcting or self-compensating function, although higher-order loops of servomechanisms can

produce complicated mixtures of asynchronous positive and negative feedback effects.

The importance of circularity in positive feedback is that it is deviation amplifying in the same way that compound interest is, by operating on an outcome of that same operation performed previously. In general, it takes relatively little in the way of circular architecture in physical or computational systems to produce highly complex and convoluted behaviors. The effect is quite general and is entirely a consequence of the global architecture or topology of the causal closure. Indeed, some degree of causal "circling back," so to speak, is inevitable in the real world because of finiteness and aggregation.

The first component of my definition, then, is that nearly every sort of phenomenon that we are inclined to agree is emergent is considered so because it manifests some variant of this kind of circular causality. What we consider emergent about such phenomena are the consequences of this circularity. What emerges are certain configurational properties that become amplified via circular topology. By "amplified" I mean something like "come to be more coherently expressed over ever more extensive scales of both space and time."

Deterministic feedback in servomechanisms or finite recurrence in recursive computations exemplify the limiting cases of this architecture. Their behaviors are often referred to as emergent, but this is so only when both phenomena are considered in abstraction. By design, such systems exhibit closure to the world of the physical details of their embodiment. In working with reliable machines we can effectively ignore any undesigned influences; and we ignore their physicalness by abstraction in computational models. So we should probably describe them as exhibiting "computational emergence" as opposed to "ontological emergence," such as is characteristic of open physical systems. Higher-order emergent behaviors of complex servo-mechanisms occur in a deterministic single-level architecture and this precisely limits their dynamics to a narrow, even if sometimes near chaotic, range. As a result, any higher-order transformation of this dynamics is not possible. For example, there is no recruitment of noise from microscopic physical tolerance variations to become "signal," except in situations where amplified perturbations exceed some physical limits of the device and it "fails." Precisely because they are idealized either in the form of representations (e.g. computation) or as engineered devices that we only operate within set tolerances (which are also embodiments of representations), they cannot actually produce ontological emergence, only highly unpredictable, non-regular behaviors.

In the natural world, however, where forms and physical interactions grade into one another, circular causal architectures can spontaneously amplify form relationships originating from anywhere at any scale and as a result of any contingency, depending on the architecture. This openness to physical influence is missing from simulations and is minimized by design in electronic or mechanical feedback devices. It is typically a source of smoothing effects because of the uncorrelated nature of such factors, including everything from extrinsic environmental changes to quantum fluctuations. But if any given contingency, no matter how minute, can be amplified in its effect by some deviation amplifying causal architecture, then a kind of unpredictability that isn't merely the result of exceeding computability can result. In this case,

even in principle, it may be impossible also to predict what specific physical contingency will become selectively amplified. However, as complexity theorists and systems theorists have shown, despite the details, there may still be convergent development to predictable attractor patterns, which will have more to do with the general causal architecture of the system in question than with which contingent event was amplified.

5 Ascent in Scale

These considerations point to one more critical feature that is typically a correlate of emergence in physical systems, and which is nearly ubiquitous with spontaneous emergence: ascent in scale. The critical role played by ascent in scale in physical emergence is that it creates the context for causal circularity and amplification. It affords the substrate for structural influences to recirculate, so to speak. It is the physical analogue of repeated sampling from the same population in statistics. Ascent in scale increases the probability that potentially reinforcing regularities will compound and that nonreinforcing regularities will wash out. Although we must design and engineer simple systems to make controlled use of the amplifying effects of circular causality, it almost inevitably occurs, in at least minimal form, with an ascent in scale that includes very large numbers of interactions. This is because the numbers of ways that various physical effects can "circle back around" to the same locus in a fixed period of time increases rapidly with size and density of the population of interacting parts, whether in a physical medium or a randomly constructed network. My implication is a simple one: an ascent in scale that correlates with greater numbers of components and their interactions increases the probability that potential canceling and amplifying effects will be manifested in that collection. A subset of these effects are what we recognize as emergent.

The causal architectonic feature I will use as my central diagnostic feature might be described as trans-scale causal recursion, that is, circles of deviation amplifying causality that develop up levels of scale. My question is this: What happens when the global, topological regularities of a locally bounded open physical system are in some way "fed back" into that system? Following Einstein's more familiar characterization of nonlinearity in terms of compound interest, then, one might describe the taxonomy of emergence I have in mind as a categorization based on a sort of "compound interest" of form as a consequence of expansion in scale.

By distinguishing the causal topology by which these influences propagate through a system, three general categories of emergence can be derived. The three highly general ways that causal architectures can be topologically recurrent across levels of scale are (1) as nonrecurrent-architectures, (2) as simple recurrent-architectures, and (3) as recurrent-architectures of recurrent architectures. These produce phenomena that can correspondingly be called first-, second-, and third-order emergence, respectively, reflecting the progressive embedding of recurrent causal effects.

6 Simple Emergence

The most basic sense of emergence can be called first-order emergence. It corresponds to the way the term "emergence" is often applied to descriptively "simple" higher-order properties of an aggregate, such as the statistically- or stochastically-determined phase states of matter I have discussed above. Statistical thermodynamics and quantum theory have provided a remarkably complete theory of how the properties of water molecules can produce liquid properties in aggregate. Thus, in one sense, these higher-order properties are considered to be fully reducible to molecular properties and their effects on molecular interactions in aggregate. But such relational properties, as opposed to intrinsic molecular properties (e.g., mass, charge, configuration of electron shells, etc.) are not symmetric across levels of description. Precisely because they are relational, they are not applicable to descriptions of water molecules in isolation, for instance, and many will not be discernable in just any large aggregate. It is only when certain of the regularities of molecular interaction relationships add up rather than cancel one another that certain between-molecule relationships can produce aggregate behaviors with ascent in scale. So in this case, although the diversity of momenta, directions of relative movement, and orientation geometries of molecular interactions cancel one another out in aggregate, the ratio of hydrogen bond strength to energy of collision and its effect on molecular distances expresses itself as large-scale aggregate properties because these features produce an average net effect.

Philosophers of science have used the term "supervenience"⁸ to refer to the dependence of higher-order phenomena on lower-order phenomena, or rather for the more subtle relationship that exists between higher-order phenomena that are not exactly just scaled-up lower-order phenomena, and yet are products of them. Though the aggregate effects that consist in liquid properties are unlike any of the molecular interaction properties, there is a direct correspondence between the summation of these and some aggregate property. Because of the cancellation effects, in the aggregate much of the detail at the micro-level can be ignored without loss of descriptive adequacy. Micro-level events don't matter with respect to interactions at the macroscopic level, such as buoyancy of certain materials, surface tension, viscosity, and so forth. The higher-order properties nevertheless "supervene" on specific lower-order interactions to the extent that the former always entail the latter but not vice versa.

Not all aggregate properties can be considered emergent, however, even in this most reduced sense. For example, consider weight or density. Weight has a descriptive adequacy almost irrespective of level of analysis. The force of gravity exerts an effect on matter irrespective of scale, even if its effect may be scale-dependent in some ways. It is scale-dependent in that it has a relatively minor contribution to make to the trajectory followed by any individual water molecule in solution, but a major effect on the behavior of large bodies of water. Nevertheless, it is the same force and the same effect at both levels of scale. Nothing more needs to be added to the account of the aggregate behavior except that aggregate properties also are affected (e.g., surface tension and surface shape).

In contrast, other liquid properties develop out of the properties of water molecules in more complex ways. This occurs in cases where the lower-order relational properties are the constitutive factor determining the higher-order

property. In these cases, the account is fundamentally incomplete without including specific "configurational" and/or distributional information concerning the relationships between the components (e.g., quantity of constituents, their relative positions and momenta, their molecular geometry, their hydrogen bond strengths in different orientations, etc.). To represent all of this would produce a description with a very high dimensionality (many degrees of freedom, roughly corresponding to the relevant properties of each molecule with respect to the others, times the number of molecules) at the lower level. At the higher level of description, this dimensional complexity can often be ignored and little descriptive adequacy will be lost.

But this is not merely a descriptive matter. The astronomically many details cancel out due to the relative linearity of the stochastic processes and the evenly distributed diversity of these details. This results in a sort of phenomenological simplification with ascent in scale. This simplification tends to be nonuniform with respect to scale, often exhibiting threshold effects, or steps, of increasing size as regularities begin to dominate the messiness. The vast amount and diversity in causal detail that is distributed among the components and their interaction histories simply has little impact on the aggregate behavior if the aggregate is very large. The micro-effects are not just buried among the vast numbers of interactions, their differences cause some to get revealed in aggregate and others to get obscured.⁹ Net effects matter, but some, like mass merely matter via the aggregation of the tiny effects, while others matter because of the specific form of the aggregate effect. Where the form of aggregate effect is critical, the form itself becomes, in some cases, an additional physical dimension. Although not independent from the details, and despite the fact that it can be given a potentially exhaustive reductive description in terms of them, it can have other form-altering effects on other systems, not because of the details, but because of the aggregate form.

Wave propagation in a fluid is such an example. Although the nature of the wave and its detailed underlying dynamical realization in each may differ depending on whether the fluid is water, air, or an electromagnetic field, the ability to propagate a wave is a first-order emergent feature they all share in common. It is this emergent commonality that enables sound waves in air, for example, to be transformed into electric current fluctuations, transmitted through space as electromagnetic fluctuations, and converted via a radio receiver back into sound waves halfway around the globe. It is for this reason that first-order emergent phenomena, although minimalistically emergent, are an essential part of the concept: they are the basis for relationships that can be far more than merely supervenient.

8 Diachronic Symmetry-Breaking

Philosophical and scientific discussions of the mind-brain mystery often invoke some implicit version of supervenient emergence to model the presumed relationship between higher-order, mental phenomena and the lower-order, cellular-molecular processes on which they depend. But, as noted above, it is clear that the analogy is not nearly adequate. At the very least it fails to capture an essential distinguishing feature: temporal development or symmetry-breaking. There is a simple self-similarity to liquid properties across time and

position that is further "smoothed" by stochastic processes. In contrast, there is a self-differentiating feature to living and mental processes, which both retains and undermines aspects of self-similarity despite (or rather ironically because of) stochastic features. This characteristic breakdown of self-similarity, or symmetry-breaking, is now recognized in numerous kinds of complex phenomena, including systems far simpler than living systems. Complex emergent phenomena of these sorts share a characteristic change of ensemble properties across time, and are often computationally unpredictable as well. So it would be useful to distinguish first-order emergence from these more complex forms of emergent phenomena in which the cumulative stochastic canceling of configurational interactions exhibited by simple entropic systems is undermined, and where this contributes to development and change of both micro- and macro-properties across time.

At first glance, it may appear as though distinguishing between synchronic and diachronic classes of emergent phenomena can capture this difference. First-order emergence is typically described in synchronic terms, and, although it is recognized that this is a short-hand descriptive trick, there is very little descriptive loss in this heuristic. But in many discussions, the synchronic and diachronic uses of the concept must be made explicit. For example, biological scientists now feel comfortable with the notion that life is an emergent property of organic matter that supervenes on the properties of its constituent molecular interactions. Because life is a continuous property of an organism, one can describe it at any one moment as synchronically supervenient on the properties of the organism's constituent molecules and their relationships. But this fails to capture the vast and convoluted temporal depth of the causal account that would be necessary to understand the configurational details that constitute structures and behaviors that we intend to capture by the notion of "life." Specifically, biological aggregate properties that we consider "functions" implicitly require reference to a detailed history of at least the micro- and macro-configurational relationships. These histories link biological function with interactions among critical substrate variables, and, implicitly, with an evolutionary ensemble of reproductively related individuals and their interactions with these substrates. In short, the synchronic shorthand account assumes far more than a simple diachronic restatement of first-order emergence.

This can be demonstrated easily by considering evolutionary examples. Consider the process by which life must have arisen from nonlife in the early history of the earth. Both a synchronic and a diachronic sense of emergence must be considered here. The higher-order properties of the molecular systems that constituted the first lifelike self-organizing proto-metabolisms can be said to have supervened on the molecular interactions concurrently involved, since the present state of these very interactions is the basis for describing the whole ensemble as having the property of being alive. But, prior to the formation of the first molecular configuration of this kind, no living phenomena existed. So, in addition to being first-order emergent by way of interactions of lower-order constituents, a higher-order emergence relationship is also involved. The configurational properties of pre-biotic systems are a part of the causal complex upon which the first biotic systems supervene-in time. In other words, evolutionary processes must be described as the successive emergence of new

emergent phenomena from old: emergent phenomena constituted of other emergent phenomena. By this I do not mean to suggest that "life" itself is more than a re-description of this supervenience, but rather to note that there is a temporal component to this kind of emergence as well as ascent in scale. There is an asymmetry in configuration across time—a development.¹⁰

This difference is not merely a matter of adding levels upon levels, as in the way solid-state physics reduces to atomic interactions that in turn reduce to elementary particle interactions and so on (presumably to quarks, etc.). Rather, it involves nonlinear systems that exhibit an internal causal recursivity and ultimately a self-undermining dynamic that causes prior states of the system to be irreversibly replaced and superseded. This recursiveness is the basis of an irreducibly diachronic asymmetry. But for this reason the distinction is more complicated than merely distinguishing between synchronic systems with self-similar, simple, entropic development (e.g., behavior of a gas) and diachronic ones with difference-amplifying development. There is more than one way to change.

8 Self-Assembly and Self-Organization

Surface tension is an emergent property of a comparatively simple type. It is characterized by predictability so long as lower-order properties fall within certain parameters. But there are conditions where this is not the case: specifically under conditions where more or less chaotic or self-organized behaviors are produced. Under chaotic conditions, for example, certain higher-order regularities become unstable, and an unpredictability of higher-order dynamics results. In chaotic systems, unpredictability derives from the fact that regularities at lower levels have become strongly affected by regularities emerging at higher levels of organization. This can happen when configurational features at the ensemble level drastically change the probabilities of certain whole classes of component interactions.

This occurs in shock wave boundaries, for example, where relative movements of whole ensembles of gas molecules exceed the rate at which energy can be exchanged by typical elastic collisions between individual molecules. So breaks in symmetry appear that affect component molecular interactions. Those on one side or the other of this interactive threshold can have an inordinate influence over ensemble behaviors. The specific nonsystematic locations of these micro-symmetry-breaks become important since they become loci of major energy transfer for whole subsystems distinguished by sharing similar global features. The result is that these specific micro-configurations have macro-configurational consequences (e.g., local energetic "cascades" and irregularities) that in turn affect future micro-configurations, and so on. This hierarchical non-linearity consequently produces a kind of self-undermining dynamic across levels that is expressed in a unique time series of configurations.

The signature feature of this complication of first-order emergence is that the configuration of individual components and unique interactions can exert an organizing effect on an entire ensemble. While in principle these examples still exhibit first-order emergence, we now must contend with an additional second-order emergence of behaviors and ensemble properties. Whereas micro-configuration can be ignored in first-order emergent systems, with minimal loss

of descriptive adequacy, this is not the case for systems exhibiting second-order emergence. Chaotic and self-organized systems are generally of this type. Chaotic phenomena have become a major focus of "complexity" theory in its many forms. Such phenomena cannot be adequately described without incorporating a detailed history of the system and its components, whereas merely first-order emergent phenomena can. Moreover, as has come to be the hallmark of complex chaotic systems, initial conditions of system history can play critical roles in the ongoing global dynamics of the system.

This is often expressed as the apocryphal "butterfly effect," in which it is supposed that a butterfly flapping its wings in a particular pattern in one part of the globe is responsible for monsoons forming in another part of the globe, years later. Although in reality this is an implausible and misleading just-so story (since the highly entropic dynamics of the atmosphere tend to quickly smooth out such perturbations rather than amplify them; see below), it does capture the possibility that large-scale configurational properties of a system could conceivably be substantially affected by extremely miniscule differences in prior micro-configurations of the same system under certain special circumstances. What are these special circumstances? Whereas atmospheric dynamics can be chaotic and unpredictable, they are not systematic. As a result, and importantly, they do not tend to consistently amplify specific perturbations. Certain other kinds of highly organized physical systems can, however, exhibit what I will call an "ampliative dynamic." This gets reflected in their critical dependency on historical contingency and complex unique individual structure. Take for example, snow crystal formation.¹¹ The structure of an individual snow crystal reflects the interaction of three major factors: (1) the hexagonal micro-structural biases of ice crystal lattice growth, inherited from water molecule symmetry, (2) the radial symmetry of heat dissipation, and (3) the unique history of changing temperature and humidity regimes as a developing crystal falls through the air. Different temperatures and humidities interact with water molecule binding regularities to generate a handful of distinct patterns of ice lattice formation. Since a snow crystal grows as it falls through different regions in a variable atmosphere, a record of the temperature and humidity differences it encounters literally gets frozen into the variant patterns of crystal structure at successive diameters.

In this way, the snow crystal is a sort of record of the conditions of its development. But it is more than merely a historical record. There is a "compound interest" effect as well. Prior patterns of crystal growth constrain subsequent patterns of crystal growth. So even identical conditions of temperature and humidity, which would otherwise determine identical lattice growth, can produce different global patterns depending on the current configuration of the crystal. The global configuration of this tiny developing system plays a critical causal role in its microscopic dynamics by excluding the vast majority of possible molecular accretions and growth points, and strongly predisposing accretion and growth at certain other sites.

Snow crystals are in this way at least minimally self-organizing (though not strictly speaking autopoietic; see below). The unique growth history of the crystal cannot be ignored. Indeed growth history is a dominant factor determining its final configuration. More importantly, these configurational

properties of the system self-amplify. Each past configuration influences all subsequent configurations as well. This influence occurs by virtue of the progressive contraction of the potential growth options that results from past growth.¹²

Both self-undermining (divergent) chaotic systems, as in turbulent flow, and self-organizing (partially convergent) chaotic systems, as in snow crystal formation, exhibit causal circularity of a kind that links higher-order and lower-order configurational properties. These feed-forward circles of cause and effect, linking reciprocally reinforcing effects at different levels of scale, are the defining features distinguishing second-order emergence from first. It is "second order" because first-order emergent properties have become self-modifying, resulting in the emergence of new emergent phenomena.

More complex second-order emergence is exhibited by "autopoietic" systems. In autopoiesis the interaction dynamics of sets of different components is constrained both by the configurational properties of the whole collection and by configuration symmetries and asymmetries that exist between the micro-configurations of the different classes of its components. In snow crystal dynamics the micro-configuration of each component is essentially the same, producing symmetric interactions and strongly constrained structural consequences. When a system is composed of different kinds of components, however, it can also exhibit a more distributed interactional reflexivity. Because of the combinatorial possibilities the resultant properties can be quite complex.

For example, molecules that interact in highly allosteric fashion, that is, which bond selectively with some but not other types of molecules, can constitute interaction sets with more elaborate self-organizing features. Both the configurations of the different classes of individual interactions and the configuration of the whole set of possible interactions become critical organizing influences. Processes of this sort can occur in a chemical "soup" that contains enough different kinds of molecules so that there is a subset of types that can each catalyze the formation of some other member in the set, constituting a closed loop of syntheses. (This is called an autocatalytic set; for a recent discussion see Kauffman 1995, 2000.) So long as sufficient energy and other raw materials are available to keep reactions going (it must, that is to say, be an open system) the "autocatalytic" character of the set will play an inordinate role in determining both what chemical reactions can take place and how the whole soup will be constituted. It is the higher-order distributed circularity of the interactions of the different classes of constituents that matters. Such a system can generate far more complex macro-dynamics and macro-dynamics than if the interactions were symmetric. The general theory of chemical reactions with these circular features was extensively investigated by Ilya Prigogine and by Manfred Eigen a generation ago. It has now become the basis for extensive research with both real and simulated chemical systems. (See Prigogine 1997, Eigen and Oswatitsch 1992 for their own accounts.) Ultimately, the metabolic molecular dynamics that constitute living cells depends on autopoietic system dynamics constituted by numerous fully and partially autocatalytic sets of molecules.

What do these examples of second-order emergent phenomena have in common? The answer is: a kind of tangled hierarchy of causality where micro-configurational particularities can be amplified to determine macro-configurational regularities, and where these in turn further constrain and/or amplify subsequent micro-configurational regularities. In such cases, it is more appropriate to call the aggregate a "system" rather than a mere collection, since the specific reflexive configurational and recurrent causal architecture is paramount. So, although these systems must be open to the flow of energy and components- which is what enables their growth and/or development-they additionally include a kind of configurational closure as well. These material flows carry configurational constraints inherited from past states of the system that constrain future behaviors of its components. As material and energy flows on through and out again, form also recirculates and becomes amplified. In one sense, this form is nothing more than a set of restrictions and biases on possible future material and energetic events. In another sense, it is what defines and bounds the higher-order individual entity that we identify as the system.

Not surprisingly, quasi-crystalline and autopoietic, second-order, emergent systems are often found linked in living systems. For example, consider the configuration-constraining role of quasi-crystal-like molecular structures, such as the polymers and membranes in living cells. These can grow by self-assembly and also tend to play a critical role in constraining the interactions critical to living, autopoietic, chemical reactions. The configuration of such a substrate may both catalyze certain individual reactions and increase the probability of causal circularity of a reflexive set of other catalytic reactions, for example, by aligning multiple species of molecules with respect to one another. This kind of interaction between emergent phenomena turns out to be important for yet higher levels of emergent interaction.

9 Evolution and Semiosis: Third-Order Emergence

There is a further difference, however, between chaotic/self-organizing emergent phenomena, like snow crystal growth or chemical autopoiesis, and evolving emergent phenomena, such as living organisms. The latter additionally involve some form of information or memory (for example, as represented in nucleic acids) that is not seen in second-order systems. The result is that specific historical moments of higher-order regularity or of unique micro-causal configurations can exert an additional cumulative influence over the entire causal future of the system. In other words, via memory, constraints derived from specific, past, higher-order states can get repeatedly re-entered into the lower-order dynamics leading to future states, in addition to their effects mediated by second-order processes.

To get a sense of this additional loop of causality, imagine the following fanciful possibility: a kind of snow crystal that, instead of just melting, contributes some bias to the formation and ongoing growth of a new snow crystal.¹³ It might perhaps provide a seed crystal, that reflected some global aspect of its geometry to the new crystal growth (e.g. the relative predominance of spire-like or plate-like geometry, the degree of structural heterogeneity, etc.). Each new generation of snow crystal would consequently be biased by the

cumulative results of what happened to some prior snow crystal, and so on. Some trace of the history of cumulative biases would be repetitively updated and re-entered as an initial bias in each new snow crystal "generation." In this world, snow crystals in different eras would likely look different, and these differences would follow various trends. But to this fanciful example we need to add one more ingredient. There is now a whole class of new contingencies that can accumulate and amplify along another gradient of scale: the expansion of lineages (defined with respect to continuity of seeding) across generation-time. A lineage is an extended aggregate in both space and time. In this way, some contingent feature represented in an individual of an aggregate can give rise to its own temporally and spatially extended aggregate of individuals. Space and time get convolved together in this way, and a new type of amplification is possible.

In this way, larger scale extrinsic contingencies and boundary conditions will matter in ways not characteristic of second-order emergence. Analogous to the ampliative architectures that occur via ascent of scale in second-order emergence, there is now the possibility of ampliative architectures via ascent in temporal scale-historical extent-because greater and greater extended chunks of contingent events can be reentrantly compounded and fed-forward. In second-order emergence, a certain synergy between boundary conditions and the relational topologies of the constituents of a system results in the self-amplification of certain features as opposed to others. With this additional twist of causality, an analogous selective amplification bias can occur also with respect to these self-amplification conditions themselves. Contingent synergies between the conditions in which these fanciful snow crystals grow can now come to bias what aspects of snow crystal geometry will be passed on and amplified in the future. In other words, this allows for a process of selection among alternative second-order emergent processes. (The term "selection" is merely a post hoc short hand way of describing the result of this spontaneous synergy-biased amplification, as it is also in evolutionary theory.)

In these systems, there is a kind of synergy between self-organization with representation. Even the topology of the ampliative architecture itself can be convoluted in unpredictable ways by a higher-order amplification in which such details come to play an inordinately important role in determining the causal architecture itself. For this kind of phenomena we must introduce a third-order of emergence, which recognizes the additional loop of recursive causality, enclosing the second-order recursive causality of self-organized systems. Capturing, condensing, and re-entering the effects of selected events distributed across a span of lineage history is what makes evolutionary processes both chaotically unpredictable in one sense and yet also historically organized, with an unfolding quasi-directionality.

Third-order emergence inevitably exhibits a developmental and/or evolutionary character. It occurs where there is both amplification of global influences on parts, but also redundant "sampling" of these influences and reintroduction of them redundantly across time and into different realizations of the same type of second-order system. Under these conditions there can be extensive amplification of lower-order emergent relationships and also lower-order chaotic/self-organizing relationships because their historical traces do

not degrade. These traces get repeatedly re-entered into new iterations of the system in new contexts. Whereas second-order emergent phenomena exhibit locally and temporally restricted whole-to-part influences, third-order emergent phenomena can exhibit amplification of these effects as well, doubly convolving the link between levels and scales of causality.

This can be imagined as a sort of self-referential self-organization, an autopoiesis of autopoieses. Amplification of complexity and of self-organizational dynamics can be enormously complex under these conditions, forming into a maze of causal circularities, because every prior state is a potentially amplifiable initial condition contributing to all later states. Moreover, because there is a "re-membered" trace of each prior "self" state contributing to the dynamics of future states, such systems can develop with respect to this prior "self," rather than just with respect to the immediately prior state of the whole. This fact contributes to the characteristic differentiation and divergence from, and convergence back toward, some "reference" state, as exhibited in organisms.

The representation relationship implicit in third-order emergent phenomena demands a combination of multi-scale, historical, and semiotic analyses for adequate description. This is why living and cognitive processes require us to introduce concepts such as representation, adaptation, information, and function in order to capture the logic of the most salient emergent phenomena. It is what makes the study of living forms qualitatively different from other physical sciences. It makes no sense to ask about the function of granite. Though the atoms composing a heart muscle fiber, or a neurotransmitter molecule, have no function in themselves, the particular configurations of the heart and its cell types, or the neurotransmitter molecule, do additionally beg for some sort of teleological assessment, some function. They do something for something. Organisms evolve and regulate the production of multiple second-order emergent phenomena with respect to some third-order phenomenon. Only a third-order emergent process has such an intrinsic identity.

So life, even in its simplest forms, is third-order emergent. This is why its products cannot be fully understood apart from either historical or functional concerns. Indeed, it may be that any third-order emergent system must be considered "alive" in some sense.¹⁴ This suggests that third-order emergence may offer something like a definition of life in its broadest sense. If this is so, then the origins of life on the earth may also mark the initial emergence of third-order emergent phenomena on the earth. More generally, it constitutes the origination of information, semiosis, and teleological relationships on earth. It is the creation of an "epistemic cut" to use Howard Pattee's felicitous phrase: the point where physical causality acquires (or rather constitutes) significance. This is a complex and subtle issue, which I can only obliquely approach in this overview. Nonetheless, a few points are relevant to the purpose of this chapter.

The first is that third-order emergence serves as a sort of minimal account of what we mean by information. It is an implicit account of what constitutes the creation of information. The physical conveyance of configurational bias from one emergent context to another to generate third-order emergence constitutes the most primitive "aboutness" relationship. It is the embodiment of

a "selection" among possible alternatives among some set of possible configurations of the same "type." Specifically it captures in a concrete way what might be called the potential amplifiability of some emergent features with respect to other emergent features and with respect to the ensemble of possible local boundary conditions. This progressively amplifies a correspondence relationship. Third-order emergence links the sampling over an ensemble of related emergent systems to sampling over an ensemble of likely boundary conditions. So in an evolutionary context we say that a given configuration of an organism is an adaptation to something. This is, implicitly, a representational relationship, even if the specific object of the representation is often somewhat vague.

Consider for example the dynamical organization of the molecular system that constitutes the ciliated pellicle, or skin, of a paramecium (the familiar sole-of-a-shoe-shaped single-celled pond organism). The ciliated pellicle of the paramecium enables it to swim around obstacles by asymmetrically depolarizing, causing the cell to tumble, when it bumps up against something. I think it is appropriate to say that in some basic sense this complex configuration of molecular modules, with their linked molecular "tails," "motors," and mechanoreceptors is a representation of a cumulative past history of obstacles and their properties in terms of how these affected the reproduction of past configurations of the whole complex. The configuration of the pellicle and its responsivity re-presents in the present system certain critical regularities of "obstacleness" that were sampled and amplified over billions of obstacle encounters and their outcomes in past generations. The "tokens" conveying this representation are ultimately genes, as interpreted with respect to the way the components tend to self-assemble and interact. The result is a system constituted so that it exhibits selective correspondent behaviors to classes of likely contextual conditions. It projects selective abstract regularities of the past onto potential future conditions. In this way the pellicle configuration is "about" obstacle avoidance. Evolution has produced the emergence of something like a functional category: 'obstacle.'

10 A Hierarchy of Emergent Phenomena

The three subcategories of emergent phenomena can be arranged into a hierarchy of increasing complexity, in which higher-order forms are composed of relationships between lower-order forms. Third-order (evolutionary) emergence contains second-order (self-organizing) emergence as a limiting case, which in turn contains first-order (supervenient) emergence as a limiting case. In this way, any given example of evolutionary emergence must also involve self-organizing and supervenient emergence as well, but not vice versa. Because higher-order emergent phenomena are dependent on and constituted by lower-order emergent phenomena, their probability of spontaneous occurrence is also substantially lower. Consequently there are vastly more examples of supervenient emergent phenomena than of self-organizing emergent phenomena and, in turn, vastly more examples of self-organizing emergent phenomena than of evolutionary emergent phenomena. So only a few generically supervenient phenomena are alive, but every living thing is a supervenient emergent phenomenon. Conversely, however, with the replication made possible by third-order emergence, the number

of emergent phenomena of all types can rapidly increase, although the production of lower-order, subordinate, emergent phenomena produced to support the third-order will always be far more numerous.

This hierarchic categorization does not exhaust the possibilities of increasingly more complex forms of emergent phenomena. Evolutionary emergent systems can further interact to form multilayer systems of exceeding complexity. Indeed, this is the nature of complex organisms that is exemplified in the ascending levels of "self" that proceed from gene to cell to organism to lineage to species, and so on, in the living world. But this logic does not lead to what might have been called fourth-order emergence. Instead, I think we must rather analyze these more complex processes as first-, second-, and third-order emergent elaborations of third order emergence, and so on, in recursive series. This is because third-order emergence includes the capacity to evolve new forms of emergence itself. This is implicit in its inherently representational or semiotic character. The introduction of referential relationship as the defining feature of third-order emergence creates a spatial and temporal boundedness that is able to encompass any physical system, and re-present any system with respect to its correspondences with another. So there is no upper, outer, past, or future bound to what can constitute a third-order emergent phenomenon. Representational capacity is ultimately unbounded.

One of the advantages of framing evolutionary processes in these quite abstract and general terms is that doing so helps us to delineate where the evolutionary and self-organizing processes of living systems can be distinguished and where they interact. For example, it makes explicit the dependence that evolutionary processes have on self-organizational processes, and provides important hints about how evolutionary processes might have themselves emerged from something simpler. So, for example, those approaches to the origins of life that emphasize the role of autocatalysis or autopoiesis as the necessary antecedent processes to living processes are consistent with this analysis. So are those that emphasize replication processes. But what might be called "replicator-first" models are not. The reason is that for replication to be ampliative it must occur as a part of and with respect to self-organization processes whose topologies are thus captured and conveyed to future iterations. Similarly, it also suggests a categorization of biological evolutionary processes that can encompass some of the complexities introduced by complex organisms and by the representational processes introduced into biology with the evolution of brains and symbolic communication.

Thus minimalist Darwinian conceptions of natural selection are a first-order variant of third-order emergence in the sense that the accumulation of the consequences of natural selection are treated as summed and integrated in individuals of succeeding lineages. There are no "top-down" effects. The adaptive self-organizing processes of the organism are not assumed to exert any feed-forward effect that biases what is subjected to selection in future generations.

However, the process described as "niche construction" and various forms of Baldwinian-like effects are more complex. (See Deacon, Chap. 5, as well as others in this volume, for a critical re-evaluation of the Baldwin effect.) In these cases, the effects of the organism's adaptive responses do feed-forward to

affect what gets exposed to or hidden from the effects of selection. As a result, a kind of self-organizational dynamic can develop in which the specific form of the behavioral or physiological adaptability can become a bias on the future range of naturally selected, adaptive capacities. So, for example, beaver dam-building and pond creation create a condition in which selection favors the evolution of aquatic adaptations, and aquatic adaptations increase the advantage of having dam-building propensities, in an ever-increasing ratcheting of interdependence. The specific details of the behavioral adaptation become the initiating biases in a compound-interest effect that both amplifies and accumulates concordant adaptations. These effects then are second-order emergent variants on the third-order emergent process of natural selection, which, like the cumulative constraints of the growing snow crystal, are incrementally amplified in succeeding stages, creating a more and more idiosyncratically integrated and distinctive result.

Finally, consider the evolution of symbolic communication in hominids. (See my detailed discussion in Deacon, Chap. 5, this volume.) This evolutionary transition marks the emergence of a new and partially decoupled evolutionary dynamic. Symbolic communication has evolved to be primarily exemplified in a highly specialized form-language-both by virtue of its influence on the evolution of the brain and by virtue of the evolution of an independently evolving and self-organizing system of symbols and their effects. Linguistic and cultural evolution exhibit their own third-order emergent dynamic. The representational nature of this adaptation exemplifies its third-order character. The effects on human physiology and human consciousness have clearly become the tail that wagged the dog in many respects.

11 Conclusions

The insistent critique from a systems-theoretic perspective of both genetic-reductionism in evolutionary theories and computational reductionism in cognitive theories can now be more precisely paraphrased. Life and mind cannot be adequately described in terms that treat them as merely supervenient because this collapses innumerable convoluted levels of emergent relationships. Life is not mere chemical mechanism. Nor is cognition mere molecular computation. These analogies miss the most salient and descriptively important dynamics of these phenomena. To collapse descriptions of living and cognitive systems in this way inevitably ignores whole classes of causal topologies that are essential to their comprehensive explanation and contribute their most robust and important characteristics. In living processes, for example, the sources of structure and function can become increasingly distributed in space and time over the course of evolution. Organizing processes may also extend to include progressively higher-level subsystems exhibiting progressively greater top-down control architecture; for instance, brains.

The capacity of evolutionary emergent processes to progressively embed evolutionary processes within one another via representations that amplify their information-handling power makes the mythical butterfly-effect story trivially simplistic. It is also what makes minds possible. The many levels of embedded evolutionary emergent processes generated in brain activities are what enable them rapidly and selectively to amplify such a vast range of possible forms of

activity. Which particular molecular system configurations will become amplified to produce subsequent states of the brain is both essentially untraceable and largely irrelevant. The most salient causal antecedents are themselves idiosyncratic global third-order emergent configurations of neural activity of an exquisitely indirect and astronomically unlikely variety. The incorporation of language into this complex convoluted hierarchy of emergent processes has further vastly amplified the emergent capacities of cognition, expanding the realm of represented forms into a truly limitless universe of emergent possibilities.

A living organism is the end product of a very elaborately convoluted history of recursive causal processes that cast a wider and wider net to capture sources of regularity and amplify them. The result is that locating any specific antecedent cause of functional organization is essentially impossible. This is one reason that evolutionary "explanations" of specific traits are inevitably vague and general in form. There is no clear causal trend from more local to more distributed, and from micro to macro scales. There are not even clear boundaries in space and time. Indeed the causes of living organization grade into molecular noise and distant past contingencies. An organism's functional properties may be currently instantiated by its molecular architecture (first-order, local). Yet this architecture emerges from a vast ensemble of molecular interactions with the world, generated by innumerable past members of a lineage (third-order, nonlocal).

Brain processes produce a further baroque convolution and temporal amplification of this logic. Brains might be characterized as "emergence machines," incessantly churning out complex high-level virtual functions, virtual environments, and virtual evolutionary lineages to track and adapt to the complexity of the world. The minds that result are marvels of high-level holistic causal loci. Subjective experience reflects this convergent holism. The experience of being a sentient agent is no less than being the locus of something that is incessantly and spontaneously emerging. This experience is itself an emerging locus at the center of a vast but only weakly constraining, weakly determinate web of semiotic and physiological influences.

Language further distributes the convoluted trace of causality that constitutes human agency. Human minds that have become deeply entangled in the evolving symbolic communicative processes of culture—as all modern human minds are—may have an effective causal locus that extends across continents and back millennia, and which grows out of the experiences of hundreds of thousands of individuals. This immense convergence of causal determination is coupled with a vast capacity for selective amplification that has itself been amplified beyond recognition by our cognitive involvement in the super-organism of symbolic culture. Human consciousness is not merely an emergent phenomenon; it epitomizes the logic of emergence in its very form: the locus of an immense confluence, condensation, amplification, and dissemination of topological constraint and bias that continually deviates from and undermines old patterns of causality and amplifies new ones: causality transcending itself.

References

Crick, Francis (1994) *The Astonishing Hypothesis: The Scientific Search for the Soul*. New York: Charles Scribner's Sons.

Eigen, Manfred and Oswatitsch, Ruth Winkler (1992) *Steps Toward Life: A Perspective on Evolution*. New York: Oxford University Press.

Holland, John (1998) *Emergence: From Chaos to Order*. New York: Oxford University Press.

Johnson, Steven (2001) *Emergence: The Connected Lives of Ants, Brains, Cities, and Software*. New York: Scribner.

Kauffman, Stuart (1995) *At Home in the Universe: The Search for the laws of Self-Organization and Complexity*. New York: Oxford University Press.

Kauffman, Stuart (2000) *Investigations*. New York: Oxford University Press.

Prigogine, Ilya (1997) *The End of Certainty: Time, Chaos, and the New Laws of Nature*. New York: Free Press.

Weber, Bruce and Deacon, Terrence (2000) *Thermodynamic Cycles, Developmental Systems, and Emergence Cybernetics & Human Knowing 7: 21-43*.

1 This paper includes material excerpted from a paper by Bruce Weber and Terrence Deacon published in *Cybernetics & Human Knowing 7: 21-43, 2000*.

2 From Ursula Goodenough, personal communication.

3 Much has been written about this notion. The interested reader can find a recent lay introduction to the general topic in *Emergence: The Connected Lives of Ants, Brains, Cities, and Software* by Steven Johnson (2001) and a slightly more sophisticated introduction in *Emergence: From Chaos to Order* by John Holland (1998).

4 Of course, this is a much more knotty problem than I can address here. To some extent, these concepts also reflect a levels-of-analysis distinction since constituent properties are to a large extent the result of lower-level topological features as well. For the sake of this discussion I will treat these as distinguishable at any given level.

5 This is a composite statement I have assembled by borrowing from a number of sources.

6 This is quite possibly an apocryphal quote, since I can find no precise citation context among the dozen or so places it is mentioned in the literature and on web sites. But I am inclined to believe it was something he could have said, given the sense of humor it conveys juxtaposed with the not immediately obvious mathematical profundity it captures. The author would be grateful to the reader who could to supply the precise context and citation, if it exists.

7 Though it might be more accurate to use the metaphor of 'spiral' causality, I will use the terms 'circular' and 'recursive' because they make it easier to visualize more complex and convoluted architectures. I have decided to use the term 'cause' in this colloquial sense throughout, but with some reservation, because it ultimately begs the questions I am trying to answer. In most cases understanding it in terms of the expansion, recursion, or propagation of topological constraints captures the more precise meaning I have in mind.

8 I can't claim to even attempt to capture all the tortuous uses this term has been subject to in philosophical writings, and here I use only one of its most basic senses.

9 Consider momentum of molecules. Usually, in aggregate, momenta in different directions cancel. But imagine a body of water that suddenly lurches upward

because a majority of the vectors of momenta of moving molecules just coincidentally happened to not be uniformly distributed.

10 This developmental asymmetry is often described as a trend that is opposed to entropy-sometimes called 'negentropy'-though I consider it more appropriate to describe it as a very rare expression of entropy.

11 In the above account of snow crystal growth I have omitted many details of the physics of this process for the sake of focusing on those most relevant to this discussion. These include the physics of the different modes of lattice formation and growth that vary with temperature and humidity and also some poorly understood aspects of surface physics including the quasi-liquid dynamics of the molecular surface of the growing crystal and how these affect the growth and symmetries of the crystal. These omissions do not substantially alter the description in any way that is relevant to this account.

12 Note that growth is not an essential factor since this same dynamic also obtains if during certain phases of snow crystal development there is periodic partial melting-a common occurrence-which is similarly historically constrained and which can lead to elaborately shaped entrapped bubbles and pits with smooth curved edges.

13 This fanciful analogy was first suggested to me by Don Favareau of UCLA.

14 As mentioned previously, with respect to physical emergence versus the virtual emergence of computational and mechanical systems, a simulation of third-order emergence is possible, but because it is a representation and the physical ground of its constituents are not constituents of the emergent dynamics, I do not consider them in themselves physical emergent. Perhaps they can be called computationally or symbolically emergent, although the principles are somewhat different. The simulation of complex systems and 'artificial-life' ecosystems can nonetheless be highly instructive tools that capture many aspects of this dynamic in representational form.