

Is Anything Ever New?

Considering Emergence

James P. Crutchfield

**Physics Department
University of California
Berkeley, California 94720**

Abstract

This brief essay reviews an approach to defining and then detecting the emergence of complexity in nonlinear processes. It is, in fact, a synopsis of Reference [1] that leaves out the technical details in an attempt to clarify the motivations behind the approach.

The central puzzle addressed is how we as scientists — or, for that matter, how adaptive agents evolving in populations — ever “discover” anything new in our worlds, when it appears that all we can describe is expressed in the language of our current understanding. One resolution — hierarchical machine reconstruction — is proposed. Along the way, complexity metrics for detecting structure and quantifying emergence, along with an analysis of the constraints on the dynamics of innovation, are outlined. The approach turns on a synthesis of tools from dynamical systems, computation, and inductive inference.

1 Emergent?

Some of the most engaging and perplexing natural phenomena are those in which highly-structured collective behavior emerges over time from the interaction of simple subsystems. Flocks of birds flying in lockstep formation and schools of fish swimming in coherent array abruptly turn together with no leader guiding the group.[2] Ants form complex societies whose survival derives from specialized laborers, unguided by a central director.[3] Optimal pricing of goods in an economy appears to arise from agents obeying the local rules of commerce.[4] Even in less manifestly complicated systems emergent global information processing plays a key role. The human perception of color in a small region of a scene, for example, can depend on the color composition of the entire scene, not just on the spectral response of spatially-localized retinal detectors.[5,6] Similarly, the perception of shape can be enhanced by global topological properties, such as whether or not curves are opened or closed.[7]

How does global coordination emerge in these processes? Are common mechanisms guiding the emergence across these diverse phenomena?

Emergence is generally understood to be a process that leads to the appearance of structure not directly described by the defining constraints and instantaneous forces that control a system. Over time “something new” appears at scales not directly specified by the equations of motion. An emergent feature also cannot be explicitly represented in the initial and boundary conditions. In short, a feature emerges when the underlying system puts some effort into its creation.

These observations form an intuitive definition of emergence. For it to be useful, however, one must specify what the “something” is and how it is “new”. Otherwise, the notion has little or no content, since almost any time-dependent system would exhibit emergent features.

1.1 Pattern!

One recent and initially baffling example of emergence is deterministic chaos. In this, deterministic equations of motion lead over time to apparently unpredictable behavior. When confronted with chaos, one question immediately demands an answer — Where in the determinism did the randomness come from? The answer is that the effective dynamic, which maps from initial conditions to states at a later time, becomes so complicated that an observer can neither measure the system accurately enough nor compute with sufficient power to predict the future behavior when given an initial condition. The emergence of disorder here is the product of both the complicated behavior of nonlinear dynamical systems and the limitations of the observer.[8]

Consider instead an example in which order arises from disorder. In a self-avoiding random walk in two-dimensions the step-by-step behavior of a particle is specified directly in stochastic equations of motion: at each time it moves one step in a random direction, except the one it just came from. The result, after some period of time, is a path tracing out a self-similar set of positions in the plane. A “fractal” structure emerges from the largely disordered step-by-step motion.

Deterministic chaos and the self-avoiding random walk are two examples of the emergence of “pattern”. The new feature in the first case is unpredictability; in the second, self-similarity. The “newness” in each case is only heightened by the fact that the emergent feature stands in

direct opposition to the systems' defining character: complete determinism underlies chaos and near-complete stochasticity, the orderliness of self-similarity. But for whom has the emergence occurred? More particularly, to whom are the emergent features "new"? The state of a chaotic system always moves to a unique next state under the application of a deterministic function. Surely, the system state doesn't know its behavior is unpredictable. For the random walk, "fractalness" is not in the "eye" of the particle performing the local steps of the random walk, by definition. The newness in both cases is in the eye of an observer: the observer whose predictions fail or the analyst who notes that the feature of statistical self-similarity captures a commonality across length scales.

Such comments are rather straightforward, even trivial from one point of view, in these now-familiar cases. But there are many other phenomena that span a spectrum of novelty from "obvious" to "purposeful". The emergence of pattern is the primary theme, for example, in a wide range of phenomena that have come to be labeled "pattern formation". These include, to mention only a few, the convective rolls of Bénard and Couette fluid flows, the more complicated flow structures observed in weak turbulence,[9] the spiral waves and Turing patterns produced in oscillating chemical reactions,[10–12] the statistical order parameters describing phase transitions, the divergent correlations and long-lived fluctuations in critical phenomena,[13–15] and the forms appearing in biological morphogenesis.[10,16,17]

Although the behavior in these systems is readily described as "coherent", "self-organizing", and "emergent", the patterns which appear are detected by the observers and analysts themselves. The role of outside perception is evidenced by historical denials of patterns in the Belousov-Zhabotinsky reaction, of coherent structures in highly turbulent flows, and of the energy recurrence in anharmonic oscillator chains reported by Fermi, Pasta, and Ulam. Those experiments didn't suddenly start behaving differently once these key structures were appreciated by scientists. It is the observer or analyst who lends the teleological "self" to processes which otherwise simply "organize" according to the underlying dynamical constraints. Indeed, the detected patterns are often *assumed* implicitly by the analysts via the statistics selected to confirm the patterns' existence in experimental data. The obvious consequence is that "structure" goes unseen due to an observer's biases. In some fortunate cases, such as convection rolls, spiral waves, or solitons, the functional representations of "patterns" are shown to be consistent with mathematical models of the phenomena. But these models themselves rest on a host of theoretical assumptions. It is rarely, if ever, the case that the appropriate notion of pattern is extracted from the phenomenon itself using minimally-biased discovery procedures. Briefly stated, in the realm of pattern formation "patterns" are guessed and then verified.

1.2 Intrinsic Emergence

For these reasons, pattern formation is insufficient to capture the essential aspect of the emergence of coordinated behavior and global information processing in, for example, flocking birds, schooling fish, ant colonies, and in color and shape perception. At some basic level, though, pattern formation must play a role. The problem is that the "newness" in the emergence of pattern is always referred outside the system to some observer that anticipates the structures via a fixed palette of possible regularities. By way of analogy with a communication channel, the

observer is a receiver that already has the codebook in hand. Any signal sent down the channel that is not already decodable using it is essentially noise, a pattern unrecognized by the observer.

When a new state of matter emerges from a phase transition, for example, initially no one knows the governing “order parameter”. This is a recurrent conundrum in condensed matter physics, since the order parameter is the foundation for analysis and, even, further experimentation. After an indeterminant amount of creative thought and mathematical invention, one is sometimes found and then verified as appropriately capturing measurable statistics. The physicists’ codebook is extended in just this way.

In the emergence of coordinated behavior, though, there is a closure in which the patterns that emerge are important *within* the system. That is, those patterns take on their “newness” with respect to other structures in the underlying system. Since there is no external referent for novelty or pattern, we can refer to this process as “intrinsic” emergence. Competitive agents in an efficient capital market control their individual production-investment and stock-ownership strategies based on the optimal pricing that has emerged from their collective behavior. It is essential to the agents’ resource allocation decisions that, through the market’s collective behavior, prices emerge that are accurate signals “fully reflecting” all available information.

What is distinctive about intrinsic emergence is that the patterns formed confer additional functionality which supports global information processing. Recently, examples of this sort have fallen under the rubric of “emergent computation”.[18] The approach here differs in that it is based on explicit methods of detecting computation embedded in nonlinear processes. More to the point, the hypothesis in the following is that during intrinsic emergence there is an increase in intrinsic computational capability, which can be capitalized on and so lends additional functionality.

In summary, three notions will be distinguished:

1. The intuitive definition of emergence: “something new appears”;
2. Pattern formation: an observer identifies “organization” in a dynamical system; and
3. Intrinsic emergence: the system itself capitalizes on patterns that appear.

2 What’s in a Model?

In moving from the initial intuitive definition of emergence to the more concrete notion of pattern formation and ending with intrinsic emergence, it became clear that the essential novelty involved had to be referred to some evaluating entity. The relationship between novelty and its evaluation can be made explicit by thinking always of some observer that builds a model of a process from a series of measurements. At the level of the intuitive definition of emergence, the observer is that which recognizes the “something” and evaluates its “newness”. In pattern formation, the observer is the scientist that uses prior concepts — e.g. “spiral” or “vortex” — to detect structure in experimental data and so to verify or falsify their applicability to the phenomenon. Of the three, this case is probably the most familiarly appreciated in terms of an “observer” and its “model”. Intrinsic emergence is more subtle. The closure of “newness” evaluation pushes the observer inside the system. This requires in turn that intrinsic emergence be defined in terms of the “models” embedded in the observer. The observer in this view is

a subprocess of the entire system. In particular, it is one that has the requisite information processing capability with which to take advantage of the emergent patterns.

“Model” is being used here in a sense that is somewhat more generous than found in daily scientific practice. There it often refers to an explicit representation — an analog — of a system under study. Here models will be seen in addition as existing implicitly in the dynamics and behavior of a process. Rather than being able to point to (say) an agent’s model of its environment, one may have to excavate the “model”. To do this one might infer that an agent’s responses are in co-relation with its environment, that an agent has memory of the past, that the agent can make decisions, and so on. Thus, “model” here is more “behavioral” than “cognitive”.

3 The Modeling Dilemma

The utility of this view of intrinsic emergence depends on answering a basic question: How does an observer understand the structure of natural processes? This includes both the scientist studying nature and an organism trying to predict aspects of its environment in order to survive. The answer requires stepping back to the level of pattern formation.

A key modeling dichotomy that runs throughout all of science is that between order and randomness. Imagine a scientist in the laboratory confronted after days of hard work with the results of a recent experiment — summarized prosaically as a simple numerical recording of instrument responses. The question arises, What fraction of the particular numerical value of each datum confirms or denies the hypothesis being tested and how much is essentially irrelevant information, just “noise” or “error”?

This dichotomy is probably clearest within science, but it is not restricted to it. In many ways, this caricature of scientific investigation gives a framework for understanding the necessary balance between order and randomness that appears whenever there is an “observer” trying to detect structure or pattern in its environment. The general puzzle of discovery then is: Which part of a measurement series does an observer ascribe to “randomness” and which part to “order” and “predictability”? Aren’t we all in our daily activities to one extent or another “scientists” trying to ferret out the usable from the unusable information in our lives?

Given this basic dichotomy one can then ask: How does an observer actually make the distinction? The answer requires understanding how an observer models data — that is, the method by which elements in a representation, a “model”, are justified in terms of given data.

A fundamental point is that *any* act of modeling makes a distinction between data that is accounted for — the ordered part — and data that is not described — the apparently random part. This distinction might be a null one: for example, for either completely predictable or ideally random (unstructured) sources the data is explained by one descriptive extreme or the other. Nature is seldom so simple. It appears that natural processes are an amalgam of randomness and order. In our view it is the organization of the interplay between order and randomness that makes nature “complex”. A complex process then differs from a “complicated” process, a large system consisting of very many components, subsystems, degrees of freedom, and so on. A complicated system — such as an ideal gas — needn’t be complex, in the sense used here. The ideal gas has no structure. Its microscopic dynamics are accounted for by randomness.

Experimental data is often described by a whole range of candidate models that are statistically and structurally consistent with the given data set. One important variation over this range of possible “explanations” is where each candidate draws the randomness-order distinction. That is, the models vary in the regularity captured and in the apparent error each induces.

It turns out that a balance between order and randomness can be reached and used to define a “best” model for a given data set. The balance is given by minimizing the model’s size while minimizing the amount of apparent randomness. The first part is a version of Occam’s dictum: causes should not be multiplied beyond necessity. The second part is a basic tenet of science: obtain the best prediction of nature. Neither component of this balance can be minimized alone, otherwise absurd “best” models would be selected. Minimizing the model size alone leads to huge error, since the smallest (null) model captures no regularities; minimizing the error alone produces a huge model, which is simply the data itself and manifestly not a useful encapsulation of what happened in the laboratory. So both model size and the induced error must be minimized together in selecting a “best” model. Typically, the sum of the model size and the error are minimized.[19–23]

From the viewpoint of scientific methodology the key element missing in this story of what to do with data is how to measure structure or regularity. (A particular notion of structure based on computation will be introduced shortly.) Just how structure is measured determines where the order-randomness dichotomy is set. This particular problem can be solved in principle: we take the size of the candidate model as the measure of structure. Then the size of the “best” model is a measure of the data’s intrinsic structure. If we believe the data is a faithful representation of the raw behavior of the underlying process, this then translates into a measure of structure in the natural phenomenon originally studied.

Not surprisingly, this does not really solve the problem of quantifying structure. In fact, it simply elevates it to a higher level of abstraction. Measuring structure as the length of the description of the “best” model assumes one has chosen a language in which to describe models. The catch is that this representation choice builds in its own biases. In a given language some regularities can be compactly described, in others the same regularities can be quite baroquely expressed. Change the language and the same regularities could require more or less description. And so, lacking prior God-given knowledge of the appropriate language for nature, a measure of structure in terms of the description length would seem to be arbitrary.

And so we are left with a deep puzzle, one that precedes measuring structure: How is structure discovered in the first place? If the scientist knows beforehand the appropriate representation for an experiment’s possible behaviors, then the amount of that kind of structure can be extracted from the data as outlined above. In this case, the prior knowledge about the structure is verified by the data if a compact, predictive model results. But what if it is not verified? What if the hypothesized structure is simply not appropriate? The “best” model could be huge or, worse, appear upon closer and closer analysis to diverge in size. The latter situation is clearly not tolerable. An infinite model is impractical to manipulate. These situations indicate that the behavior is so new as to not fit (finitely) into current understanding. Then what do we do?

This is the problem of “innovation”. How can an observer ever break out of inadequate model classes and discover appropriate ones? How can incorrect assumptions be changed? How

is anything new ever discovered, if it must always be expressed in the current language?

If the problem of innovation can be solved, then, as all of the preceding development indicated, there is a framework which specifies how to be quantitative in detecting and measuring structure.

4 Where is Science Now?

Contemporary physics does not have the tools to address the problems of innovation, the discovery of patterns, or even the practice of modeling itself, since there are no physical principles that define and dictate how to measure natural structure. It is no surprise, though, that physics does have the tools for detecting and measuring complete order — equilibria and fixed point or periodic behavior — and ideal randomness — via temperature and thermodynamic entropy or, in dynamical contexts, via the Shannon entropy rate and Kolmogorov complexity.

For example, a physicist can analyze the dynamics of a box of gas and measure the degree of disorder in the molecular motion with temperature and the disorganization of the observed macroscopic state in terms of the multiplicity of associated microstates; that is, with the thermodynamic entropy. But the physicist has no analogous tools for deducing what mechanisms in the system maintain the disorder.

Then again, the raw production of information is just one aspect of a natural system's behavior. There are other important contributors to how nature produces patterns, such as how much memory of past behavior is required and how that memory is organized to support the production of information. Information processing in natural systems is a key attribute of their behavior and also how science comes to understand the underlying mechanisms.

The situation is a bit worse than a lack of attention to structure. Physics does not yet have a systematic approach to analyzing the complex information architectures embedded in patterns and processes that occur between order and randomness. This is, however, what is most needed to detect and quantify structure in nature.

The theories of phase transitions and, in particular, critical phenomena do provide mathematical hints at how natural processes balance order and randomness in that they study systems balancing different thermodynamic phases. Roughly speaking, one can think of crystalline ice as the ordered regime and of liquid water as the (relatively) disordered regime of the same type of matter (H_2O). At the phase transition, when both phases coexist, the overall state is more complex than either pure phase. What these theories provide is a set of coarse tools that describe large-scale statistical properties. What they lack are the additional, more detailed probes that would reveal, for example, the architecture of information processing embedded in those states; namely, the structure of those complex thermodynamic states. In fact, modern nonequilibrium thermodynamics can now describe the dominance of collective “modes” that give rise to the complex states found close to certain phase transitions.[24,25] What is still needed, though, is a definition of structure and way to detect and to measure it. This would then allow us to analyze, model, and predict complex systems at the “emergent” scales.

5 A Computational View of Nature

One recent approach is to adapt ideas from the theory of discrete computation, which has developed measures of information processing structure.[26] Computation theory defines the notion of a “machine” — a device for encoding the structures in discrete processes. It has been argued that, due to the inherent limitations of scientific instruments, all an observer can know of a process in nature is a discrete-time, discrete-space series of measurements. Fortunately, this is precisely the kind of thing — strings of discrete symbols, a “formal” language — that computation theory analyzes for structure.

How does this apply to nature? Given a discrete series of measurements from a process, a machine can be constructed that is the best description or predictor of this discrete time series. The structure of this machine can be said to be the best approximation to the original process’s information-processing structure, using the model size and apparent error minimization method discussed above. Once we have reconstructed the machine, we can say that we understand the structure of the process.

But what kind of structure is it? Has machine reconstruction discovered patterns in the data? Computation theory answers such questions in terms of the different classes of machines it distinguishes. There are machine classes with finite memory, those with infinite one-way stack memory, those with first-in first-out queue memory, and those with infinite random access memory, among others. When applied to the study of nature, these machine classes reveal important distinctions among natural processes. In particular, the computationally distinct classes correspond to different types of pattern or regularity.

Given this framework, one talks about the structure of the original process in terms of the complexity of the reconstructed machine. This is a more useful notion of complexity than measures of randomness, such as the Kolmogorov complexity, since it indicates the degree to which information is processed in the system, which accords more closely to our intuitions about what complexity should mean. Perhaps more importantly, the reconstructed machine describes *how* the information is processed. That is, the architecture of the machines themselves represents the organization of the information processing, that is, the intrinsic computation. The reconstructed machine is a model of the mechanisms by which the natural process manipulates information.

6 Computational Mechanics: Beyond Statistics, Toward Structure

Reference [1] reviews how a machine can be reconstructed from a series of discrete measurements of a process. Such a reconstruction is a way that an observer can model its environment. In the context of biological evolution, for example, it is clear that to survive agents must detect regularities in their environment. The degree to which an agent can model its environment in this way depends on its own computational resources and on what machine class or language it implicitly is restricted to or explicitly chooses when making a model. Reference [1] also shows how an agent can jump out of its original assumptions about the model class and, by induction, can leap to a new model class which is a much better way of understanding its

environment. This is a formalization of what is colloquially called “innovation”. The inductive leap itself follows a hierarchical version of machine reconstruction.

The overall goal, then, concerns how to detect structures in the environment — how to form an “internal model” — and also how to come up with true innovations to that internal model. There are applications of this approach to time series analysis and other areas, but the main goal is not engineering but scientific: to understand how structure in nature can be detected and measured and, for that matter, discovered in the first place as wholly new innovations in one’s assumed representation.

What is new in this approach? Computation theorists generally have not applied the existing structure metrics to natural processes. They have mostly limited their research to analyzing scaling properties of computational problems; in particular, to how difficulty scales in certain information processing tasks. A second aspect computation theory has dealt with little, if at all, is measuring structure in stochastic processes. Stochastic processes, though, are seen throughout nature and must be addressed at the most basic level of a theory of modeling nature. The domain of computation theory — pure discreteness, uncorrupted by noise — is thus only a partial solution. Indeed, the order-randomness dichotomy indicates that the interpretation of any experimental data has an intrinsic probabilistic component which is induced by the observer’s choice of representation. As a consequence probabilistic computation must be included in any structural description of nature. A third aspect computation theory has considered very little is measuring structure in processes that are extended in space. A fourth aspect it has not dealt with traditionally is measuring structure in continuous-state processes. If computation theory is to form the foundation of a physics of structure, it must be extended in at least these three ways. These extensions have engaged a number of workers in dynamical systems recently, but there is much still to do.[26–32]

7 The Calculi of Emergence

Ref. [1] focuses on temporal information processing and the first two extensions — probabilistic and spatial computation — assuming that the observer is looking at a series of measurements of a continuous-state system whose states an instrument has discretized. The phrase “calculi of emergence” in its title emphasizes the tools required to address the problems which intrinsic emergence raises. The tools are (i) dynamical systems theory with its emphasis on the role of time and on the geometric structures underlying the increase in complexity during a system’s time evolution, (ii) the notions of mechanism and structure inherent in computation theory, and (iii) inductive inference as a statistical framework in which to detect and innovate new representations.

First, Ref. [1] defines a complexity metric that is a measure of structure in the way discussed above. This is called “statistical complexity”, and it measures the structure of the minimal machine reconstructed from observations of a given process in terms of the machine’s size. Second, it describes an algorithm — ϵ -machine reconstruction — for reconstructing the machine, given an assumed model class. Third, it describes an algorithm for innovation — called “hierarchical machine reconstruction” — in which an agent can inductively jump to a new model class. Roughly speaking, hierarchical machine reconstruction detects regularities in a *series* of

increasingly-accurate models. The inductive jump to a higher computational level occurs by taking those regularities as the new representation. The bulk of Ref. [1] analyzes several examples in which these general ideas are put into practice to determine the intrinsic computation in continuous-state dynamical systems, recurrent hidden Markov models, and cellular automata. It concludes with a summary of the implications of this approach to detecting and understanding structure in nature.

The goal throughout is a more refined appreciation of what “emergence” is, both when new computational structure appears over time and when agents with improved computational and modeling ability evolve. The interplay between computation, dynamics, and induction emphasizes a trinity of conceptual tools required for studying the emergence of complexity; presumably this is a setting that has a good chance of providing empirical application.

8 Discovery versus Emergence

The arguments and development turn on distinguishing several different levels of interpretation: (i) a system behaves, (ii) that behavior is modeled, (iii) an observer detects regularities and builds a model based on prior knowledge, (iv) a collection of agents model each other and their environment, and (v) scientists create artificial universes and try to detect the change in computational capability by constructing their own models of the emergent structures. It is all too easy to conflate two or more of these levels, leading to confusion or, worse, subtle statements seeming vacuous.

It is helpful to draw a distinction between discovery and emergence. The level of pattern formation and the modeling framework of computational mechanics concern discovery. Above it was suggested that innovation based on hierarchical machine reconstruction is one type of discovery, in the sense that new regularities across increasingly-accurate models are detected and then taken as a new basis for representation. Discovery, though, is not the same thing as emergence, which at a minimum is dynamical: over time, or over generations in an evolutionary system, something new appears. Discovery, in this sense, is atemporal: the change in state and increased knowledge of the observer are not the focus of the analysis activity; the products of model fitting and statistical parameter estimation are.

In contrast, emergence concerns the *process* of discovery. Moreover, intrinsic emergence puts the subjective aspects of discovery *into* the system under study. In short, emergence pushes the semantic stack down one level. In this view analyzing emergence is more objective than analyzing pattern formation in that detecting emergence requires modeling the dynamics of discovery, not just implementing a discovery procedure.

The arguments to this point can be recapitulated by an operational definition of emergence. A process undergoes emergence if at some time the architecture of information processing has changed in such a way that a distinct and more powerful level of intrinsic computation has appeared that was not present in earlier conditions.

It seems, upon reflection, that our intuitive notion of emergence is not captured by the “intuitive definition” given in the first section. Nor is it captured by the somewhat refined notion of pattern formation. “Emergence” is meaningless unless it is defined within the context of processes themselves; the only well-defined notion of emergence would seem to be intrinsic

emergence. Why? Simply because emergence defined without this closure leads to an infinite regress of observers detecting patterns of observers detecting patterns This is not a satisfactory definition, since it is not finite. The regress must be folded into the system, it must be immanent in the dynamics. When this happens complexity and structure are no longer referred outside, no longer relative and arbitrary; they take on internal meaning and functionality.

9 Evolutionary Mechanics

Where in science might a theory of intrinsic emergence find application? Are there scientific problems that at least would be clarified by the computational view of nature outlined here?

In several ways the contemporary debate on the dominant mechanisms operating in biological evolution seems ripe. Is anything ever new in the biological realm? The empirical evidence is interpreted as a resounding “yes”. It is often heard that organisms today are more complex than in earlier epochs. But how did this emergence of complexity occur? Taking a long view, at present there appear to be three schools of thought on what the guiding mechanisms are in Darwinian evolution that produce the present diversity of biological structure and that are largely responsible for the alteration of those structures.

Modern evolutionary theory continues to be governed by Darwin’s view of the natural selection of individuals that reproduce with variation. This view emphasizes the role of fitness selection in determining which biological organisms appear. But there are really two camps: the Selectionists, who are Darwin’s faithful heirs now cognizant of genetics, and the Historicists, who espouse a more anarchistic view.

The Selectionists hold that structure in the biological world is due primarily to the fitness-based selection of individuals in populations whose diversity is maintained by genetic variation.[33] In a sense, genetic variation is a destabilizing mechanism that provides the raw diversity of structure. Natural selection then is a stabilizing dynamic that acts on the expression of that variation. It provides order by culling individuals based on their relative fitness. This view identifies a source of new structures and a mechanism for altering one form into another. The adaptiveness accumulated via selection is the dominant mechanism guiding the appearance of structure.

The second, anarchistic camp consists of the Historicists who hold fast to the Darwinian mechanisms of selection and variation, but emphasize the accidental determinants of biological form.[34,35] What distinguishes this position from the Selectionists is the claim that major changes in structure can be and have been nonadaptive. While these changes have had the largest effect on the forms of present day life, at the time they occurred they conferred no survival advantage. Furthermore, today’s existing structures needn’t be adaptive. They reflect instead an accidental history. One consequence is that a comparative study of parallel earths would reveal very different collections of life forms on each. Like the Selectionists, the Historicists have a theory of transformation. But it is one that is manifestly capricious or, at least, highly stochastic with few or no causal constraints. For this process of change to work the space of biological structures must be populated with a high fraction which are functional.

Lastly, there are the Structuralists whose goal is to elucidate the “principles of organization” that guide the appearance of biological structure. They contend that energetic, mechanical,

biomolecular, and morphogenetic constraints limit the infinite range of possible biological form.[16,36–40] The constraints result in a relatively small set of structure archetypes. These are something like the Platonic solids in that they pre-exist, before any evolution takes place. Natural selection then plays the role of choosing between these “structural attractors” and possibly fine-tuning their adaptiveness. Darwinian evolution serves, at best, to fill the waiting attractors or not depending on historical happenstance. Structuralists offer up a seemingly testable claim about the ergodicity of evolutionary processes: given an ensemble of earths, life would have evolved to a similar collection of biological structures.

The Structuralist tenets are at least consistent with modern thermodynamics.[24,25] In large open systems energy enters at low entropy and is dissipated. Open systems organize largely due to the reduction in the number of active degrees of freedom caused by the dissipation. Not all behaviors or spatial configurations can be supported. The result is a limitation of the collective modes, cooperative behaviors, and coherent structures that an open system can express. The Structuralist view is a natural interpretation of the many basic constraints on behavior and pattern indicated by physics and chemistry. For example, the structures formed in open systems such as turbulent fluid flows, oscillating chemical reactions, and morphogenetic systems are the product of this type of macroscopic pattern formation. Thus, open systems offer up a limited palette of structures to selection. The more limited the palette, the larger the role for “principles of organization” in guiding the emergence of life as we know it.

What is one to think of these conflicting theories of the emergence of biological structure? In light of the preceding sections there are several impressions that the debate leaves an outsider with.

1. Natural selection’s culling of genetic variation provides the Selectionists with a theory of transformation. But the approach does not provide a theory of structure. Taking the theory at face value, in principle one can estimate the time it takes a *given* organism to change. But what is the mean time under the evolutionary dynamic and under the appropriate environmental pressures for a hand to appear on a fish? To answer this one needs a measure of the structure concerned and of the functionality it does or does not confer.
2. The Historicists also have a theory of transformation, but they offer neither a theory of structure nor, apparently, a justification for a high fraction of functionality over the space of structures. Perhaps more disconcerting, though, in touting the dominance of historical accident, the Historicists advocate an antiscientific position. This is not to say that isolated incidents do not play a role; they certainly do. But it is important to keep in mind that the event of a meteor crashing into the earth is extra-evolutionary. The explanation of its occurrence is neither the domain of evolutionary theory nor is its occurrence likely ever to be explained by the principles of dynamics: it just happened, a consequence of particular initial conditions. Such accidents impose constraints; they are not an explanation of the biological response.
3. In complementary fashion, the Structuralists do not offer a theory of transformation. Neither do they, despite claims for the primacy of organization in evolutionary processes, provide a theory of structure itself. In particular, the structure archetypes are not analyzed in terms of their internal components nor in terms of system-referred functionality. Considering

these lacks, the Structuralist hope for “deep laws” underlying biological organization is highly reminiscent of Chomsky’s decades-long search for “deep structures” as linguistic universals, without a theory of cognition. The ultimate failure of this search[41] suggests a reconsideration of fundamentals rather than optimistic forecasts of Structuralist progress.

The overwhelming impression this debate leaves, though, is that there is a crying need for a theory of biological structure and a qualitative dynamical theory of its emergence.[42] In short, the tensions between the positions are those (i) between the order induced by survival dynamics and the novelty of individual function and (ii) between the disorder of genetic variation and the order of developmental processes. Is it just an historical coincidence that the structuralist-selectionist dichotomy appears analogous to that between order and randomness in the realm of modeling? The main problem, at least to an outsider, does not reduce to showing that one or the other view is correct. Each employs compelling arguments and often empirical data as a starting point. Rather, the task facing us reduces to developing a synthetic theory that balances the tensions between the viewpoints. Ironically, evolutionary processes themselves seem to do just this sort of balancing, dynamically.

The computational mechanics of nonlinear processes can be construed as a theory of structure. Pattern and structure are articulated in terms of various types of machine classes. The overall mandate is to provide both a qualitative and a quantitative analysis of natural information processing architectures. If computational mechanics is a theory of structure, then innovation via hierarchical machine reconstruction is a computation-theoretic approach to the transformation of structure. It suggests one mechanism with which to study what drives and what constrains the appearance of novelty. The next step, of course, would be to fold hierarchical machine reconstruction into the system itself, resulting in a dynamics of innovation, the study of which might be called “evolutionary mechanics”.

10 The Mechanics

By way of summarizing the main points, let’s question the central assumption of this approach to emergence.

Why talk about “mechanics”? Aren’t mechanical systems lifeless, merely the sum of their parts? One reason is simply that scientific explanations must be given in terms of mechanisms. Explanations and scientific theories without an explicit hypothesis of the underlying causes — the mechanisms — are neither explanations nor theories, since they cannot claim to entail falsifiable predictions.[43] Another, more constructive reason is that modern mathematics and physics have made great strides this century in extending the range of Newtonian mechanics to ever more complex processes. When computation is combined with this, one has in hand a greatly enriched notion of mechanism.

It might seem implausible that an abstract “evolutionary mechanics” would have anything to contribute to (say) biological evolution. A high-level view at least suggests a fundamental, if indirect, role. By making a careful accounting of where the observer and system-under-study are located in various theories of natural phenomena, a certain regularity appears which can be summarized by a hierarchy of mechanics. The following list is given in the order of increasing attention to the context of observation and modeling in a classical universe. The first two are

already part of science proper; the second two indicate how computation and innovation build on them.

1. Deterministic mechanics (dynamical systems theory): The very notions of cause and mechanism are defined in terms of state space structures. This is Einstein's level: the observer is entirely outside the system-under-study.
2. Statistical mechanics (probability theory): Statistical mechanics is engendered by deterministic mechanics largely due to the emergence of irreducible uncertainty. This occurs for any number of reasons. First, deterministic mechanical systems can be very large, too large in fact to be usefully described in complete detail. Summarizing the coarse, macroscopic properties is the only manageable goal. The calculus for managing the discarded information is probability theory. Second, deterministic nonlinear systems can be chaotic, communicating unseen and uncontrollable microscopic information to affect observable behavior.[44] Both of these reasons lead to the necessity of using probabilistic summaries of deterministic behavior to collapse out the irrelevant and accentuate the useful.
3. Computational mechanics (theory of structure for statistical mechanics): As discussed at some length, it is not enough to say that a system is random or ordered. What is important is how these two elements, and others, interact to produce complex systems. The information processing mechanisms distinguished by computation theory give a (partial) basis for being more objective about detecting structure, quantifying complexity, and the modeling activity itself.
4. Evolutionary mechanics (dynamical theory of innovation): As noted above, evolutionary mechanics concerns how genuine novelty occurs. This is the first level at which emergence takes on its intrinsic aspect. Building on the previous levels, the goal is to delineate the constraints guiding and the forces driving the emergence of complexity.

A typical first question about this hierarchy is, Where is quantum mechanics? The list just given assumes a classical physical universe. Therefore, quantum mechanics is not listed despite its undeniable importance. It would appear, however, either as the most basic mechanics, preceding deterministic mechanics, or at the level of statistical mechanics, since that is the level at which probability first appears. In a literal sense quantum mechanics is a theory of the deterministic dynamics of complex "probabilities" that can interfere over spacetime. The interference leads to new phenomena, but the goals of and manipulations used in quantum mechanics are not so different from that found in stochastic processes and so statistical mechanics. My own prejudice in these issues will be resolved once a theory of measurement of nonlinear processes is complete. There are several difficulties that lie in the way. The effect of measurement distortion can be profound, for example, leading to irreducible indeterminacy in completely deterministic systems.[31]

So is anything ever new? I would answer "most definitely". With careful attention to the location of the observer and the system-under-study, with detailed accounting of intrinsic computation, with quantitative measures of complexity, we can analyze the patterns, structures, and novel information processing architectures that emerge in nonlinear processes. In this way, we demonstrate that something new has appeared.

11 Acknowledgments

Many thanks are due to Melanie Mitchell for a critique of Ref. [1] that led to the present essay. The author is also indebted to Dan McShea for his thoughts on the evolutionary trichotomy. Comments from Lisa Borland, Don Glaser, Jim Hanson, Blake LeBaron, and Dan Upper are gratefully acknowledged.

Bibliography

1. J. P. Crutchfield. The calculi of emergence: Computation, dynamics, and induction. *Physica D*, 1994. in press; Santa Fe Institute Report SFI-93-11.
2. C. W. Reynolds. Flocks, herds, and schools: A distributed behavioral model. *Computer Graphics*, 21:25 – 34, 1987.
3. B. Holldobler and E. O. Wilson. *The Ants*. Belknap Press of Harvard University Press, Cambridge, Massachusetts, 1990.
4. E. F. Fama. Efficient capital markets II. *J. Finance*, 46:1575 – 1617, 1991.
5. E. Land. The retinex. *Am. Scientist*, 52:247 – 264, 1964.
6. B. A. Wandell. Color appearance: The effects of illumination and spatial pattern. *Proc. Nat. Acad. Sci.*, 10:2458 – 2470, 1993.
7. I. Kovacs and B. Julesz. A closed curve is much more than an incomplete one - effect of closure in figure ground segmentation. *Proc. Nat. Acad. Sci.*, 90:7495 – 7497, 1993.
8. J. P. Crutchfield, N. H. Packard, J. D. Farmer, and R. S. Shaw. Chaos. *Sci. Am.*, 255:46, 1986.
9. H. L. Swinney and J. P. Gollub, editors. *Hydrodynamic Instabilities and the Transition to Turbulence*, Berlin, 1981. Springer Verlag.
10. A. M. Turing. The chemical basis of morphogenesis. *Trans. Roy. Soc., Series B*, 237:5, 1952.
11. A. T. Winfree. *The Geometry of Biological Time*. Springer-Verlag, Berlin, 1980.
12. Q. Ouyang and H. L. Swinney. Transition from a uniform state to hexagonal and striped turing patterns. *Nature*, 352:610 – 612, 1991.
13. H. E. Stanley. *Introduction to Phase Transitions and Critical Phenomena*. Oxford University Press, Oxford, 1971.
14. P. Bak and K. Chen. Self-organized criticality. *Physica A*, 163:403 – 409, 1990.
15. J. J. Binney, N. J. Dowrick, A. J. Fisher, and M. E. J. Newman. *The Theory of Critical Phenomena*. Oxford University Press, Oxford, 1992.
16. D. W. Thompson. *On Growth and Form*. Cambridge University Press, Cambridge, 1917.
17. H. Meinhardt. *Models of Biological Pattern Formation*. Academic Press, London, 1982.
18. S. Forrest, editor. *Emergent Computation: Self-organizing, collective, and cooperative behavior in natural and artificial computing networks: Introduction to the Proceedings of the Ninth Annual CNLS Conference*, Amsterdam, 1990. North-Holland.
19. J. G. Kemeny. The use of simplicity in induction. *Phil. Rev.*, 62:391, 1953.

20. C. S. Wallace and D. M. Boulton. An information measure for classification. *Comput. J.*, 11:185, 1968.
21. J. Rissanen. Modeling by shortest data description. *Automatica*, 14:462, 1978.
22. J. P. Crutchfield and B. S. McNamara. Equations of motion from a data series. *Complex Systems*, 1:417, 1987.
23. J. Rissanen. *Stochastic Complexity in Statistical Inquiry*. World Scientific, Singapore, 1989.
24. G. Nicolis and I. Prigogine. *Self-Organization in Nonequilibrium Systems*. Wiley, New York, 191977.
25. H. Haken. *Synergetics, An Introduction*. Springer, Berlin, third edition, 1983.
26. J. P. Crutchfield and K. Young. Inferring statistical complexity. *Phys. Rev. Lett.*, 63:105, 1989.
27. S. Wolfram. Computation theory of cellular automata. *Comm. Math. Phys.*, 96:15, 1984.
28. L. Blum, M. Shub, and S. Smale. On a theory of computation over the real numbers. *Bull. AMS*, 21:1, 1989.
29. M. G. Nordahl. Formal languages and finite cellular automata. *Complex Systems*, 3:63, 1989.
30. J. E. Hanson and J. P. Crutchfield. The attractor-basin portrait of a cellular automaton. *J. Stat. Phys.*, 66:1415, 1992.
31. J. P. Crutchfield. Unreconstructible at any radius. *Phys. Lett. A*, 171:52 – 60, 1992.
32. C. Moore. Real-valued, continuous-time computers: A model of analog computation, part I. Technical Report 93-04-018, Santa Fe Institute, 1993.
33. J. Maynard-Smith. *Evolutionary Genetics*. Oxford University Press, Oxford, 1989.
34. J. Monod. *Chance and Necessity: An Essay on the Natural Philosophy of Modern Biology*. Vintage Books, New York, 1971.
35. S. J. Gould. *Wonderful Life*. Norton, New York, 1989.
36. C. H. Waddington. *The Strategy of the Genes*. Allen and Unwin, London, 1857.
37. B. Goodwin. Evolution and the generative order. In B. Goodwin and P. Sanders, editors, *Theoretical Biology: Epigenetic and Evolutionary Order from Complex Systems*, pages 89 – 100, Baltimore, Maryland, 1992. Johns Hopkins University Press.
38. S. A. Kauffman. *Origins of Order: Self-Organization and Selection in Evolution*. Oxford University Press, New York, 1993.
39. W. Fontana and L. Buss. What would be conserved if the tape were played twice? *Proc. Nat. Acad. Sci.*, 91:757 – 761, 1994.
40. W. Fontana and L. Buss. “The Arrival of the Fittest”: Toward a theory of biological organization. *Bull. Math. Bio.*, 56:1 – 64, 1994.
41. R. A. Harris. *The Linguistic Wars*. Oxford University Press, New York, 1993.
42. B. Goodwin and P. Sanders, editors. *Theoretical Biology: Epigenetic and Evolutionary Order from Complex Systems*, Baltimore, Maryland, 1992. Johns Hopkins University Press.
43. K. R. Popper. *The Logic of Scientific Inquiry*. Basic Books, New York, 1959.
44. R. Shaw. Strange attractors, chaotic behavior, and information flow. *Z. Naturforsch.*, 36a:80, 1981.