Editorial for the Special Issue on ‘Emergent Properties of Complex Systems’

Emergence and Levels of Abstraction

R.I. Damper

Image, Speech and Intelligent Systems (ISIS) Research Group

Department of Electronics and Computer Science

University of Southampton

Southampton S017 1BJ

United Kingdom

Tel: +44 1703 594577

FAX: +44 1703 594498

Email: `rid@ecs.soton.ac.uk`

Abstract

The notions of *emergence* and *emergent properties* have a long history in science, and have recently regained popularity in systems science fuelled largely by the growth of computer simulation as an exploratory and investigative tool. Unfortunately, the notions and terms are not especially well defined: ideas of evolution, self-organization, collective (‘systemic’) properties and cooperative behaviour are all involved to a greater or lesser extent. It is often claimed that emergent properties arise at a particular level of system description by virtue of the interaction of relatively simple lower-level components – between themselves and with the environment – but cannot be explained at this lower level. Yet there are obvious scientific and philosophical problems with a definition based on an *inability* to explain observable effects in particular terms. This editorial outlines the history of emergence as a scientific concept, and reviews attempts to refine and qualify the term.

1 Introduction

This Special Issue of the *Journal* focuses on a topic of considerable and growing interest in systems science – emergent properties of complex systems. So what is an emergent property, and why should scientists in general, and systems scientists in particular, be interested in such a thing?

According to de Mey (1982, p. 149), Karl Popper "stands out as the strongest defender of a view which sees science as the product of an inexorable methodology". An important part of this methodology (e.g. Popper 1965) is the use of *models* (or theories) which are judged by objective criteria such as their parsimony, ability to explain the results of experiment or careful observation, and their predictive power which thereby offers possibilities to falsify the model and/or discover new empirical knowledge. Probably most present-day scientists hold – unthinkingly if not consciously – to this view. This emphasises the importance to the methodology of exploring the predictive consequences of scientific models and theories. But how is this to be done . . . methodically?

One of the most notable changes in the practice of science over the last decade or two has been the increasing availability of high-power digital computers. Whereas earlier generations of investigators had little but their insight and imagination (plus pen and paper) with which to work out the predictive capabilities of their theories, today's scientists can readily explore the implications of their models by computer simulation. This almost revolutionary aspect of the changing face of science has recently been well reviewed by Casti (1997). Indeed, it could be argued that the science of *complex systems* only really came into existence with easy access to computers, allowing high-speed simulation of reasonably large collections of interacting and/or cooperating components or subsystems.

In this editorial, I first introduce the basic notions of emergence before sketching a brief history of the subject. A central tenet is that emergence has to do with qualitatively different kinds of description being appropriate to different levels of abstraction. This raises the question of how levels relate to each other and what is *fundamental*, which is dealt with next. The qualitatively different kinds of behaviour and phenomena which emerge at higher levels of the hierarchy are frequently described in terms of surprise or unpredicability. The next section argues that this is an unsatisfactory description and lays the foundations for the review of recent attempts to refine and qualify notions of emergence which follows. A recurrent theme is that emergence as a principle seems opposed to reductionism – a widely accepted doctrine which has proved enormously fruitful in science. Accordingly, the relation between the two is explored before concluding.

2 Emergence: Basic Ideas

To introduce the topic of emergence, let us consider (to cite Schroeder 1991, p. 35) "one of the most surprising instances of a power law in the humanities", namely Zipf's (1949) law – according to which the frequency of occurrence, f, of words is (approximately) inversely proportional to their rank[1], r, for many natural languages. That is:

$$f \propto \frac{1}{r} \quad \text{and} \quad \log f = -\log r + \text{const.}$$

Hence, a plot of f against r on double logarithmic axes yields a straight line with negative slope[2].

Certain properties of other complex systems of human origin, as well as language, were also found to satisfy a form of this law: For instance, the size of cities is similarly related to

---

[1] By the *rank* of a word, I mean that $r = 1$ for the most frequent word, $r = 2$ for the next most frequent word, etc.

[2] Obviously, this kind of research – which typically involves trawling through masses of data – can be done manually but is greatly facilitated by computers.

their rank. Such a surprise was Zipf's law to its discoverer that he imagined it must distinguish behavioural and social systems from those of purely physical origin, and reflected the operation of some 'intelligent' least effort principle which would not obtain in the inanimate world. The assumption of an underpinning intelligence was, however, demolished by Mandelbrot (1961) who was able to show that a 'language' composed by randomly striking typewriter keys also obeyed Zipf's law. Today, this law is generally recognised as just one of a number of scaling or power laws occurring widely in the natural world (Schroeder 1991; Gell-Mann 1994; Bak 1996; Casti 1997).

For many authors, Zipf's law is a prime example of what today we would call an *emergent property*. For instance, Casti (1997, p. 128–9) writes that this law:

> "… is not a pattern that can be seen in the individual words … but rather emerges from the interaction of the words to form sentences … This is the essence of what is meant by an emergent property … "

This view of emergence is largely shared and supported by other commentators. For instance, in his study of the simple and the complex in nature, Gell-Mann (1994, pp. 99–100) writes more generally:

> "Scientists … are trying hard to understand the ways in which structures arise without the imposition of special requirements from outside. In an astonishing variety of contexts, apparently complex behaviors arise from systems characterized by simple rules. These systems are said to be self-organized and their properties are said to be emergent."

We have made much of the 'surprise' aspect of Zipf's law which is held by some to be indicative of emergence. However, a very simple proof of Mandelbrot's result (for a random ‘language’), relying on nothing other than elementary probability theory, has been given by Li (1992), who states (p. 1844):

> “Zipf’s law is not a deep law in natural language as one might first have thought. It is very much related to the particular representation one chooses, i.e., rank as the independent variable.”

Li goes on to state (presumably because she was able to show its simple origins) that Zipf’s law is somehow different from other power and scaling laws. This seems unwarranted. Rather, I prefer to take Li’s result as indicating that ‘surprising’ emergent phenomena can indeed be explained, sometimes very simply.

Does this hold for other phenomena thought to be ‘emergent’, or is Zipf’s law rather too simple and so just not a very good example of the genre? The views of Crick on this question are relevant (1994, p. 11):

> “... there are two meanings of the term *emergent*. The first has mystical overtones. It implies that the emergent behavior cannot in any way, even in principle, be understood as the combined behavior of its component parts ... The scientific meaning of emergent ... assumes that, while the whole may not be the simple sum of the separate parts, its behavior can, at least in principle, be *understood* from the nature and behavior of its parts *plus* the knowledge of how these parts interact.”

Crick’s first meaning seems to be what Horgan has in mind when he writes (1996, p. 192): “Emergence ... is a hoary idea, related to holism, vitalism, and other antireductionist creeds that date back to the last century at least. Certainly Darwin did not think that natural selection could be derived from Newtonian mechanics."[3] With this perspective – and wishing to avoid "mystical overtones" in favour of "scientific meaning" – a better example than Zipf's law of an emergent phenomenon might be locomotion. Arguably, locomotion in animals is not a property of individual neurons, or muscles or bones, but can be understood by the way these separate parts work together. Typically, this understanding will be enhanced by – or even dependent on – the creation of new and appropriate vocabulary to describe properties of "the whole" which is more than "the simple sum of the parts". That is, a satisfying explanation of walking – as of other emergent phenomena – relies on getting the *level of abstraction* right.

3 A Brief History of Emergence

It would, of course, be wrong to infer from the above, introductory remarks that *emergence* only became an issue in the computer age. Perhaps the earliest hint of the notion comes in David Hume's *Dialogues Concerning Natural Religion* (1779). Hume was concerned with the celebrated and influential *argument from design* (e.g. Matson 1965; Popper 1978; Dawkins 1986), recently paraphrased by Dennett (1995, p. 28) as follows:

> "... among the effects that we can objectively observe in the world, there are many that are not (cannot be ...) mere accidents; they must have been designed to be as they are, and there cannot be design without a Designer; therefore, a Designer, God, must exist ... as the source of all these wonderful effects."

Chief among these "wonderful effects" is the phenomenon we call *life*. Certainly, this is altogether a more complex issue than Zipf's law! For centuries, thinkers have puzzled why some matter is inanimate while other matter displays the attributes of life, in its many diverse

---

[3] One should probably not infer from this that Horgan is antipathetic to the notion of emergence. He goes on to give a more balanced discussion of the pros and cons.

forms. One of Hume's purposes in his treatise (published posthumously for fear of persecution) was to show the logical fallacy of the argument from design as an explanation of order in the universe. He writes: "For aught we know, *a priori*, matter may contain the source, or spring, of order originally, within itself . . ." (cited by Wilson 1999, p. 24). Now what is this other than Gell-Mann's complex behaviour arising "without the imposition of special requirements from outside"? Thus, Hume brilliantly anticipates[4] the scientific revolution to be unleashed 80 years later by Darwin in *The Origin of Species* (1859). It is now firmly established[5] that species arise and evolve by a purely mechanistic process of natural selection, driven by interaction with the environment. There is no logical necessity for "special requirements from outside" in the sense of some global controller, vitalist force, creationist Designer or whatever.

As far as I am aware, Darwin never used the term *emergence* (at least in any specialised scientific or technical sense as implied by our discussion here). It was only in the early decades of the 19th century that the subject started to attract the attention of scientists and philosophers (e.g. Broad 1919; 1925; Pepper 1926). Quite apart from evolutionary biology and biochemistry, where ideas of emergence have been influential (e.g. Kauffman 1987; Eigen 1996; Stein and Varela 1993), the notion has proved popular in such diverse fields as the study of mind and consciousness (Popper 1978; Dennett 1991; Franklin 1995), self-reproducing automata and artificial life (Kauffman 1984; Langton 1986; Cariani 1991; Steels 1991; Levy 1992; Adami 1998), speech and language (Hawkins and Gell-Mann 1992; Guenter and Gjaja 1996; MacWhinney 1998; Sussman, Fruchter, Hilbert, and Sirosh 1998), the theory of computation (Forrest 1990 *et seq.*), robotics (McFarland and Bösser 1993; Arkin 1998) and the

---

[4] This description might have been questioned by Popper (1978, p. 341) who wrote "Hume . . . attacked, *somewhat feebly*, the argument from design" [italics added].

[5] To cite Dennett (1995, p. 21): "The fundamental core of contemporary Darwinism, the theory of DNA-based reproduction and evolution, is now beyond dispute among scientists".

branch of physics that Haken (1997) has called *synergetics*[6].

These are all relatively specialised areas, however. More widely, a major impact on science has been made in the last 10-15 years by connectionist models and thinking. After several years in the doldrums, interest in connectionism and neural nets was revived by authors such as Kohonen (1977), Hopfield (1982) and Rumelhart and McClelland (1986) who have all, in various ways, stressed the notion of emergent, 'collective' properties. (See especially Clark 1989 and Kohonen 1997 in addition.) Yet, as Quinlan (1991, p. 224) has pointed out:

> "... it is difficult to be sure, sometimes, what the intended meaning of the term 'emergent properties' is in the ... connectionist literature. The most neutral interpretation is that nets exhibit properties that are either, in some sense, counterintuitive or that they behave in a way that could not have been predicted at the outset ... The phrase 'emergent properties' can also be interpreted relative to nets exhibiting graceful degradation and automatic generalization. Both ... are observed to emerge out of the intrinsic dynamics of a massively parallel and distributed processing system."

In essence, Quinlan – like Crick – is offering two alternative meanings. Clearly, the first of these (his "most neutral interpretation") has something of Crick's "mystical overtones" about it. For what can it mean that something "could not have been predicted"? (See below for discussion of this question.) The second meaning is much closer to Crick's preferred interpretation whereby a new level of abstraction/description with new vocabulary ('generalisation', 'graceful degradation', ...) is found useful in understanding and exploiting connectionist systems. And what is good for connectionism, I suggest, is good for the whole gamut of complex systems and non-linear science. For it is non-linearity that decrees that the whole

---

[6] According to Haken (1997, p. 190), "one may consider synergetics as a theory dealing with the emergence of new qualities on macroscopic scales".

may exceed the sum of the parts, and complexity that means that collective properties can arise from an ensemble of many parts.

4 What is Fundamental? Physics, Chemistry and Biology

Of course, the idea that emergence is centrally concerned with the role of different levels of abstraction or hierarchies of description – i.e. properties arise at one level which are not discernible at a lower level – is not new. It has been pervasive throughout the history of the topic. For instance, according to Pepper (1926, p. 241) writing many years ago:

> "The theory of emergence involves three propositions: (1) that there are levels of existence . . . (2) that there are marks which distinguish these levels from one another . . . (3) that it is impossible to deduce marks of a higher level from those of a lower level . . . "

Some years later, inspired by Denbigh (1975), Popper (1978, p. 342) wrote:

> "I think that science suggests to us (tentatively of course) a picture of the universe that is inventive or even creative; of a universe in which *new things* emerge, on *new levels*."

This raises the question, considered at some length by Gell-Mann (1994, Chap. 9), of what is the most basic level: *what is fundamental*? He suggests (p. 109) ". . . that science A is more fundamental than science B when

1. The laws of science A encompass in principle the phenomena and laws of science B.

2. The laws of science A are more general than those of science B . . . "

According to this scheme, quantum electrodynamics (QED) is more fundamental than chemistry, because the laws of the latter can (in principle) be derived from the former "provided the additional information describing suitable chemical conditions is fed into the equations; moreover, these conditions are special – they do not hold throughout the universe" (p. 111). This gives a useful way to think about emergence. Phenomena and natural laws at a given level emerge from the operation of laws at a more fundamental level (see also Anderson 1972). Not only this, the lower (more fundamental) levels constrain the higher levels (cf. Holland 1998, pp. 186–188). Successive levels of representations and models (each new level meeting the constraints of the lower levels) depend upon macro-descriptions emerging from the next lowest level (cf. Holland 1998, pp. 195–198).

Currently, string theory is under serious consideration as a 'most fundamental' description of nature. In *The Elegant Universe*, Greene (1999, p. 139) writes:

> "... numerous features of the standard model *[of physics]* – features that had been painstakingly discovered over the course of decades of research – *emerged naturally* and simply from the grand structure of string theory. As Michael Green[7] has said, '... almost all of the major developments in physics over the last hundred years emerge – and emerge with such elegance – from such a simple starting point ... ' " [my italics]

In the case of the standard model of physics, however, there is little necessity to provide Gell-Mann's special conditions in the form of "additional information" because physics remains largely fundamental. This is not the case with chemistry, and even less so with biology. Echoing Anderson's (1972) aphorism "More is different", Georgi (1989, p. 447–8) states "biology is not a branch of physics". He argues:

---

[7] Green is a pioneer of string theory.

> "It *is* true that in chemistry and biology one does not encounter any new physical principles. But the systems on which the old principles act differ in such drastic and qualitative way that in the different fields that it is simply not *useful* to regard one as a branch of another . . . indeed, 'principles' of new kinds must be developed."

But what are we to make of Pepper's assertion "that it is impossible to deduce marks of a higher level from those of a lower level"? Do we not recognise here the "mystical overtones" of Crick and/or the "hoary . . . antireductionist creed" of Horgan?

5 Emergence and Surprise

Many authors have attempted to explain emergence in terms of *surprise*, e.g. Casti (1997, Chap. 3) and the several quotations above, or at least made it a central tenet. The problem here is that surprise is in the eye of the beholder. What was deep and unfathomable to Zipf is merely a simple consequence of probability theory to Li. What might be surprising on first acquaintance or at a particular stage of scientific knowledge tends to become commonplace, trite or predictable after intensive, lengthy study.

Dennett (1995, pp. 412–419) entertainingly addresses what it means for properties to be unpredictable or surprising in his 'Tale of Two Black Boxes'. In this allegory, and omitting some details, Dennett tells the story of scientists trying to explain the workings of a system in which pressing button α on box A causes a red lamp to light on box B while pressing button β on box A causes a green lamp to light on box B – yet the wire connecting the two boxes transmits a *different* bit string each time. After much study, and in spite of the obvious causality at the macrolevel of buttons and lamps, some of the scientists were tempted to call the system properties *emergent* since they appeared "*unpredictable in principle* from . . . analysis of the microproperties of the strings themselves". Cariani's rhetorical question (1991, p. 776) is apposite: "If we randomly come across a computer simulation and we have no clue as to its purpose, can we tell if its computations are emergent?"

Eventually, however, the secret is revealed. Pressing button α causes one of a large number of facts expressed in English (or implications generated from them) to be selected at random from a knowledge base, encoded in ASCII and transmitted to box B which then checks for 'truth' against its own knowledge base (containing almost identical facts but differently expressed and encoded) and lights the red lamp in the case of 'true strings'. Conversely, pressing button β causes a 'false string' to be transmitted, so lighting the green lamp. According to Dennett: "The point of the fable is simple. There is no substitute for the intentional stance; either you adopt it, and explain the pattern by finding the semantic level facts, or you will forever be baffled by the regularity – the causal regularity – that is manifestly there" (p. 421). In other words, we must find the right level of abstraction – that at which the properties become semantically interpretable by virtue of the new concepts and corresponding vocabulary that we employ. And once we have settled on the right, semantically interpretable level of abstraction, our surprise evaporates.

But if surprise evaporates, there is a tendency to feel almost 'cheated'. Critics often say[8] when claims of emergent behaviour are made: "What else could the system do?" But is this criticism as valid as it seems? Franklin (1995, note 10, p. 207) offers what, to me, is a very satisfactory riposte to this kind of remark. Citing the objection[9] to Wilson's (1985) artificial creature 'Animat' that: "He seems … to have been jury-rigged to do what his creators wanted!", Franklin responds:

> "Yes, of course. Every autonomous agent is 'jury-rigged,' that is, designed and/or evolved to couple with its environment or to learn to do so."

---

[8] I have certainly suffered this criticism in anonymous peer reviews of my own work.

[9] Credited to David Lee Larom.

John Holland (personal communication, September 13 1999) has pointed out that von Neumann's demonstration of a self-reproducing machine "nicely refutes the use of *surprise* as part of the definition of emergence". "It's still a great example of emergence", he says, "even though we now know it in detail." And rather than feeling 'cheated' once we know the details, there should be "… no diminution in wonder. Instead a new realm [is] opened, offering new wonders and new questions" Holland (1998, p. 13).

6 Towards a Definition of Emergence

It should be obvious from the foregoing discussion that the term *emergence* does not easily admit of a precise definition. For if we recognise emergence principally by the necessity to define new categories, concepts and descriptive terms appropriate to a semantically interpretable level of abstraction, how will we avoid doing this afresh for each and every case? As Holland (1990, p. 108) writes "… we still have little more than rules of thumb when it comes to the model-*building* process itself". More recently, Holland (1998, p. 3), has described this difficulty in the following words:

> "Despite its ubiquity and importance, emergence is an enigmatic and recondite topic, more wondered at than analyzed … It is unlikely that a topic as complicated as emergence will submit meekly to a concise definition."

In spite of the problems, many authors have attempted definitions or, at least, tried to classify different kinds of emergence. In this section, we review some of the more recent of these attempts.

According to Cariani (1991, p. 775): "The problem of emergence classically involved the origins of qualitatively new structures and functions which were not reducible to those already in existence". He goes on to identify three current tracts of thought on emergence, calling them “computational”, “thermodynamic” and “relative to a model”. Computational emergence is *formally-based*, i.e. it is related to the manifestation of new global forms, such as flocking behaviour and chaos, from local interactions. Thermodynamic emergence is *physically-based* and is concerned with issues such as the origins of life, where order emerges from noise. The emergence relative to a model concept is *functionally-based* and deals with situations where observers need to change their model in order to keep up with a system’s behaviour. The problem that I have with these so-called different forms of emergence is in distinguishing them! The formal/physical/functional division appears to be no more than different views of the same thing. How are the first two kinds distinguished from the last? If, for instance, we have “new global forms” as in computational emergence, this surely means that “observers have to change their model” (e.g. by the creation of new descriptive terms) as in emergence relative to a model. Thus, the functionally-based ‘relative to a model’ view seems to me to subsume the other two.

Steels (1991, p. 451) emphasises the functionally-based view when he writes: “Emergent functionality means that a function is not achieved directly by a component or a hierarchical system of components, but indirectly by the interaction of more primitive components among themselves *and with the world* [my italics]”. Later (Steels 1994), he refers to ongoing processes which produce results invoking vocabulary not previously involved in the description of the system’s inner components – “new descriptive categories”. The “with the world” qualification is potentially important and often forgotten (although not by Gell-Mann or Franklin). For interaction with the world is what means that chemistry is not physics, and biology is not chemistry. And this is why intelligent systems (‘agents’) might learn things that are useful to us – because they inhabit the same world.

Slightly more recently, Stephan (1998, p. 640) writes: “It is controversial what the criteria are by which emergent properties are to be distinguished from non-emergent properties. Some criteria are very strong ... Other criteria are inflationary in that they count many, if not all, system properties as emergent." He goes on to distinguish three types of emergence:

synchronic – a property is emergent if it is *irreducible* in terms of the arrangement and properties of the system's parts;

diachronic – properties are considered emergent if they cannot, in principle, be predicted "before their first instantiation";

weak – on which the other two are based.

The latter, it is said, "pervades emergentist theorizing mainly in connectionism and theories of self-organization". Emergent properties (in the weak sense) are, among other things, *systemic*. A property is systemic if and only if a system possesses it, but no part of a system possesses it. This seems to be synonymous with the term 'collective' as used, for instance, by Hopfield (1982).

Stephan argues that the problem with 'weak' emergence is that almost all systemic properties come into this category. On the other hand, the problem with his other two, stronger forms is that very little seems to satisfy them! Also, many commentators would doubtless have problems with his use of the term *irreducible* (shades of Horgan's "hoary . . . antireductionist creed"). Stephan mentions qualia as a 'candidate' for synchronic emergence but to cite Crick (1994, p. 9) on qualia: "This is a very thorny issue". It would not be productive to delve deeper into this thorny issue here – the interested reader is referred to Churchland (1989) for more extensive discussion on qualia. Suffice it to say that an issue (qualia) that generates great debate among philosophers is unlikely to offer us much insight about another controversial issue (emergence).

7 Emergence and Reductionism

Throughout this editorial, there has been a tension between emergence and the philosophy of reductionism. The latter holds that it is possible – indeed valuable – to explain a complex phenomenon or system by the interaction among its parts. Anderson (1972, p. 393) says of reductionism "… among the great majority of active scientists I think it is accepted without question" while Crick (1994, pp. 8–9) writes "… it is the main theoretical method that has driven the development of physics, chemistry and molecular biology". So how can this pervasive and obviously effective doctrine be squared with a notion in which it is "impossible to deduce marks of a higher level from those of a lower level" (cf. Pepper) or that "a property is emergent if it is *irreducible* in terms of the arrangement and properties of the system's parts" (cf. Stephan)? To some extent, we have downplayed, if not actually rejected, this formulation of emergence in favour of one in which the litmus test is the appearance of new, qualitatively different phenomena at higher levels of the hierarchy of abstraction/description. Yet the tension remains because the two creeds seem to work in opposite directions. If reductionism is so powerful and useful, who needs emergence?

Part of the resolution, I think, rests with Dennett's (1995, pp. 80–82) warning that reductionism "has no fixed meaning" and his distinction between *bland* reductionism (which "no sane scientist refutes") and *greedy* reductionism, which pushes the notion too far. An expression of how it is possible to push reductionism too far is given by Anderson (1972, p. 313) who outlines what he calls the "constructivist" fallacy: "to reduce everything to simple fundamental laws does not imply the ability to start from those laws and reconstruct the universe". A concrete instance of this inability could be the classical many-bodies problem of dynamical systems whereby all the relevant physical laws and corresponding equations are well known, but exact methods of solution appear beyond us. Perhaps, as suggested by Hofstadter's fictional Achilles character (1979, p. 312), "there is a larger context into which both holistic and reductionistic explanations fit".

8 Concluding Remarks

This Special Issue contains seven papers on diverse aspects of emergence in complex systems illustrating, we hope, the ubiquity and generality of the topic. I would not presume to state what I think each of the individual authors means by emergence. Rather, my goal has been to provide a backdrop against which the various contributions can be more effectively read and appraised. I have tried to argue that emergence is best considered from the perspective of the *understanding* which can stem from viewing complex phenomena and systems at different levels of abstraction – as opposed to the *difficulty* or *impossibility* of so doing.

Acknowledgements

Thanks to Alastair Channon, Stan Franklin, Stevan Harnad, John Hogden, John Holland, Luis Rocha and Adam Prügel-Bennett who kindly commented on an earlier draft of this editorial, and provided valuable pointers to improvement. I have not been able to accommodate and answer all of their criticisms in this final version. The outstanding imperfections remain as testimony to my still limited understanding and to the difficulty and ubiquity of the topic of emergence.

References

Adami, C. (1998). *Introduction to Artificial Life*. New York, NY: Springer-Verlag.

Anderson, P. W. (1972). More is different: Broken symmetry and the nature of the hierarchical structure of science. *Science 177*, 393–396.

Arkin, R. C. (1998). *Behavior-Based Robotics*. Cambridge, MA: MIT Press.

Bak, P. (1996). *How Nature Works: The Science of Self-Organized Criticality*. New York, NY: Copernicus/Springer-Verlag.

Broad, C. D. (1919). Mechanical explanation and its alternatives. *Proceedings of the Aristotelian Society 19*, 86–124.

Broad, C. D. (1925). *The Mind and its Place in Nature*. London, UK: Kegan Paul, Trench, Trubner and Co.

Cariani, P. (1991). Emergence and artificial life. In C. G. Langton, C. Taylor, J. D. Farmer, and S. Rasmussen (Eds.), *Artificial Life II*, pp. 775–797. Redwood City, CA: Addison-Wesley.

Casti, J. L. (1997). *Would-Be Worlds: How Simulation is Changing the Frontiers of Science*. New York, NY: John Wiley.

Churchland, P. M. (1989). *A Neurobiological Perspective: The Nature of Mind and the Structure of Science*. Cambridge, MA: Bradford Books/MIT Press.

Clark, A. (1989). *Microcognition, Philosophy, Cognitive Science and Parallel Distributed Processing*. Cambridge, MA: MIT Press.

Crick, F. (1994). *The Astonishing Hypothesis*. London, UK: Simon and Schuster.

Darwin, C. (1859). *The Origin of Species by Means of Natural Selection*. London, UK: John Murray.

Dawkins, R. (1986). *The Blind Watchmaker*. Harlow, UK: Longman.

de Mey, M. (1982). *The Cognitive Paradigm*. Dordrecht, The Netherlands: D. Reidel.

Denbigh, K. G. (1975). *The Inventive Universe*. London, UK: Hutchinson.

Dennett, D. C. (1991). *Consciousness Explained*. Boston, MA: Little, Brown and Company.

Dennett, D. C. (1995). *Darwin's Dangerous Idea: Evolution and the Meanings of Life*. London, UK: Penguin.

Eigen, M. (1996). *Steps Towards Life: A Perspective on Evolution*. Oxford, UK: Oxford University Press. With R. Winkler-Oswatitsch; translated by P. Woolley.

Forrest, S. (1990). Emergent computation: Self-organization, collective, and cooperative phenomena in natural and artificial computing networks. *Physica D 42,* 1–11. Special Issue on Emergent Computation.

Franklin, S. (1995). *Artificial Minds*. Cambridge, MA: Bradford Books/MIT Press.

Gell-Mann, M. (1994). *The Quark and the Jaguar: Adventures in the Simple and the Complex*. London, UK: Little, Brown and Company.

Georgi, H. M. (1989). Effective quantum field theories. In P. Davies (Ed.), *The New Physics,* pp. 446–457. Cambridge, UK: Cambridge University Press. Pagination refers to 1996 paperback edition.

Greene, B. R. (1999). *The Elegant Universe: Superstrings, Hidden Dimensions and the Quest for the Ultimate Theory*. London, UK: Jonathon Cape.

Guenter, F. H. and M. N. Gjaja (1996). The perceptual magnet effect as an emergent property of neural map formation. *Journal of the Acoustical Society of America 100*, 1111–1121.

Haken, H. (1997). Visions of synergetics. *International Journal of Bifurcation and Chaos 7*(9–10), 1927–1951.

Hawkins, J. A. and M. Gell-Mann (Eds.) (1992). *The Evolution of Human Languages*. Reading, MA: Addison-Wesley.

Hofstadter, D. R. (1979). *Gödel, Escher, Bach: An Eternal Golden Braid*. Hemel Hempstead, UK: Harvester Press. (Pagination refers to 1980 Penguin edition).

Holland, J. H. (1990). Emergent models. In A. Scott (Ed.), *Frontiers of Science*, pp. 107–125. Oxford, UK: Blackwell.

Holland, J. H. (1998). *Emergence: From Chaos to Order*. Reading, MA: Addison-Wesley.

Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective abilities. *Proceedings of the National Academy of Sciences, USA 79*, 2554–2558.

Horgan, J. (1996). *The End of Science: Facing the Limits of Knowledge in the Twilight of the Scientific Age*. Boston, MA: Little, Brown and Company.

Kauffman, S. A. (1984). Emergent properties of random complex automata. *Physica D 10*, 145–156.

Kauffman, S. A. (1987). *Origins of Order: Self-Organization and Selection in Evolution*. Oxford, UK: Oxford University Press.

Kohonen, T. (1977). *Associative Memory: A System Theoretic Approach*. Berlin: Springer-Verlag.

Kohonen, T. (1997). Emergence of invariant-feature detectors in the adaptive-subspace self-organizing map. *Biological Cybernetics 75*(4), 281–291.

Langton, C. G. (1986). Studying artificial life with cellular automata. *Physica D 22*, 120–149.

Levy, S. (1992). *Artificial Life: A Report From the Frontier Where Computers Meet Biology*. New York, NY: Pantheon/Random House.

Li, W. (1992). Random texts exhibit Zipf's-law-like word frequency distribution. *IEEE Transactions on Information Theory IT-38*(6), 1842–1845.

MacWhinney, B. (1998). Models of the emergence of language. *Annual Reviews of Psychology 49*, 199–227.

Mandelbrot, B. B. (1961). On the theory of word frequencies and related Markovian models of discourse. In R. Jacobson (Ed.), *Structures of Language and its Mathematical Aspects*. New York, NY: American Mathematical Society. Pagination unknown.

Matson, W. I. (1965). *The Existence of God*. Ithaca, NY: Cornell University Press.

McFarland, D. and T. Bösser (1993). *Intelligent Behavior in Animals and Robots*. Cambridge, MA: Bradford Books/MIT Press.

Pepper, S. C. (1926). Emergence. *Journal of Philosophy 23*, 241–245.

Popper, K. (1965). *The Logic of Scientific Discovery*. New York, NY: Basic Books.

Popper, K. (1978). Natural selection and the emergence of mind. *Dialectica 32*, 339–355.

Quinlan, P. (1991). *Connectionism and Psychology: A Psychological Perspective on New Connectionist Research*. Hemel Hempstead, UK: Harvester Wheatsheaf.

Rumelhart, D. E. and J. L. McClelland (Eds.) (1986). *Parallel Distributed Processing: Explorations in the Microstructure of Cognition (2 Volumes)*. Cambridge, MA: Bradford Books/MIT Press.

Schroeder, M. (1991). *Fractals, Chaos, Power Laws: Minutes from an Infinite Paradise*. New York, NY: W. H. Freeman.

Steels, L. (1991). Towards a theory of emergent functionality. In J.-A. Meyer and S. W. Wilson (Eds.), *From Animals to Animats: Proceedings of the First International Conference on Simulation of Adaptive Behavior*, Cambridge, MA, pp. 451–461. Bradford Books/MIT Press.

Steels, L. (1994). The artificial life roots of artificial intelligence. *Artificial Life Journal 1*(1), 89–125.

Stein, W. D. and F. J. Varela (Eds.) (1993). *Thinking About Biology: An Invitation to Current Theoretical Biology*. Reading, MA: Addison-Wesley.

Stephan, A. (1998). Varieties of emergence in artificial and natural systems. *Zeitschrift für Naturforschung C-A Journal of Biosciences 53*(7–8), 639–656.

Sussman, H. M., D. Fruchter, J. Hilbert, and J. Sirosh (1998). Linear correlates in the speech signal: The orderly output constraint. *Behavioral and Brain Sciences 21*, 241–299.

Wilson, A. N. (1999). *God's Funeral*. London, UK: John Murray.

Wilson, S. W. (1985). Knowledge growth in an artificial animal. In *Proceedings of 1st International Conference on Genetic Algorithms and their Applications*, Hillsdale, NJ, pp. 16–23. Lawrence Erlbaum.

Zipf, G. K. (1949). *Human Behavior and the Principle of Least Effort*. Cambridge, MA: Addison-Wesley.