Testing for Emergence in Artificial Life

Edmund M. A. Ronald,[1] Moshe Sipper,[2] and Mathieu S. Capcarrère[2]

[1] Centre de Mathématiques Appliquées, Ecole Polytechnique,
91128 Palaiseau Cedex, France. E-mail: eronald@cmapx.polytechnique.fr.
[2] Logic Systems Laboratory, Swiss Federal Institute of Technology, IN-Ecublens,
CH-1015 Lausanne, Switzerland. E-mail: {name.surname}@epfl.ch.

Abstract. The field of artificial life (Alife) is replete with documented instances of *emergence*, though debate still persists as to the meaning of this term. In the absence of a formal definition, researchers in the field would be well served by adopting an emergence certification mark which would garner approval from the Alife community. We propose an *emergence test*, consisting of three criteria—design, observation, and surprise—for conferring the emergence label.

1 Introduction

When a bank's accounting program goes seemingly independent and does its own thing, the programmer scratches his head, sighs, and prepares for doing overtime with the debugger. But when a society of agents does something surprising, Alife researchers may solemnly document this "emergent behavior," and move on to other issues without always seeking to determine the cause of their observations. Indeed, overly facile use of the term emergence has made it controversial. Arkin recently observed that:

> Emergence is often invoked in an almost mystical sense regarding the capabilities of behavior-based systems. Emergent behavior implies a holistic capability where the sum is considerably greater than its parts. It is true that what occurs in a behavior-based system is often a surprise to the system's designer, but does the surprise come because of a shortcoming of the analysis of the constituent behavioral building blocks and their coordination, or because of something else? [1](page 105)

Altogether, it seems the emergence tag has become a great attention grabber, thanks to the striking behaviors demonstrated in artificial-life experiments. We do not think, however, that emergence should be diagnosed *ipso facto* whenever the unexpected intrudes into the visual field of the experimenter; nor should the diagnosis of emergence immediately justify an economy of explanation. Such abuse and overuse of the term will eventually devalue its significance, and bring work centered on emergence into disrepute. Therefore, we contend that in the absence of an acceptable definition, researchers in the field would be well served by adopting an emergence certification mark which would garner approval from the Alife community.

Motivated by this wish to standardize the tagging task, we propose an *emergence test*, namely, criteria by which one can justify conferring the emergence label. Our criteria are motivated by an examination of published work in the field of Alife [10].

The emergence test is presented in the next section, followed in Section 3 by four case studies demonstrating its applicability. Finally, in Section 4, we discuss a number of issues pertaining to our test.

2 An operant definition of emergence for Alife researchers

The difficulties we face in adopting a definition of the concept of emergence are reminiscent of the complications faced by early Artificial Intelligence (AI) researchers in defining intelligence.[1] Nonetheless, where the equally elusive concept of intelligence is concerned, Alan Turing found a way to cut the Gordian knot, by means of an *operant definition* which is useful *within the limited context of man-machine interaction* [14]. Debate concerning the concept of intelligence is unlikely to subside in the foreseeable future, and the same, we believe, holds for emergence. We deem, however, that viewing the world through Turing-colored glasses might improve our vision as regards the concept of emergence—at least where modern-day Alife practice is concerned.

Alife is a constructive endeavor: some researchers aim at evolving patterns in a computer, some seek to elicit social behaviors in real-world robots, others wish to study life-related phenomena in a more controllable setting, while still others are interested in the synthesis of novel life-like systems in chemical, electronic, mechanical, and other artificial media. Alife is an experimental discipline, fundamentally consisting of the observation of run-time behaviors, those complex interactions generated when populations of man-made, artificial creatures are immersed in real or simulated environments. Published work in the field usually relates the conception of a model, its instantiation into real-world or simulated objects, and the observed behavior of these objects in a collection of experiments.

The Turing Test focuses on a human experimenter's incapacity at discerning human from machine when holding what we would now call an Internet chat session. Our emergence test centers on an observer's avowed incapacity (amazement) to reconcile his perception of an experiment in terms of a global world view with his awareness of the atomic nature of the elementary interactions.

[1] On the difficulties in defining emergence, Emmeche, Køppe, and Stjernfelt recently remarked: "One reason for the widespread scepticism against the word [emergence] is a historical load of confusion surrounding the metaphysical aspects of the concept, reflected in the fact that it has been used in a long series of different ways, apparently making it impossible to use it as a clearly defined term..." [5](page 84)

Assume that the scientists attendant upon an Alife experiment are just two: a system designer and a system observer (both of whom can in fact be one and the same), and that the following three conditions hold:

(i) Design. The system has been constructed by the designer, by describing *local* elementary interactions between components (e.g, artificial creatures and elements of the environment) in a language $\mathcal{L}_1$.

(ii) Observation. The observer is *fully aware* of the design, but describes *global* behaviors and properties of the running system, over a period of time, using a language $\mathcal{L}_2$.

(iii) Surprise. The language of design $\mathcal{L}_1$ and the language of observation $\mathcal{L}_2$ are distinct, and the causal link between the elementary interactions programmed in $\mathcal{L}_1$ and the behaviors observed in $\mathcal{L}_2$ is *non-obvious* to the observer—who therefore experiences surprise. In other words, there is a cognitive dissonance between the observer's mental image of the system's design stated in $\mathcal{L}_1$ and his contemporaneous observation of the system's behavior stated in $\mathcal{L}_2$.

The above three clauses relating design, observation, and surprise describe our conditions for diagnosing emergence, i.e., for accepting that a system is displaying emergent behavior.

When assessing the surprise clause of our test one should bear in mind that as human beings we are quite easily surprised (as any novice magician will attest). The question reposes rather on how *evanescent* the surprise effect is, i.e., how easy (or strenuous) it is for the observer to bridge the $\mathcal{L}_1$–$\mathcal{L}_2$ gap, thus reconciling his global view of the system with his awareness of the underlying elementary interactions. One can draw an analogy with the concept of intelligence and the Turing test: the chatty terminal might at first appear to be carrying on like an intelligent interlocutor, only to lose its "intelligence certificate" once the tester has pondered upon the true nature of the ongoing conversation.

Some of the above points deserve further elaboration, or indeed invite debate. Before treating these issues in Section 4, we wish to demonstrate the application of our test to four cases.

3 Administering the emergence test: Four case studies

In this section we administer the emergence test to four examples, thus demonstrating its application (additional examples are given in [10]). Each example ends with a "test score," constituting our own assertion as observers of whether we are indeed surprised, that is, of whether emergent behavior is indeed displayed—or not.

1. Emergence of a nest structure in a simulated wasp colony, from the interactions taking place between individual wasps [13].

(i) The design language $\mathcal{L}_1$ is that of local wasp interactions, including movement on a three-dimensional cubic lattice and placement of bricks. A wasp's decision is based upon a local configuration of bricks, which lie in its "visual" field. Actions to be taken are prewired under the form of a lookup table with as many entries as there are stimulating configurations.
(ii) The observation language $\mathcal{L}_2$ is that of large-scale geometry, as employed to describe nest architectures.
(iii) While fully aware of the underlying wasp interaction rules, the observer nonetheless marvels at the sophistication of the constructions and their striking similarity to naturally occurring nests.
Diagnosis: emergent behavior is displayed by the nest-building wasps.

2. Emergence of a "highway" created by the artificial Langton ant, from simple movement rules [12].

(i) The design language $\mathcal{L}_1$ is that of single moves of a simple, myopic ant. The ant starts out on the central cell of a two-dimensional, rectangular lattice, heading in some selected direction. It moves one cell in that direction and looks at the color of the cell it lands on—black or white. If it lands on a black cell, it paints it white and turns 90 degrees to the left; if it lands on a white cell, it paints it black and turns 90 degrees to the right. These simple rules are iterated indefinitely.
(ii) The observation language $\mathcal{L}_2$ is that of global behavioral patterns, extended over time and space (i.e., tens of thousands of single ant moves, spanning thousands of cells). Specifically, the ant was observed to construct a "highway," i.e., a repeating pattern of fixed width that extends indefinitely in a specific direction (Figure 1a).
(iii) While fully aware of the very simple ant rules, the observer is nonetheless surprised by the appearance of a highway.
Diagnosis: emergent behavior is displayed by the highway-constructing ant.

3. Emergence of flocking behavior in simulated birds, from a set of three simple steering behaviors [9].

(i) The design language $\mathcal{L}_1$ is that of local bird interactions, the three rules being: *separation:* steer to avoid crowding local flockmates; *alignment:* steer toward the average heading of local flockmates; *cohesion:* steer to move toward the average position of local flockmates. A bird's decision is based upon its nearby neighbors, i.e., those that are in its "visual" field.
(ii) The observation language $\mathcal{L}_2$ is that of flocking behaviors, such as the flock's parting smoothly when faced with an obstacle, and "flowing" around it—to then reunite again (Figure 1b).
(iii) While fully aware of the underlying bird interaction rules, the observer nonetheless marvels at the lifelike flocking behaviors.

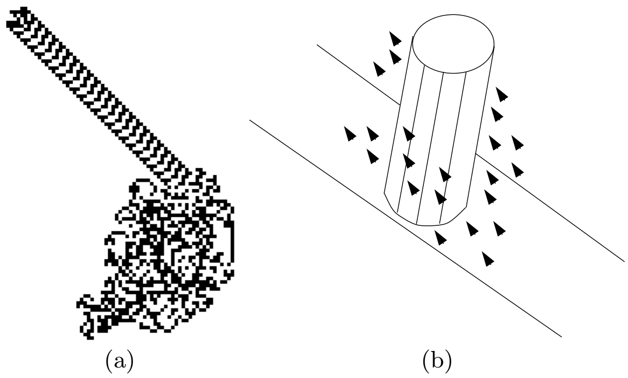

Fig. 1. Examples of emergence. (a) The trail created by the highway-constructing, Langton ant. (b) A flock of simulated birds parts smoothly when faced with an obstacle, and "flows" around it—to then reunite again (after Reynolds [9]).

Diagnosis: the flocking behavior exhibited by the artificial birds was considered a clear case of emergence when it was first reported upon in 1987. However, one could now maintain that it no longer passes the emergence test, since wide-spread use of this technique in computer graphics has obviated the element of surprise. This example demonstrates that the diagnosis of emergence is contingent upon the sophistication of the observer.

4. Emergence of wall-following behavior in an autonomous, mobile robot, from the simultaneous operation of two simple behavior systems: obstacle avoidance and wall seeking [11].

(i) The design language $\mathcal{L}_1$ is that of simple robot behaviors, including—in this case—obstacle avoidance and wall seeking.

(ii) The observation language $\mathcal{L}_2$ is that of more elaborate robot behaviors, consisting—in this case—of wall following.

(iii) Steels wrote that "Wall following is emergent in this case because the category 'equidistance to the (left/right) wall' is not explicitly sensed by the robot or causally used in one of the controlling behavior systems." [11](page 92)

Diagnosis: Steels diagnosed emergence in this case as it accords with his own definition, namely, that a behavior is emergent if it necessitates the use of new descriptive categories that are not needed to describe the behavior of the constituent components [11]. While thus alluding to the language dichotomy rendered explicit by our definition (i.e., the existence of two distinct languages—that of design and that of observation), we maintain that the surprise element is missing: the wall-following behavior can be quite readily deduced by an observer aware of the two underlying simpler behaviors. We thus conclude that emergent behavior is *not* displayed by the wall-following robot.

4 Discussion

We now discuss the various components of our test.

The operant nature of the test. In this we have drawn our inspiration from Turing, who—concerning intelligence—opted for an operant, informal, "social" definition, deliberately eschewing rigor. Turing's definition still serves the AI community well—almost half a century after its publication [14].

Emergence as a property of artificial systems. In his book *Emergence: From Chaos to Order*, Holland wrote that "*Emergence occurs in systems that are generated*" [7](page 225). Reviewing Holland's book, Mallot also opined that "In this context [the construction of artificial systems], the problem of emergence may actually be a genuine one." [8] These views accord with our own view, namely, that the diagnosis of emergence should be considered (and hence our test applied) within domains such as Alife, which are inherently constructive endeavors. This view naturally gives rise to clause (i) of our test, i.e., the existence of a designer—and of a design language.

The existence of an observer. Artificial systems are constructed to be beheld—one does not usually build one's system, to then walk away nonchalantly without ever looking back. Hence, there exists an observer *ipso facto* (who need not necessarily be the constructor himself), a fundamental aspect which has not escaped researchers in the field. In a paper discussing emergence and artificial life, Cariani wrote that "The interesting emergent events that involve artificial life simulations reside not in the simulations themselves, but in the ways that they change the way we think and interact with the world." [2](page 790) He goes on to say that "computer simulations are catalysts for emergent processes in our own minds..." [2](page 790)

Another author, Emmeche, in an introductory monograph on artificial life, examines the case for emergence "in the eye of the beholder." [4](page 145) Also, Crutchfield, in an article devoted to the subject of emergence, asks: "But for whom has the emergence occurred? More particularly, to whom are the emergent features 'new'?... The newness in both cases is in the eye of an observer..." [3](page 517)

Holland brings up the issue of the observer circuitously, when writing that "*The whole is more than the sum of the parts in these generated systems*... Said another way, there are regularities in system behavior that are not revealed by direct inspection of the laws satisfied by the components." [7](page 225) One may ask *direct inspection by whom?* Why, by the observer of course![2] Clearly, the existence of an observer is a *sine qua non* for the issue of emergence to arise at all.

[2] Holland also cites a passage from Gell-Mann's book *The Quark and the Jaguar* [6], which also brings up indirectly the role of the observer: "In an astonishing variety of contexts, apparently complex structures or behaviors emerge from systems

Surprise. By bringing the observer's *emotion* of surprise into play, our emergence test widens the focal beam of discussion, now shining both on the *system's behavior* as well as on the experimenter and her *internalized expectations*. This relates to Cariani's nutshell description of emergence relative to a model as "the deviation of the behavior of a physical system from an observer's model of it." [2](page 779) An author subscribing to said deviation-from-model view would wish to document her *a priori* expectations before diagnosing emergence and abandoning attempts at explanation. Our emergence test might then be reformulated as *Design (Expectations), Observation, Surprise.*

To summarize, the three clauses of our emergent test are grounded in previous work: the design clause expresses our wish to restrict the test to artificially constructed systems; the observation clause reflects the necessity of there being an observer for emergence to arise at all; and the surprise clause embodies both the deliberation and the emotion implied by human judgments of value.

Acknowledgments

We thank Bernard Besnier, Dario Floreano, Daniel Mange, Eduardo Sanchez, Marco Tomassini, and the anonymous reviewers for their helpful remarks and suggestions.

References

1. R. C. Arkin. *Behavior-Based Robotics.* The MIT Press, Cambridge, MA, 1998.
2. P. Cariani. Emergence and artificial life. In C. G. Langton, C. Taylor, J. D. Farmer, and S. Rasmussen, editors, *Artificial Life II*, volume X of *SFI Studies in the Sciences of Complexity*, pages 775–797, Redwood City, CA, 1992. Addison-Wesley.
3. J. P. Crutchfield. Is anything ever new? Considering emergence. In G. Cowan, D. Pines, and D. Melzner, editors, *Complexity: Metaphors, Models, and Reality*, pages 515–537. Addison-Wesley, Reading, MA, 1994.
4. C. Emmeche. *The Garden in the Machine: The Emerging Science of Artificial Life.* Princeton University Press, Princeton, N.J., 1994.
5. C. Emmeche, S. Køppe, and F. Stjernfelt. Explaining emergence: Towards an ontology of levels. *Journal for General Philosophy of Science*, 28:83–119, 1997.
6. M. Gell-Mann. *The Quark and the Jaguar: Adventures in the Simple and the Complex.* Freeman, New York, 1994.
7. J. H. Holland. *Emergence: From Chaos to Order.* Addison-Wesley, Reading, Massachusetts, 1998.
8. H. Mallot. Life is like a game of chess: review of "Emergence: From Chaos to Order" by John Holland. *Nature*, 395:342, September 1998.

---

characterized by simple rules." [7](page 238) Gell-Mann's use of the qualifier "apparently" suggests that the quality in question necessitates a judgment call—that is, an observer.

9. C. W. Reynolds. Flocks, herds, and schools: A distributed behavioral model. *Computer Graphics*, 21(4):25–34, July 1987.
10. E. M. A. Ronald, M. Sipper, and M. S. Capcarrère. Design, observation, surprise! A test of emergence. *Artificial Life*, 1999. (to appear).
11. L. Steels. The artificial life roots of artificial intelligence. *Artificial Life*, 1(1/2):75–110, 1994. The MIT Press, Cambridge, MA.
12. I. Stewart. The ultimate in anty-particles. *Scientific American*, 271(1):104–107, July 1994.
13. G. Theraulaz and E. Bonabeau. Coordination in distributed building. *Science*, 269:686–688, August 1995.
14. A. M. Turing. Computing machinery and intelligence. *Mind*, 59:433–460, 1950.