

# 1

## A Turing test for emergence

Fabio Boschetti and Randall Gray

### 1.1 Introduction

Dealing with *complex* systems present a particular challenge to many traditional engineering approaches. The pertinent assumption inherent in these approaches is that component parts of a system can be neatly partitioned and that their interactions have limited, predictable effects. This assumption is not always tenable, and has an impact both on degree of overall control attainable and on the robustness of the resulting systems. A traffic controller does not need to give exact instructions to each vehicle on the road, and a Treasurer does not need to control each single business in a country; rather they both provide general guidelines which aim at a desired global outcome; they both rely on the local organisation inherent in road traffic and business interactions to account for local details. Similarly we would like a designer to specify broad guidelines in order for a complex system to act according to a general requirement. Since the inherent organisation we wish to exploit is often a dynamical and stable configuration, a system designed to capitalise on this may also display a robustness and adaptivity which is currently beyond our engineering abilities.

In the parlance of complex system science, the global outcomes arising from broad guidelines on a system's components, including robustness and adaptivity, are often defined as emergent features. Because design inevitably requires a trial and error process, it is natural to expect that our community will need to develop methods to:

- detect emergent features when they arise;
- categorise them in order to understand what classes of processes arise as a result of different initial conditions;
- experiment with various configurations in order to optimise the emergent processes.

Experimentation is something which is often carried out via computer simulation; however, computers are a perfect example of a 'traditional' engineering apparatus, and consequently display the very same features (lack of robustness and adaptivity,

and a requirement for detailed instructions) which we are trying to circumvent. In this chapter we explore this apparent paradox and ask what kind of emergent features can be generated (and thus modelled) in a computational framework. We will show that this question directly relates to the other two items listed above, that is to the experimental detection of emergence and its classification.

## 1.2 Background

The concept of emergence evolved to capture our intuition that when a large number of entities interact, the resulting system can display features and behaviours which are not displayed by the individual constituents. The human body possesses behaviours and functions which are not expressed by our individual cells; metals show properties not displayed by individual atoms; societies undergo dynamics which transcend individuals. Basic examples can also be modelled very easily on a computer; famously, Conway's Game of Life (18) shows how very simple local rules generate features whose dynamics is not explicitly coded in the algorithm. Examples are so ubiquitous in Nature that some scientists suspect that all structures we see 'emerge' from underlying simpler levels (34).

Nevertheless, emergence raises considerable intellectual and scientific challenge. Despite a vast literature, going back several decades (13), no agreement can be found on a definition, nor on a framework for its study, nor on whether emergence is a 'real' natural phenomenon or merely a by-product of our perception, or a convenient way to make sense of processes otherwise too hard to comprehend (15; 35).

Why should a process which appears so obvious and easy to model prove so hard to define and conceptualise? The fundamental reason is that in an emergent process it is very hard to discriminate 'who does what'. When I decide to listen to music, is it my 'emergent' self which takes the decision or my cells? My body depends on cellular activity for its functioning, so cells must be the controlling entities. However, no cell decides to listen to music since listening to music is not something cells 'do'. This leads straight into old and unsolved philosophical problems of causality, determinism and freewill.

Crucially, this is also a technological problem. Today, probably for the first time in history, technological developments in many applications depend on the understanding of emergent phenomena. Advances in Information Technology, Epidemiology, Ecosystem Management, Health Science, just to name a few, depend on approaches which go beyond traditional reductionism and address the understanding of how emergent properties arise, what they 'do' and how they can be controlled.

It thus seems natural that when we ask whether emergence is 'real' or merely lies 'in the eyes of the observer', or whether emergence is a distinct process of its own or encompasses different processes among which we are not yet able to discriminate, the answer needs to account for what these processes 'do'. In other words, we need to account for causal relationships and causal power. It may appear that we are trying to address a slippery problem (emergence) via one which is even more slippery (causality). This does not need to be so if we constrain what we mean by causality

and we adopt an ‘operative’ definition. Following Pattee (1997) and Pearl (2000) we associate causal power with control: a process has causal power if, by acting upon it, we can change the effects it produces. Pattee (1997) describes this very simply: *“Useful causation requires control. .... Clearly it is valuable to know that malaria is not a disease produced by “bad air” but results from Plasmodium parasites that are transmitted by Anopheles mosquitoes. What more do we gain in these examples by saying that malaria is caused by a parasite ..? I believe the common, everyday meaning of the concept of causation is entirely pragmatic. In other words, we use the word cause for events that might be controllable. In the philosophical literature controllable is the equivalent of the idea of power. In other words, the value of the concept of causation lies in its identification of where our power and control can be effective. For example, while it is true that bacteria and mosquitos follow the laws of physics, we do not usually say that malaria is caused by the laws of physics (the universal cause). That is because we can hope to control bacteria and mosquitos, but not the laws of physics.”*

Building from this observation and from the work of Shalizi (2001), Crutchfield (1994b; 1994a), Rabinowitz (2005), among others (6; 5; 10; 19; 1; 21; 45; 39; 4; 16; 17), we propose to discriminate between three types of emergence, depending on increasing level of ‘causal’ power: pattern formation, intrinsic emergence and causal emergence.

Despite the philosophical halo of the above discussion, our aim is utterly practical. In a scientific culture in which understanding is increasingly synonymous with computer modelling, we ask what forms of emergence can be studied by simulation and what we can gain from doing so. We will see that computational and ‘causal’ barriers are strongly related. This may lead to new insights into the limitations and future of the computer modelling of complex processes.

### 1.3 Formal Logic and Computation

There exists an equivalence between the workings of formal grammars, logical systems and computation (11; 42; 29; 32). All these start from some fundamental set of strings (starting symbols, axioms or input data), a set of rewriting rules (production rules, rules of inference, computer instructions), and they generate outputs (strings, theorems and computational results) which are obtained by transforming the a priori set via the rewriting rules.

In a formal system, true statements are almost always either theorems or tautologies (Kurt Gödel demonstrated that there are true statements which are not accessible from the axioms and rules of logic. These true statements are not theorems since they are not derived from the axioms). This is so because, given a set of axioms and inference rules, these statements are necessarily true (they are true for all possible scenarios and cannot be otherwise). Given a set of axioms and inference rules, these statements necessarily follow and are true for all possible scenarios and cannot be otherwise. Consequently, these statements do not provide any information about the real world (any information such a string may seem convey is a result of correspondences we see (or think we see) between the real world and the fundamental system, and is wholly dependent on our

perception of these correspondences). An example clarifies the concept: the statement ‘*it’s raining*’ may be true today and may or may not be true tomorrow; it depends on its agreement with the vagaries of the real world. Assessing whether the statement is true or not provides information about the real world. Pythagoras’ theorem, in contrast, is true independently of Nature’s vagaries, it must be true and always will. The fact that Pythagoras’ theorem is useful to us and matches our perception of reality is due to the clever choices of the basic axioms of geometry. It is because of the appropriateness of axioms collected in Euclid’s work that the properties of triangles match our perception of reality.

Given Euclid’s axioms and our rules of mathematical reasoning, Pythagoras’ theorem is an inevitable consequence. It helps us to understand Nature better by simplifying geometrical considerations, by putting place holders in our geometrical thinking so we do not always need refer back to the axioms, and it helps us to communicate this understanding, but it does not provide any information which is not already implicit in the our axioms and rewriting rules. Theorems are transformations of information, not new information. In some sense, all the theorems of Euclidean geometry could be compressed, with no loss of information, into the basic axioms and inference rules (this is formally proved in (11) and is the base for Kolmogorov/Chaitin’s complexity measure). It could reasonably be argued, though, that any decompressor which could reproduce the theorems of Euclidean geometry through its decompression would need to be at least as complex as a mathematician and, like a mathematician, would stand a reasonable chance of passing the Turing Test (which we discuss later).

The PCs on our desks are equivalent to a finite tape Turing Machine (TM), an abstract and general computational device commonly employed in theoretical computer science. Because the execution of a TM is equivalent to the application of production rules in a formal grammar, and to proof in a formal system (42), it follows that the result of running a TM is equivalent to a theorem or a valid string: the results are independent of reality.

It thus also follows that the outputs of any of our computer models are similarly dictated by their initial state and the rewriting rules embodied by the program (technically, this is correct provided the PC does not allow for interaction with the outside world, see (45; 24)); a computer model transforms the information contained in its input via its coded algorithm, but does not generate information. Clearly, a model’s output helps our finite mental capability to see consequences of what we coded (which at times we cannot envisage), but its truth status and relevance to the real world is limited to the truth and relevance of the user code and the input fed to the computation. No actual information about the real world is produced by a simulation. Information is generated solely by the writing of the code and the choice of the input. In this way, our choices about how we model a system are much like Euclid’s choices and the comparison of the results of our simulations to what we observe in Nature tells us about the appropriateness of the rules we implement and the input we choose.

## 1.4 Algorithms and Physical Laws

In our perception of reality, causality manifests itself as physical laws (conversely, a physical law can represent both causal relations and mere correlations, from which it arises the philosophical dilemma behind causality. For the purpose of our discussion it is important to stress that causality can be represented only as a physical law, such as "for every action there is a corresponding equal and opposite reaction"). Our computational representation of physical laws involves algorithms which are essentially transformation rules (sequences of instructions). Since we have seen that transformation rules of this sort are constrained to produce results which are members of a set which is totally determined by these rules and the initial conditions, we need to conclude that the running of algorithms which represent physical laws can only produce similarly deterministic results. Any physical law (rule) which an algorithm can generate must already be implicit in the physical laws (rules) represented in the coded algorithm. No new physical law (or representation of it) can be generated by modelling.

When faced with the question "can genuinely novel causal laws emerge from lower level causal laws?" or "can causal laws which transcend the causal power of their constituents exist in Nature?" we can envisage two possible answers:

1. either emergent, genuinely novel, causal laws can not exist and are only apparent and perceived as such because of the limitations in the representation we use;
2. or emergent causal laws must arise via natural processes which are non-algorithmic, fundamentally different from the workings of a formal logic system and consequently not computable in classical sense.

## 1.5 Three Levels of Emergence

In this section we examine three levels of emergence, often discussed in the literature. Our analysis focuses on the relative causal power of the emergence features they can generate.

### 1.5.1 Pattern Formation and Detection

Pattern formation captures the most intuitive view of emergence. The interaction of low level simple entities, leading to symmetry breaking, generates a coordinate behaviour; this is expressed by patterns which are novel and identifiable as such by an external observer. "The patterns do not appear to have specific meaning within the system, but obtain a special meaning to the observer once (and if) he/she is able to detect them. When this happens, the patterns become part of the tool-box the observer can employ to describe and study the process" (15). Examples include the Game of Life discussed above, spiral waves in oscillating chemical reactions, convective cells in fluid flow and fractal structures in fractured media.

For the purpose of our discussion, pattern formation does not, in itself, imply causal power. Let's consider the Game of Life and the emergent gliders. Detecting their presence is useful for an observer to comprehend the effect of the local rules, to highlight

the potentially universal computational capability of the system and possibly to devise a language able to compress their description (37; 35). The question relevant to our discussion is whether the gliders can ‘do’ something or are simply ‘passive’ expressions of internal dynamics; can we exert any causal control on the gliders? What should we do to affect the behaviour of the gliders?

The obvious answer is that we could manipulate the Cellular Automata (CA) local rules. This however acts at the lower level (the CA cells) not at the level of the gliders. By doing so, gliders are still merely a representation of our manipulation of the local rules. Can we act on the gliders themselves? We believe that this could happen only via re-writing the CA code, that is via an external intervention and a complete redesign of the system. We will discuss this more extensively in Section 1.6.3. For now we suggest that pattern formation, per se, does not imply causal power.

### 1.5.2 Intrinsic Emergence

Intrinsic emergence refers to features which are important within the system because they confer additional functionality on the system itself. These emergent features may support global coordination-computation-behaviour like the motion of a flock of birds or stock market pricing (15). Examples with immediate relevance to modelling are Minority Game models (2; 3): agents must take local decisions on actions which result in an economic outcome but they are not able to communicate, so they have no information about other agents’ behaviour. If they identify an emergent feature, providing information about the global dynamics of the population’s economy, then they can use this measure to decide what actions to take (7). This feature now acts as an avenue for global information processing and provides to the system the possibility for coordinated behaviour. Clearly, the agents’ behaviour influences the global measure, but now the global measure affects the behaviour of the agents by determining their future actions. Self-referentiality becomes a fundamental ingredient for complex dynamics and intrinsic emergence.

Discriminating whether intrinsic emergence implies causal control is more challenging and is surely not as clear cut as for pattern formation. In a real world we could externally affect the stock market (with some sort of governmental intervention, for example) thereby changing indirectly the dynamics of the agents, who would respond to the sudden external change by altering their future behaviour. This intervention is not possible in the case of pattern formation described above, since we cannot intervene on a convective cell (for example) without acting directly on the molecules’ motion. In the case of a simulation, we could affect the future behaviour of the model by changing the values of the emergent feature (market), without having to re-program the code. However, this is not fully satisfactory since, in a classic Turing Machine, no interaction with the computation is allowed and, consequently, the distinction between algorithm and input data is blurred.

### 1.5.3 Causal Emergence

The relation between emergence and causality has been studied under the term ‘downward causation’ or ‘strong emergence’ (19; 6; 20). Roughly, ‘a feature is emergent

if it has some sort of causal power on lower level entities'. Like all topics involving causality, this is a subject open to considerable controversy (see (35)). Here we refer to it as 'causal emergence' to highlight the fact that we employ the weaker definition of causality involving control and consequently our conclusions do not necessarily generalise to the global problem of downward causation. Another suitable name could be emergence of control.

With causal emergence we define the arising of structures on which we can exert direct control without manipulating, nor concerning ourselves with, the lower level constituents. As an example, we assume again that the ultimate cause of human behaviour lies in the biochemical process arising at a molecular and cellular level. Suppose I want to ask my friend Jim to play some music for me. I can do so by addressing him directly, for example by speaking or writing a message. Once a message is received, my friend will employ his biological machinery to accept the invitation, but I do not need to concern myself with it. I do not need to re-program complicated instructions into Jim's cellular sub-stratum. For all practical concerns, my friend acts as an entity with emergent causal power.

## 1.6 Modelling Causal Emergence

In the previous section we proposed to subdivide emergence into three classes depending on the causal power of the features they can generate, ranging from pattern formation, which generates features with no causal power, to intrinsic emergence, displaying limited, indirect causal power to causal emergence, empowered with full causal power.

In Section 1.4 we claimed that the generation of causal power cannot be modelled, since an algorithm cannot produce novel rules. If this statement is correct, then we deduce that while we can model pattern formation, and we may or may not be able to model intrinsic emergence (depending on whether we allow for interaction with data rather than instructions), we should not be able to model causal emergence. If true, this is quite a bad piece of news, since a large component of research on Complex Systems is today carried out via computer modelling and emergence is considered to be a crucial ingredient of complex systems.

This is a potentially important claim. For a claim to be meaningful, however, it needs to be relevant and falsifiable. In this section we discuss why this claim is relevant to current scientific investigation by addressing applications to biological and ecological modelling, Artificial Intelligence, Artificial Life and the mining of large scientific data sets. This will lead us along the difficult path of falsification via a variant of the Turing test, applicable to emergence processes. A full discussion of the falsification of this claim requires addressing much subtler issues of the philosophy of science and meta-mathematics which are beyond the scope of this paper, but which we touch upon briefly in the final Discussion.

### 1.6.1 Is This Relevant?

Pattern formation is usually considered the most trivial form of emergence. Nevertheless, its relevance to our scientific enquiry is beyond doubt. An inspiring exposition

on the relevance of intrinsic emergence to the understanding of Nature can be found in (15), to which we refer the reader. Here we discuss the possible relevance of the concept of causal emergence. As mentioned above we distinguish causal emergence from Downward Causation in this work. Ample discussion of the related concept of Downward Causation can be found in (1).

Following our discussion in Section 1.4, the relevance of causal emergence depends essentially on whether we believe uncomputability can be found in Nature. On this topic, the scientific community is broadly divided into two groups. The first group, by far the largest, believes that uncomputability exists only in the abstract world of formal logic and pure mathematics, not in the natural world. A common justification of this view is that no example of uncomputability has so far been detected in Nature nor is there a specific need to include it in our descriptive tools. A smaller community believes that uncomputability can be found in Nature. Among these we can cite Penrose's famous claims about the super computability of the human brain (33; 32; 39; 22). According to Penrose we can easily envisage real implementations of the abstract concept of a Turing Machine, so there is no reason to believe that uncomputability cannot be generated in Nature. For a further discussion on this topic see also (12; 9). A natural observation for supporters of the latter view is that, if all tools enabling us to study Nature are based on computation (i.e. algorithms), then it follows that no uncomputable process can be detected. This observation leads to a possible third view of the problem, according to which Nature may or may not include uncomputable processes, but we will never be able to detect or access them because of the inherent limitation in the language we use to interpret it. We will come back to this possibility later.

### 1.6.2 Biological/Ecological Modelling

The idea underlying any computer modelling is to create a virtual laboratory where a researcher can perform experiments and scenario testing which would be impossible, impractical or too costly to carry out in the real world. The relevance of these experiments depends on how well the virtual laboratory resembles the real world. Nineteenth century physics has taught us that perfect accuracy is beyond our reach (Heisenberg Uncertainty Principle, for example), and this teaching is today well accepted. Nineteenth century mathematics has taught more fundamental concerns (Gödel Theorem, for example), which, curiously, are more easily dismissed.

A considerable experience in engineering, physics and chemistry has shown immense practical benefits and, when the general limitations are carefully accounted for, has proved how useful computer modelling can be. When porting the approach to biological and ecological modelling it becomes tempting to employ the same method for studying processes like evolution, adaptation, and creation of novelty and diversity. However, we believe that these processes involve the same causal emergence we discussed above and it thus becomes necessary to ask whether the virtual laboratory has a similar functional relation to the real world to that enjoyed by physical systems. The same question can be framed as follows: to what extent can a biological agent be modelled within the same framework used to model non-living objects and processes?



To give a practical example of where the challenge may reside, it is useful to remember that a crucial concept in biological and ecological studies is the existence of multiple levels of organisation (cells → organs → individuals → communities → species → ecologies, etc.).

According to our current understanding, these structures self-organise (they do not follow explicit external direction templates (21)) and are linked by two-way (upward and downward) interactions. Many real world problems (ecological and renewable resource management for example) depend on our understanding (i.e., modelling) of this supposedly spontaneous generation of organisation and two-way interactions. Obviously, the more complex the questions we ask, the more complicated the models we need to develop, and the more levels of organisation we may need to include in the model. For example, in a fishery management problem we may want to study how individual fish organise themselves in schools or how individual fishers organise a fishing fleet. This represents one level of organisation. If the specifics (or the scale) of the problem requires so, we may also need to model how schools of different fish interact, or how a school of fish interacts with a fishing fleet; this represents a second level of organisation. In our model we can design a set of rules (a module) which controls the behaviour of the individual fish and vessels and a set of modules for the behaviour of fish schools and the fleet. However, if our purpose is to understand how these multiple levels arise and interact, then we would like the dynamics of the different levels to be shared or at least related. In principle we may want to code a single module (of the lower level) and see how higher levels of organisation arise as a result; after all, this is what we conjecture happens in Nature.

Here, in our opinion, a fundamental discontinuity is revealed. In order to model this nesting of organisation, the schools and fleets need to be more than mere patterns arising from the lower level; they need to be able to ‘do’ something. In particular, they need to be able to causally interact with other entities. Following our discussion in Section 1.6, this equates to asking whether we can exert control on the system without needing to ‘refer back’ or manipulate the rules controlling the individual fish and vessels. In other words, as we are able to ask our friend to play some music (without needing to concern ourselves about his ‘lower level’, local, biochemical rules) by merely interacting with him at a higher level similarly, we would like to be able to exert control on a school or fleet without having to concern ourselves with the lower level rules governing them. If we cannot do that, then we must conclude that the school/fleet system is merely a pattern, which we can identify and analyse, but which does not have any causal power. Thus the question is, can we exert such control?

### 1.6.3 Artificial Life and Artificial Intelligence - A Turing Test for Emergence

We believe the answer to the previous question is negative. We also believe this is merely a conjecture and that it cannot be proved. We also believe this issue is highly debatable since it mostly depends on potentially different interpretations of causal control, as we discuss in this section.

In Section 1.5.1 we expressed our opinion that the gliders in the Game of Life are mere patterns with no causal power and we asked ourselves whether we can interact

with the gliders without re-coding their local rules. Answering this is not trivial, mostly because it depends on how much ‘purpose’ is placed in the original local rules.

We explain what we mean by ‘purpose’ with an example. Let’s take a flocking model (Reynolds). Birds fly in flocks by ensuring they maintain certain constraints on the position between each other. Suppose we now place an obstacle on the route of the flock. The flock will circumvent the obstacle. It thus appears that we were able to exert control on the behaviour of the flock; the flock appears to have causal power. However, we ask ourselves whether the flock has actually done anything which was not explicitly coded in the lower level rules. After all, all the flock did was to maintain flight by following a lead bird which avoided the obstacle. Is there any emergent behaviour in this? Is there any causal power which was not purposely coded.

Can we causally control the flock in any different way? Reasonable arguments can be given in both the affirmative and negative in answering this question. Interestingly, this is not particularly important. Let’s consider once again my friend Jim playing some music. It is beyond our current understanding to discriminate to what extent our verbal instruction interacts with his underlying biochemistry. Similarly, it is beyond our current knowledge to grasp how our higher level invitation (spoken request) is processed at his lower (biochemical) level for the task to be carried out. For our discussion, what matters is only our perception of causal control on Jim’s behaviour. By analogy, we are led to conclude that in the Game of Life or flocking example, what matters is the perception to which we believe we can exert causal control over the higher level emergent features. Does it look as if those features possess causal control? Does it look like they do more than the limited number of behaviours purposely encoded in the local rules? Do system entities behave as if they were autonomously interacting with external processes and respond accordingly?

These new questions have the flavour of a ‘Turing test for emergence’. Famously, the Turing test (for related variations, see the series of paper contained in (36; 40)) was designed to circumvent the difficult question of defining what intelligence is and to detect when a computer can be said to have achieved it (one of the original purposes of Artificial Intelligence at its very conception). Turing suggested as testing whether a human (an intelligent agent) was able to discriminate blindly between another human (another intelligent agent) and a computer. Should he/she not be able to, then we should conclude that the computer and the human act as intelligently as each other, and therefore they are both similarly intelligent. Following an analogous reasoning, we conceive an ‘emergent’ version of the test and we ask whether a process empowered with autonomous causal emergent properties (a human) can discriminate between another causal emergent process and a computer program. Should he/she not be able to do so, then we should conclude that the computer displays causal emergence.

We are not actually suggesting that the test be carried out in earnest. Rather, we would like to refer to and build upon the vast body of work (both conceptual and practical) carried out on the Turing test over several decades and extend some of the conclusions which may be relevant to the study of emergent processes and computer modelling. In this regard, notice that intelligence is itself often considered an emergent feature of the processing occurring in a nervous system. If we accept this view, then the ‘Turing test for emergence’ can be seen as a generalisation of the traditional Turing

test. Consequently extending the discussion of the traditional Turing test to emergence becomes more than merely exploiting an imaginary analogy.

The traditional Turing test has been subjected to considerable theoretical discussion and criticism. Nevertheless, practical implementations of the Turing test are carried out annually in the form of the Loebner prize (Rosenzweig). So far, it is widely accepted that improvement in the test performance over the years has not been particularly significant and ‘passing the test’ does not seem to be a likely short term outcome. The entire artificial intelligence community has, therefore, revisited its own role, scope and measure of success. Far from being a proof, this observation does somehow reinforce our conjecture that modelling causal emergence via computer simulation should, at the very least, not be taken for granted.

On a more positive side, this suggests a reason for the Complex System Sciences (CSS) community building more closely on the extensive experience accumulated along the difficult path followed by artificial intelligence. After a few decades of pessimism, a new breeze of optimism can be felt in both the artificial life and artificial intelligence community. This renewed confidence is not based on the infrastructure of logical programming or the complications of expert systems (as in the past), nor on hopes of super computability brought to us by quantum computing. Rather it depends on more down-to-earth, often biologically inspired, approaches. As an example, in a series of papers (45; 43; 26; 25; 24; 23; 44), van Leeuwen and Wiedermann show formally that agents interacting with their environment have computational capabilities which supersede classic computation. There are a number of reasons why interacting agents can achieve these acrobatics: they run indefinitely (as long as the agent is alive), they continuously receive input from a (potentially infinite) environment and from other agents (unlike a classic machine for which the input is determined and fixed at the beginning of the calculations), they can use the local environment to store and retrieve data and they can adapt to the environment. In particular, the agents’ adaptation to their environment means that the ‘algorithm’ within the agents can be updated constantly and in (26) it is shown how super computability can arise from the very evolution of the agents. Also, in an interactive machine, the traditional distinction between data, memory and algorithm does not apply, which results in more dynamical and less specifiable computational outcomes (27). Other classes of relatively down-to-earth machines which seem to guarantee to break classic computation barriers include fuzzy Turing machines (44).

Today human-computer interactions are standard in a large number of applications. Usually, these are seen as enhancing human capabilities by providing the fast computation resources available to electronic machines. Should we see the interaction in the opposite direction, as humans enhance the computational capabilities of electronic machines? In (23) it is speculated that today personal computers, connected via the web to thousands of machines world wide, receiving inputs via various sensors and on-line instructions from users, are already beyond classic computers. Today sensors monitor several aspects of the environment routinely and some have even been installed on animals in the wilderness (38). Can we envisage a network computing system, in which agents (computers) interact with the environment via analog sensors, receive data from

living beings, and instructions from humans to deal with unexpected situations? Could this be the way forward to understand emergence?

More intriguingly, could these systems potentially already sit on our desks?

#### 1.6.4 Data Explosion and Scientific Data Mining

In a recent issue of *Nature* (8; 28; 41), the picture was drawn of a near future when improved instrumentation and extensive sensing will provide us with exponentially increasing quantities of data for scientific enquiry. This implies more information but at a considerable cost. It promises more and better information about a vast range of things, from space to ocean depths, from ecologies to the human body, from genomes to social behaviour. However, the data explosion may go beyond our ability to process and analyse it. Unravelling new mysteries of Nature will then be jeopardised by something as mundane as lack of time and resources. It is hypothesised that this will be circumvented by clever software able to supervise the instrumentation, detection of new interesting patterns and possibly use rule extraction algorithms that uncover new processes and biophysical or social laws; a very difficult task, but (supposedly) merely a technological one.

This picture relies on 2 assumptions:

1. that all natural processes we may wish to study or detect are algorithmic;
2. that the process which allows us to understand and study Nature is also algorithmic.

Neither of these assumptions has been proved and both are open to debate. The first statement has been discussed above. The second one requires some clarification. First, a computational system which scans a data set in order to find patterns of interest must be algorithmic, by definition. Similarly, a system which, upon detecting a pattern, performs some rule extraction in order to attract our attention and suggest an interpretation also needs to be algorithmic. It seems evident that any algorithm capable of sifting through a stream of data and picking out just those novel patterns which are of interest to a human being, is more than a few steps along the way to passing the Turing test. Similarly, the second of the two systems bears a remarkable resemblance to the Halting Problem. An algorithmic system cannot, by definition, process a non-recursive language, from which it follows that if Nature displays a non-algorithmic process, this will not be detected by a fully automated computational system.

It is interesting to note that the rigors of formal logic apply not only to computational systems, but to the broader scientific method as well. The scientific method requires experiments to be reproducible. This implies that an experiment needs to follow a quite detailed and rigorous procedure in order to be replicated by different observers under inevitably different experimental settings. Basically, an experiment is reduced to an algorithm (39, page 122), and consequently scientific experimentation suffers the very same limitation of formal logic and computer systems, and thus is, by itself, unable to detect truly emergent processes. Curiously, the same desire for a rigorous, quasi-algorithmic approach affects scientific communication, with scientific journals often requiring a quasi-algorithmic way of writing. However, it is often suspected that

the large leaps in scientific understanding are fired by a brilliance which may be non-algorithmic. While further considerable work needs to be done to understand this creative process, it seems that over-relying on formal logic not only to model, but also to detect and analyse Nature may come with the risky consequence of preventing us from seeing the very processes we want to discover.

## 1.7 Conclusions

Our aim is by no means to suggest that computer modelling is a purposeless activity. Rather, that clarity is needed to discriminate the means (computer modelling as a tool) from the aim (acquiring knowledge about Nature). In this framework, confusing the means with the aim equates to carrying out a scientific program (including experimental and formal analysis) in the virtual world of a computer model as if this was the ‘real world’ and then extend the ‘virtual’ results to the ‘real’ natural world, under the assumption that the two are, to some degree, isomorphic. Here the old Chinese saying “if all you have is a hammer, everything looks like a nail” nicely highlights possible dangers and could be translated as ‘if all you have is a computer, everything looks computational’.

We can thus summarise our proposed guidelines as follows:

1. Care should be used to discriminate among: the processes which are ‘naturally’ amenable to computer modelling; the processes which are numerically or theoretically intractable (large combinatorial problems, NP-hard problems, chaotic problems) but for which useful approximations can be found (either in terms of non optimal solutions or large scale approximations); and processes which may be fundamentally intractable.
2. The widely accepted conjecture that intractable problems do not exist in Nature should (at least) be carefully studied, rather than accepted dogmatically.
3. The rigors of the algorithmic approach do not apply only to the world of formal systems and computer languages. Scientific investigation (the iterative testing of hypotheses) is also subject to these constraints due to its algorithmic nature. Recently, a new scientific tendency is to call for a more free and creative way of reporting and discussing science. Complex System Science, which naturally mixes experts ranging from pure mathematics to social science, seems to be in a particularly fortunate development for thorough exploration of the potential for reintroduction of artistic and other creative contributions to science.

## 1.8 Acknowledgments

This research was carried out as a part of the CSIRO Emergence Interaction Task, [http://www.per.marine.csiro.au/staff/Fabio.Boschetti/CSS\\_emergence.htm](http://www.per.marine.csiro.au/staff/Fabio.Boschetti/CSS_emergence.htm).

## References

- [1] Andersen, P., Emmeche, C., Finnemann, N., and Christiansen, P. (2000). *Downward Causation*. Aarhus University Press, Aarhus.
- [2] Arthur, W. (1994). Inductive behaviour and bounded rationality. *The American Economic Review*, 84:406–411.
- [3] Arthur, W. (1998). Modeling market mechanism with evolutionary games. *Europhys. News*, 29:5154.
- [4] Atay, F. and Josty, J. (2003). *On the emergence of complex systems on the basis of the coordination of complex behaviors of their elements*. Working paper, 04-02-005, Santa Fe Institute.
- [5] Bedau, M. (1997). Weak emergence. In Tomberlin, J., editor, *Philosophical Perspectives: Mind, Causation, and World*, vol. 11, pages 375–399. Blackwell Publishers.
- [6] Bickhard, M. H. (2000). Emergence. In Andersen, P. B., Emmeche, C., Finnemann, O., and Christiansen, V., editors, *Downward Causation*, pages 322–348. University of Aarhus Press, Aarhus, Denmark.
- [7] Boschetti, F. (2005). Improved resource exploitation by collective intelligence. In Zenger, A. and M., A. R., editors, *MODSIM05: International Congress on Modelling and Simulation: advances and applications for management and decision making*. Canberra, ACT.
- [8] Butler, D. (2006). 2020 computing: Everything, everywhere. *Nature*, 440(7083):402–405.
- [9] Calude, Campbell, Svozil, and Stefanescu (1995). Strong determinism vs. computability. In W. Depauli-Schimanovich, E. Koehler, F. S., editor, *Downward Causation*, pages 115–131. Kluwer Academic, Dordrecht.
- [10] Campbell, D. (1974). Downward causation in hierarchically organized biological systems. In Ayala, F. and Dobzhansky, T., editors, *Studies in the Philosophy of Biology*, pages 179–186. Macmillan Press.
- [11] Chaitin, G. (1997). *The limits of mathematics: a course on information theory & limits of formal reasoning*. Springer, New York.
- [12] Cooper, B. and Odifreddi, P. (2003). Incomputability in nature. In Cooper, S. B. and Goncharov, S. S., editors, *Computability and Models*, pages 137–160. Kluwer Academic, Dordrecht.
- [13] Corning, P. (2005). The re-emergence of emergence: a venerable concept in search of a theory. In *Holistic Darwinism: Synergy, Cybernetics, and the Bioeconomics of Evolution*. Univ. of Chicago Press, Chicago.
- [14] Crutchfield, J. (1994a). Is anything ever new? considering emergence, in complexity: Metaphors, models, and reality. In Cowan, G., Pines, D., and Melzner, D., editors, *SFI Series in the Sciences of Complexity XIX*, pages 479–497. Addison-Wesley, Redwood City.
- [15] Crutchfield, J. P. (1994b). The calculi of emergence: Computation, dynamics, and induction. *Physica D*, 75:11–54.
- [16] Darley, V. (1994). Emergent phenomena and complexity. In Brooks, R. and Maes, P., editors, *Proceedings of Artificial Life IV*, Cambridge, MA. MIT Press.

- [17] Emmeche, C., Koppe, S., and Stjernfelt, F. (2000). Levels, emergence, and three versions of downward causation. In Andersen, P. B., Emmeche, C., Finnemann, O., and Christiansen, V., editors, *Downward Causation*, pages 13–34. University of Aarhus Press, Aarhus, Denmark.
- [18] Gardner, M. (1970). Mathematical games: The fantastic combinations of john conway’s new solitaire game ”life”. *Scientific American*, 223:120–123.
- [19] Goldstein, J. (2002). The singular nature of emergent levels: Suggestions for a theory of emergence. *Nonlinear Dynamics, Psychology, and Life Sciences*, 6(4).
- [20] Heylighen, F. (1991). Modelling emergence. ftp papers from Principia Cybernetical server, also for ”world Futures: the Journal of General Evelopment”, special issue on creative evolution, G. Kampis (ed).
- [21] Kauffman, S. (2000). *Investigations*. Oxford University Press.
- [22] Kellett, O. (2006). *A multi-faceted attack on the busy beaver problem*. Phd thesis, Rensselaer Polytechnic Institute Troy, new York.
- [23] Leeuwen, J. v. and Wiedermann, J. (2000). *The Turing machine paradigm in contemporary computin*. Technical report uu-cs-2000-33, Institute of Information and Computing Sciences, Utrecht University.
- [24] Leeuwen, J. v. and Wiedermann, J. (2001a). Beyond the turing limit - evolving interactive systems. In Ruzicka, L. P. . P., editor, *Theory and Practice of Informatics*, pages 90–109. Springer-Verlag, Berlin.
- [25] Leeuwen, J. v. and Wiedermann, J. (2001b). *A computational model of interaction in embedded systems*. Technical report uu-cs-2001-02, Institute of Information and Computing Sciences, Utrecht University.
- [26] Leeuwen, J. v. and Wiedermann, J. (2003). The emergent computational potential of evolving artificial living systems. *AI Communications*, 15:205–215.
- [27] Milner, R. (1993). Elements of interaction. *Communications of the ACM archive*, 36(1):78–89.
- [28] Muggleton, S. (2006). 2020 computing: Exceeding human limits. *Nature*, 440(7083):409–410.
- [29] Ord, T. (2002). *Hypercomputation: Computing More Than the Turing Machine*. Technical report, University of Melbourne.
- [30] Pattee, H. (1997). Causation, control, and the evolution of complexity. In Andersen, P. B., Emmeche, C., Finnemann, O., and Christiansen, V., editors, *Downward Causation*. University of Aarhus Press, Aarhus, Denmark.
- [31] Pearl, J. (2000). *Causality: models, reasoning and inference*. MIT Press, Cambridge, Mass.
- [32] Penrose, R. (1989). *The emperor’s new mind : concerning computers, minds, and the laws of physics*. Vintage, London, Melbourne.
- [33] Penrose, R. (1994). *Shadows of the mind: a search for the missing science of consciousness*. Oxford University Press, Oxford.
- [34] R., L. and D., P. (2000). The theory of everything. *Proc of the National Academy of Science*, 97(1):28–31.
- [35] Rabinowitz, N. (2005). *CEmergence: an algorithmic formulation*. Phd thesis, The University of Western Australia.
- [Reynolds] Reynolds, C. Boids. <http://www.red3d.com/cwr/boids/>.

- [Rosenzweig] Rosenzweig, V. Wikipedia. [http://en.wikipedia.org/wiki/Loebner\\_prize](http://en.wikipedia.org/wiki/Loebner_prize).
- [36] Saygin, A., Cicekli, I., and Akman, V. (2000). Turing test: 50 years later. *Minds and Machines*, 10(4):463–518.
- [37] Shalizi, C. (2001). *Causal Architecture, Complexity and Self-Organization in Time Series and Cellular Automata*. Phd thesis, University of Michigan.
- [38] Simonite, T. (2005). Seals net data from cold seas. *Nature*, 438:402–403.
- [39] Stannett, M. (2003). Computation and hypercomputation. *Minds and Machines*, 13(1):115–153.
- [40] Sterrett, S. (2000). Turing’s two tests for intelligence. *Minds and Machines*, 10(4):541–559.
- [41] Szalay, A. and J., G. (2006). 2020 computing: Science in an exponential world. *Nature*, 440(7083):409–410.
- [42] Turing, M. A. (1936). On computable numbers, with an application to the entscheidungsproblem. *Proceedings of the London Mathematical Society*, 42:230–265.
- [43] Verbaan, Leeuwen, v., and Wiedermann (2004). Lineages of automata - a model for evolving interactive systems. In et al, J., editor, *Theory is for ever*, pages 268–281. Springer-Verlag, Berlin.
- [44] Wiedermann, J. (2000). *Fuzzy Computations Are More Powerful than Crisp Ones*. Technical report v-828, Prague University.
- [45] Wiedermann, J. and van Leeuwen, J. (2002). The emergent computational potential of evolving artificial living systems. *AI Communications*, 15(4):205–216.