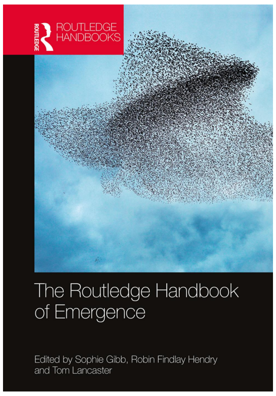

The Routledge Handbook of Emergence

Sophie Gibb, Robin Findlay Hendry, Tom Lancaster

Computational Emergence

Publication details
https://www.routledgehandbooks.com/doi/10.4324/9781315675213-10
Mark Pexton
Published online on: 22 Mar 2019

9

COMPUTATIONAL EMERGENCE

Weak and strong

Mark Pexton

Introduction

Imagine you have a piece of square graph paper and a marker pen. You colour in some of the squares in the first column of the paper, leaving others blank. You then move to the second column, and colour in squares in this according to some rule relating to the position of coloured squares in the first column. For example, the rule could be that for every coloured square in column 1, the equivalent square in column 2 is left blank, and every blank in column 1 becomes coloured in column 2. Imagine you now repeat this process for several thousand iterations to see what pattern of coloured squares arises. This setup captures in essence what is known as a cellular automaton in computer science. A cellular automaton is a grid with initial conditions and a set of rules for propagating those initial conditions across the grid. Cellular automata are the basis behind Conway's famous Game of Life (Conway 1970). For the vast majority of initial conditions and rules the Game of Life is uninteresting, either completely random or highly predictable distributions are produced. But for some combinations of rules and initial conditions, striking, almost life-like, distributions arise. These distributions are complex, existing in the hinterland between randomness and regularity. Even though all the rules of the Game of Life, and all the initial conditions, are known perfectly, structures can still appear which defy prediction from that perfect knowledge of the system.

This lack of predictability flies in the face of the reductionist intuition behind Laplace's demon. The standard demon story is that if we had perfect knowledge of the laws of physics and starting conditions, we could predict everything that follows (or objectively specify the probability of everything that follows in a quantum universe). Examples such as the Game of Life, and the growing fields of complexity and chaos science, suggest that even for the demon some systems defy predictability. This realisation has led to a new approach to explicating emergence based on notions derived from computer science. This *computational emergence* has two mutually exclusive forms. The first, developed primarily by Bedau (1997, 2002, 2008) and expanded upon by Huneman (2008, 2012), is called weak emergence. Newman (1996) has also discussed similar ideas (in the context of chaos theory but with a different formalism), as have Boschetti & Gray (2013). Weak emergence is a particular way of classifying a special set of complex resultant phenomena. The second approach is pancomputational emergence, first suggested by Davies (2004, 2007, 2010); Walker, Cisneros & Davies (2012); and Pexton (2015). Pancomputational emergence is a strong form of emergence which identifies higher-level aggregate entities as necessary for the universe to "calculate" the behaviour of complex systems.

Weak emergence

Bedau developed his account of weak emergence inspired by cellular automata, and he makes two claims of them. First, that cellular automata can themselves display weak emergence. Second, that cellular automata are a good model for many real-world systems, and these natural systems also display weak emergence. As stated, in certain scenarios, despite perfect knowledge, cellular automata can produce a confluence of highly context-specific circumstances which produce complex patterns that could not be predicted in advance. This lack of predictability is the hallmark of weak emergence. Imagine we have a macrostate *M*, with a corresponding microstate *m*, at time t_1. Then the system evolves at time t_2 to a state (*M**, *m**). If there is no way in principle to predict from *m* what *m** will be other than running through each intermediate stage between t_1 and t_2, then we have weak emergence. Bedau dubs this running through each stage simulation, since one must use a computer simulation which runs through each step to explain the system's behaviour in microphysical terms.

Weak emergence stands as a subset of resultant behaviour and is perfectly compatible with reduction. It is incompatible with a system displaying strong emergence. This is because all causation in weakly emergent systems is ultimately microphysical causation. All higher-level causal attributions are merely summaries of the aggregate effect of the "micro-causal web". The special category of resultant phenomena that weak emergence captures are those systems that we have to let "crawl the causal web" in order to know what will happen.

> Each concrete physical embodiment of weak emergence is ontologically nothing more than some kind of aggregation of smaller embodied objects . . . each example of macro-level weak emergence is ontologically and causally reducible to micro-level phenomena. However, in practice, typically nobody can understand or follow such a micro-causal reduction unless they simulate the micro-causal web on a computer, because the micro-level causal web is so complex. In a wealth of interesting cases, studied in fields like soft artificial life, computer simulations make it possible to crawl the causal web.
>
> *(Bedau 2008, p. 448)*

Bedau expresses the need to simulate in two very different ways: algorithmic incompressibility and explanatory incompressibility. Explanatory incompressibility is the less precise idea. The basic intuition is that to explain the behaviour of a weakly emergent system, we cannot do better than let it develop and observe its pathway through state space. Explanatory incompressibility is not well defined though, and any definition will no doubt depend on what account of explanation one begins with. Fortunately, Bedau's other way of specifying the necessity to simulate, algorithmic incompressibility, is borrowed from information theory and is more precise.

Algorithmic incompressibility is defined in terms of Kolmogorov complexity. Kolmogorov complexity, or algorithmic complexity, was independently proposed by Solomonoff (1964), Kolmogorov (1965) and Chaitin (1969). Imagine that one wishes to describe an object using a string of binary. The object can have many different possible descriptions, but we are interested in the shortest possible description. The length of the description gives us a measure of the object's intrinsic complexity. We can use computational algorithms to make this more precise: the Kolmogorov complexity of an object X is the length (in bits) of the shortest possible computer program that prints X and then halts. (See Grünwald & Vitányi 2008 for a comprehensive review.)

So in a weakly emergent system the Kolmogorov complexity is at a maximum, and we cannot compress our algorithm at all. There is no shortcut to get from *m* to *m**. By contrast, in a non-weakly emergent system, there are enough regularities to mean that the Kolmogorov complexity of the system is much reduced. We can therefore compress our algorithm and provide a shortcut to extrapolate *m** from *m*.

The metaphysics of weak emergence

What then of the metaphysical status of weak emergence? As we have seen, weak emergence is a particular form of resultant behaviour; it is not "emergent" at all in the traditional usage of the term. Moreover, it is incompatible with strong emergence. Any strongly emergent system would not be fully determined by microphysical events. Hence, such a strongly emergent system could not in principle also be weakly emergent.

Bedau claims that cases of epistemic emergence merely reflect an epistemic agent's ignorance of microphysical facts. But weak emergence is not like that – even the demon could not compress the algorithms and predict the outcome of a weakly emergent system. In his 1998 paper *Is Weak Emergence Just in the Mind?* Bedau is keen to stress the objectivity of weak emergence. It is not that we happen not to be able to compress the algorithms, it is that no such compression is possible in principle, because of the informational/causal structure of the system.

> It is presumably not an accident that one sort of micro-causal structure is incompressible and another sort is compressible. So, though weak emergence [is a kind of] "epistemological" emergence, weak emergence is not merely epistemological. It is not just in the mind. Instead, weak emergence results from incompressible macro-level structure in the network of micro-level causal connections.
>
> *(Bedau 2008, p. 451)*

And:

> The micro-causal web is real and objective, and the incompressible causal pathways of weak emergent phenomena have a distinctive epistemological consequence. Note that the explanatory incompressibility that defines weak emergence applies to the explanations of any naturalistic epistemic agent, in principle. Just like us, any non-human epistemic agent will have to work through the objective complexity of the local micro-causal interactions. Thus, weak emergence is not merely in the mind, but refers to objective complexity in the objective natural world that is in principle irreducible in practice.
>
> *(Bedau 2008, p. 453)*

As the Game of Life shows, it is not ignorance of laws and starting conditions that prevents predictability, it is just a fact about the complex, context-sensitive, open nature of the system that means one must simulate it. It is the objective micro-causal structure of the system that makes it weakly emergent.

Open questions for weak emergence

The problem of predictability

The first challenge for weak emergentist claims concerns predictability. Predictability is an inherently epistemic concept; hence, weak emergence cannot avoid being a species of epistemic emergence. In order to predict something or not, one must first specify a relevant timescale and level of precision the prediction is to be judged by. Even a chaotic system is predictable on very short timescales (if it weren't, it couldn't be simulated) and on very long timescales. For instance, consider the UK weather system – undoubtedly a chaotic system. Yet in some respects, it is very predictable. If it is sunny one second, it is highly likely it is sunny one second later, and similarly it is highly likely there will be a period of winter at least once a year.

Note that it is impossible to provide a general rule for proving if an algorithm is incompressible in principle but not practice (see Ming & Vitányi 1997). The only way of having a general proof would be to actually compress the string; hence, we can never know for sure whether a string that hasn't been compressed could be compressed or not. However, we can use the fact that attempts to compress an algorithm have failed as empirical evidence that an algorithm is incompressible. So when Bedau claims that no observer in principle could compress an algorithm of a particular system, then this is a provisional statement. We cannot in general say by looking at the micro-causal structure of a system that it is impossible to compress it, only that we have evidence that it may be incompressible. Moreover, such empirically based claims have more weight for artificial systems such as the Game of Life. These are constructed, and we have perfect information concerning every aspect of the system, yet they defy our compression attempts. But for any real system we are not in this position; it is more of an open question whether perfect information concerning starting conditions would allow compression or not. That is not to say that it is not plausible that there are weakly emergent systems in nature, of course.

It is not only timescales that matter but also precision. Bear in mind the distinction between precision and accuracy here. Two people can use the same ruler, where the smallest division marked is 1 mm. One person may be good at using the ruler, the other bad, and so their measurements will have different accuracies. But both measurements will have the same precision: 1 mm. The ruler is not capable of measuring to a greater precision. Of course, one can always use a different ruler, with finer divisions, and this will have a greater precision. There is no mind-independent level of precision; it is an epistemic concept. Now for cellular automata, the level of precision is picked at the point of construction, so all observers agree in advance by design. The "graph paper" is constructed with a certain density of squares. But for real systems, there is no automatic level of precision to view them at or one that all observers must agree to in advance. Now whether a given system is predictable or not depends on the precision one picks to predict it to. So if we are interested in very coarse-grained structural properties, then many chaotic/complex systems will be predictable; that is why there is a science of chaos/complexity in the first place.

Going to the other extreme, if we are fine grained enough, then many ordinary resultant systems cease to be predictable. Consider a very predictable system: a simple pendulum. If we know the pendulum's length, it is a straightforward matter to predict its period to within a second or millisecond, for instance. But say we change our precision and want to predict the period to within a femtosecond or less. For this degree of precision, each actual physical pendulum is completely unpredictable. We might say a pendulum is predictable, but the weather is not, so therefore the weather is weakly emergent. But we can only make this claim because we have implicitly assumed a precision metric. For a different metric both systems would be predictable or both weakly emergent. Which scale we as humans think is relevant is an epistemic concern.

Bedau is correct that whether all epistemic agents would agree on compressibility *once they had agreed on a precision metric* is objective. Compressibility reflects the conjunction of observer-independent facts about the system and the observer-dependent chosen metric. But this objectivity that all *views from somewhere* share is different from weak emergence reflecting a view from nowhere. There is no sense in which an *observerless* universe could be said to have weakly emergent systems.

The problem of randomness

This last point leads to another potential difficulty with weak emergence: Does it adequately distinguish between random states and complex states? As the pendulum example shows, at a fine enough level of precision, nearly all systems contain random elements. This has led McAllister (2013) to contend that all systems are algorithmically incompressible. This is because all systems are made up of structured patterned components and random components. But to compress an algorithm, one must reproduce everything exactly; it is no good replacing a long algorithm with a shorter algorithm that is only nearly the same. Compression requires it to be exactly the same. So even a small random element in a system means that, according to McAllister, it cannot be compressed.

If McAllister is correct, then one cannot use algorithmic incompressibility as the measure of weak emergence. One could still use explanatory incompressibility, but it isn't clear this notion has any rigorous definition. It certainly doesn't have a definition which would apply regardless of which account of explanation one uses. Regardless of this objection, basing compression on explanatory concerns suggests weak emergence is even more epistemically contextualised. As an observer one can simply ignore the aspects of the system that are not compressible because one is uninterested and summarise the rest in some coarse-grained, higher-level rule, but it is difficult to claim that that is not an epistemic activity. Why must all observers agree about which aspects of the system they think are explanatorily salient or not? This will depend on what they wish to explain and what level of detail they wish that explanation to provide.

Even if one rejects McAllister's argument, this only allows weak emergence to distinguish complex systems from regular systems. But weak emergence still faces the challenge of distinguishing complex systems with patterns in randomness from truly random systems. A completely random system is definitely incompressible, yet such systems do not fit into the spirit of what one might think weak emergence should be about. A random data set of pure noise might be incompressible and not predictable, but is it weakly emergent? Emergence implies that there is some pattern – it is defined positively by what the system has, not negatively by what the system lacks (compressibility, predictability, etc.). Complex systems are a balance between order and randomness, where unexpected patterns emerge. Yet characterising this as incompressibility means both complex systems and simply random systems are weakly emergent. Complexity requires randomness, but also patterns and weak emergence, as incompressibility alone says nothing about this structured component of complex systems.

Strong emergence

In contrast to weak emergence there are accounts of strong/ontological emergence based on computational ideas, specifically pancomputationalism. Pancomputationalism is usually accredited to Konrad Zuse, as summed up by "Zuse's Thesis" (ZT):

> ZT: The world is a computer and physical processes are algorithms that are executed on that computer.

The computer could be a cellular automaton as argued by Zuse himself (Zuse 1969), as well as Von Neumann & Burks (1966), Fredkin (2003) and Wolfram (2002). Or it could be a universal Turing machine (Schmidhuber 1997) or a quantum computer (Lloyd 2002).

If the universe is viewed as a computer, then it is natural to consider the physical limitations on that computer. This approach is exemplified by Landauer:

> The laws of physics are essentially algorithms for calculation. These algorithms are significant only to the extent that they are executable in our real physical world. Our usual

> laws of physics depend on the mathematician's real number system. With that comes the presumption that given any accuracy requirement, there exists a number of calculational steps, which if executed, will satisfy that accuracy requirement. But the real world is unlikely to supply us with unlimited memory or unlimited Turing machine tapes. Therefore, continuum mathematics is not executable, and physical laws which invoke that cannot really be satisfactory.
>
> *(Landauer 1996, p. 192)*

Calculations in our world require physical instantiation: computers are real objects bounded by practical limitations. It is the laws of physics that determine the amount of information that a physical system can register (the number of bits) and the number of elementary logic operations that a system can perform (number of operations, or ops).

Information and computation

There are two ways of investigating the limits on computation: 1) make an estimate of the amount of information in a pancomputational universe and 2) make an estimate of how hard the algorithms are that a pancomputational universe executes.

The first of these has been pursued by Lloyd (Lloyd 2002), Davies (2004, 2007, 2010) and Gough (2013) independently. Lloyd has estimated the total number of bits and the total number of operations on those bits since the universe began. Lloyd's calculation is based on looking at some of the physical limitations on computation. These include the constraints on information processing due to entropy, energy and the limit on signalling to be no faster than the speed of light. Lloyd calculates that there are no more than 10^{120} ops on 10^{90} bits (or 10^{120} bits if gravitational degrees of freedom are taken into account) in the universe. A similar number is calculated by Davies from considering the holographic principle and by Gough by considering the current distribution of baryonic matter.

Consideration of these informational limits suggests that reductionism in a pancomputationalist universe may be implausible. The reason is that many higher-level special science systems seem to exceed the limits of microphysical computation. For example, Davies (2004) suggests that the combinatorics of many special science systems exceed Lloyd's information bound. One example is the enzymatic efficacy of a protein (see Luisi 2002). There are 20 varieties of amino acids: a peptide chain of *n* amino acids can be arranged in 20^n different ways, and each sequence can adopt an enormous number of different conformations. If we assume each amino acid can adopt five different orientations (Fasman 1989) then the total number of conformations is 5^n. Combining 20^n and 5^n means that we have about 10^{2n} different molecular structures. Now *n* is typically of the order of 100 for a small protein. This means there are 10^{40} times more bits of information needed to simulate the protein step by step at the level of microphysics then there are bits in the accessible universe.

For Davies, it is striking that the upper bound on information processing is exceeded by the complexity of biological systems. Davies argues that this provides good evidence that biological systems are strongly emergent. Davies expresses this form of emergence in causal terms: a pancomputationalist universe cannot be causally closed at the microphysical level, since there is not enough information to compute living systems using informational resources from the microphysical level alone. Davies concludes the informational limits of the universe suggest the causal closure of physics (CCP) is violated.

In a paper co-authored with Walker and Cisneros (Walker, Cisneros & Davies 2012) Davies also suggests that top-down causation can be understood in informational terms. Walker, Cisneros and Davies use major evolutionary transitions in biology as a case study, such as the transition from prokaryotic life to eukaryotic life. They argue that such transitions can be understood as transitions from bottom-up causation to top-down causation, in which higher-organising principles determine the system's behaviour. They argue that such a reversal of the direction of causal power can be identified with a reversal in the flow of information. Bottom-up causation requires information to flow from lower levels to higher levels, whereas top-down causation requires information flow from higher levels to lower levels.

Whilst Davies has identified a new approach to understanding strong emergence, there are challenges to be faced by his account. The most severe is with his identification of information flow with causal determination. There are many different accounts of causation in the philosophical literature. Causation can be based on a set of privileged physical interactions, counterfactual reasoning or powers/capacities. Yet Davies does not make connection to any of these, and it would be desirable to have an explicit account of exactly what is meant by claiming information flow is a form of causal determination.

Algorithmic complexity and computation

The second constraint on computation comes from considering the inherent difficulty of the algorithms themselves (Pexton 2015). In computer science the difficulty of executing an algorithm is evaluated by imagining how difficult such a problem would be to solve for a Turing machine. A Turing machine is a type of thought experiment: an ideal computer invented by Alan Turing (Turing 1937) to probe the limits of calculability. There are two basic constraints on a Turing machine: time and space. Now in computation, each problem, each question, needs to be coded, and some problems will take more code to specify them. Let n be the input size of a problem: we then say a Turing machine can solve this problem in *polynomial time* if the time taken to solve this, $T(n)$, is only a function of n raised to various powers. The problem is dominated by its largest term, so, for instance, if:

$$T(n) = n^2 + 3n + 4$$

then we say that this problem is solvable in polynomial time of order n^2.

Polynomial time contrasts with exponential time. In this it might take on the order of 2^n steps to solve a problem. If we have $n = 100$ and say 10^{12} operations per second, then this exponential time would be on the order of the age of the universe. These limits of ideal computing define different complexity classes as shown in Figure 9.1. It is an open question whether some of these classes are equivalent or not. The aim here is not to rest too much philosophically on the precise boundaries as currently understood. Instead, it is to make the case that some computational problems are difficult in *principle*. The difficulty of these problems is not a limitation of actual technology, but rather an in-principle limitation of what an ideal computer can compute, given that such a computer is a physical object of some kind or another and is therefore not able to use infinite amounts of time or space.

In algorithmic complexity theory, one of the most studied elements of the hierarchy is the relation between the classes P and NP. The class P is the set of all problems that can always be solved by using a polynomial time algorithm. The class NP is the set of problems for which, if we are given a solution, we can verify that solution in polynomial time using a Turing machine. The class P is contained within NP, but it is an open question whether P = NP (mainstream opinion holds that P does not equal NP). The category NP-Hard is the set of problems that are at least as hard as any NP problem. NP-Complete problems are those that are both NP-Hard

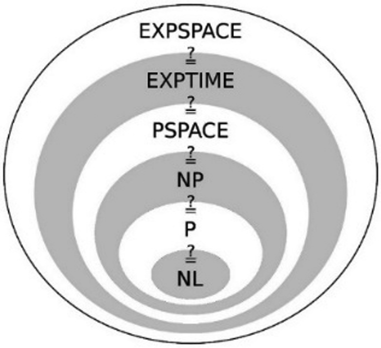

Figure 9.1 The hierarchy of complexity classes, starting with the easiest, NL, P and moving to the most difficult, EXPTIME and EXPSPACE. The question marks reflect the open question of what the relationships between these categories are.

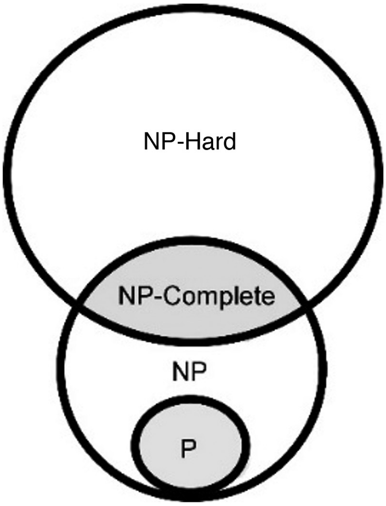

Figure 9.2 Class of NP-Complete for P not equal to NP

and in NP, that is, they are the hardest problems in NP (see Figure 9.2). So although if we are given a solution to an NP-Complete problem we can verify it relatively easily, there is no efficient way to find a solution in the first place. Efficiency is the key concept: for an NP-Complete problem the time required to solve the problem grows quickly as the size of the string inputting the problem increases.

So the lesson from computer science is that some calculations must in principle take longer than the age of our universe in principle (or any physical universe that shared some basic common features such as conservation of energy).

It is possible that some problems in other sciences, when expressed algorithmically, are NP-Complete. For example, Barahona (1982) suggests that spin glass models of phase transitions represent NP-Complete problems, while Welsh (1993) suggests that certain problems in knot theory relating to DNA are potentially NP-Complete or even PSPACE-Complete. One particularly interesting candidate system is protein folding. Fraenkel (1993), Crescenzi et al. (1998) and Berger & Leighton (1998) have each proposed that protein folding is an example of a physical system solving an NP-Complete problem. The status of these examples is controversial though; see Aaronson (2005) for arguments that these cases are not NP-Complete.

So the basic idea is that these problems *expressed in microphysical terms alone could not be calculated given the space, time and informational constraints in a pancomputationalist universe.* But given that these processes definitely occur, this means that a pancomputationalist universe must utilise higher-level terms to reduce both the difficulty of the problems and the amount of information processing needed to solve them.

The Kolmogorov complexity of the algorithms capturing various microphysical processes are reducible if nature can make use of higher-level structural features of those processes. This compression is not merely the epistemic compression of an observer wishing to predict the outcome of a process efficiently. Instead, the compression must be something a pancomputationalist world does itself; otherwise, it does not have the calculational resources to determine the outcome of those microphysical calculations. By using higher-level terms a pancomputationalist world can vastly simplify the informational requirements to perform calculations, thereby allowing it to circumvent the informational limits calculated by Lloyd and to compress algorithms to find solutions to problems that couldn't otherwise be efficiently solved. By applying a pancomputationalist model to the laws of physics, it suggests that the robust higher-level variables of special sciences are not just the epistemic creations of a particular way of viewing the world, but are in fact an essential feature of how nature itself structures the world for calculational purposes.

So we have strong pancomputationalist emergence when:

i We have a system that cannot be informationally compressed at the microphysical level.

And

ii That system finds a solution in a timescale that exceeds the calculational limits it has at the microphysical level.

Since the system's evolution is "computed", then, given the previously stated constraints, the universe must compress the algorithm (reduce its Kolmogorov complexity) by utilising structural features/terms that are at a higher scale as informational processing resources.

These considerations lead to an information-theoretic definition of strong emergence in terms of the scale-relative compressibility of the Kolmogorov complexity of algorithms. The algorithms can only be compressed to be executable within the salient calculational limits if structural features/kind terms applicable to a higher level are used to reduce the Kolmogorov complexity of the algorithms.

Scale Relative Compressibility-emergence (SRC-emergence)

> A system is SRC-emergent if it must use higher-level structural relations and terms to reduce the Kolmogorov complexity of the algorithms that represent that system (such that those algorithms only then become executable given the calculational constraints on the system).

So the basic idea of SRC-emergence is that a pancomputationalist universe uses higher-level kind terms and structural facts as elements in algorithms. By using these higher-level features, those algorithms can be compressed. One way of appreciating this idea is to consider minimum message length encoding, or MML (see Twardy, Gardner & Dowe 2005). In MML, instead of compressing raw data, we look for the most compressed version of a two-part algorithm: the first part states a theory capturing structural features of the data set, and the second part states the data under the assumption that the theory is true. Observational data sets can be thought of as encoding a set of events. Optimal coding theory states that such a sequence will be most efficiently encoded if we use code words of length-log(p) to encode events with probability p. Essentially, the more likely an event is, the shorter the code word we give it; then when we write the whole sequence of events down, the sequence will be compressed. It is knowledge of the structure of a problem that can allow us to allocate the most efficient coding scheme and compress the data.

Let's illustrate this with an example. Andrew is a scientist who wants to test a die with sides A, B, C and D. He rolls it 20,000 times and wishes to send the data set to his friend Jacqui for analysis. The record of the die rolls is an empirical data set with plenty of random noise. Each result can be encoded in binary, say: 00 for A, 01 for B, 10 for C and 11 for D. If McAllister is correct this data set cannot be compressed and the most efficient way of sending the set is just to list the results themselves; therefore, it would take 20,000 bits for Andrew to send Jacqui his results.

But optimal coding theory allows Andrew to do better in certain circumstances. For instance, imagine Andrew suspects the die is biased so that $\Pr(A) = 1/2$, $\Pr(B) = 1/4$, $\Pr(C) = 1/8$ and $\Pr(D) = 1/8$. Now using optimal coding theory Andrew can give each outcome a code name of length $-\log(p)$. So outcome A will have code of length 1, 2 for B and 3 for C and D: for example, 0 stands for A, 10 for B, 110 for C and 111 for D. So Andrew now sends Jacqui the set encoded like this and a piece of code relating the theory about the die. If Andrew's theory about the bias of the die is correct, then he has now compressed the data set. Instead of 20,000 bits it will now only take 17,500 bits to send the data to Jacqui.

Remarkably, even if Andrew's theory is only approximately correct, he will still have compressed the data. Let's say the actual probabilities were $\Pr(A) = 4/9$, $\Pr(B) = 3/9$, $\Pr(C) = 1/9$ and $\Pr(D) = 1/9$. Now the expected length for sending the data is 17,778 bits; it is not as compressed as if Andrew's theory were exact, but it is still a huge compression in comparison to sending the original data.

How does MML apply to emergence in the context of pancomputationalism? In essence, if we think of laws of nature as algorithms which are executed, then SRC-emergence implies that a pancomputationalist universe must use some special science regularities to allow it to calculate certain physical outcomes. So, for instance, certain higher-level structural facts about the energy landscape of protein folding might be involved in allowing nature to compute which optimal shape a protein should fold to in only 1 sec (rather than the billions of years it would take if it were merely a random microphysical search). Notice that in MML the theory doesn't need to be exactly correct, and hence our special science law doesn't need to be exact and exceptionless – it only has to capture enough facts about probability/counterfactual dependencies to allow a pancomputational universe to compress a calculation. Notice also that MML means that there is no problem with randomness at the microphysical level. The particulars of each concrete microphysical situation (the random "noise" in the data set) do not prevent compression using higher-level algorithmic laws, since these only need to be approximate. (This is contra McAllister 2013.)

Strong and weak pancomputationalism

There are two different strengths of pancomputationalism with different metaphysical commitments:

i The universe can be adequately modelled as a computer (*weak pancomputationalism*)
ii The universe literally is a computer (*strong pancomputationalism*)

Davies subscribes to strong pancomputationalism, whereas Pexton (2015) contends that only weak pancomputationalism need be plausible for strong computational emergence to be legitimate. Davies' position is:

> [W]hen it comes to the information bound on the universe, one is forced to confront the status of information: is it ontological or epistemological? If information is simply a description of what we know about the physical world . . . there is no reason why Mother Nature should care about the limit [on the total amount of information]. But if information really does underpin physical reality – if it, so to speak, occupies the ontological basement . . . then the bound on the universe represents a fundamental limitation on all reality, not merely on states of the world that humans perceive.
>
> *(Davies 2007, p. 10)*

Although perfectly defensible, Davies' advocacy of strong pancomputationalism is unnecessary: merely subscribing to weak pancomputationalism still can tell us interesting facts about how the world is constrained. Even though we are only *modelling* the world in computational terms, that modelling step gives us access to physical constraints that would otherwise be obscured. It is not that a weak pancomputationalist universe is constrained by algorithmic features directly; rather, thinking of the world in terms of algorithms replicates certain structural constraints that would be hard to state in non-informational/computational terms.

The Platonic/anti-Platonic debate in mathematics provides a parallel with the debate about whether information theoretic models should be viewed realistically or not. There is little doubt that mathematics is theoretically indispensable to physics – almost all models in physics are expressed mathematically. Certain models could not work without mathematics, and the truths those models reveal to us about the world would be hidden without mathematics.[1] There are two responses to the indispensability of mathematics: the first is Platonism first articulated in modern terms by Frege (1953), in which mathematical objects are real, and, in some versions of Platonism, underpin the physical world and constrain it (see Baker 2009, for an advocate of mathematical constraints).

However, an alternative view is fictionalism about mathematical objects (Balaguer 1996). In this view, mathematical objects do not exist; instead, they are an epistemic construct of mathematicians. However, the fact mathematical objects are fictions does not prevent mathematical models from helping us to express facts about the world. The reason is that there are structural similarities between the constructed mathematical realm and the physical realm. This mapping from mathematics to the real world allows mathematics to capture physical features of the world (Pincock 2004; Leng 2002). Mathematics is applicable in the way a street map is. The map shares structural similarities to the city it corresponds to, allowing a user to identify facts about the city from facts about the map. Note that the map doesn't have to capture every relation in the city, it only has to capture some salient ones. So, for instance, we can model tyres as circles because circles and tyres share some structural similarities, without having to believe that ideal circles exist and constrain the tyre in any way.

Bueno and Colyvan (2011) have extended the mapping idea into an inferentialist account of mathematics. In this, mathematics works not just by a straightforward mapping of features but also by allowing inferences to be made in two stages. The idea is that since certain mathematical objects can share certain structural features of the world, we can then represent the world in abstract mathematical terms. Once we have constructed the mathematical representation, we can use the relation between those directly mapping pieces of mathematics and other (non-directly mapping) pieces of mathematics (dubbed surplus structures) to make mathematical inferences which we can then apply back to the physical world (see Bueno & French 2012; Pexton 2014). The inferences this process allows may otherwise be impossible to make.

Note, however, that for the inferentialist account of mathematics to work, we do not need to believe the world is underpinned by real mathematical objects and laws that constrain physical reality. Mathematical constraints are placeholders for physical constraints that might otherwise have never been revealed without the mathematical modelling step. Information and weak pancomputationalism can be viewed in similar inferentionalist terms. That is, there are structural relations in the physical world which map onto computational relations. By modelling the world in computational terms, we can then make inferences using secondary computational relations and apply them back to the real world. The computational constraints on the universe are then to be viewed as a way of summarising physical constraints on the world. If such an inferentialist view of pancomputationalism is defensible then, contra Davies, we do not have to defend strong pancomputationalism in order to use information theoretic models to explore issues such as emergence. Although strong pancomputationalism may turn out to be true, weak pancomputationalism is much easier to establish. In addition, as with the map and city example, weak pancomputationalism does not have to model every aspect of the world, just certain relations that are salient to the question of emergence.

Note

1 Field (1989) has attempted to construct physical theories without mathematics, but his account faces severe difficulties; see Pincock (2007).

References

Aaronson, S. (2005). Guest column: NP-complete problems and physical reality. *ACM Sigact News*, *36*(1), 30–52.

Baker, A. (2009). Mathematical explanation in science. *The British Journal for the Philosophy of Science*, *60*(3), 611–663.

Balaguer, M. (1996). A fictionalist account of the indispensable applications of mathematics. *Philosophical Studies*, *83*(3), 291–314.

Barahona, F. (1982). On the computational complexity of Ising spin glass models. *Journal of Physics A: Mathematical and General*, *15*(10), 3241.

Bedau, M. A. (1997). Weak emergence. *Noûs*, *31*(s11), 375–399.

Bedau, M. A. (2002). Downward causation and the autonomy of weak emergence. *Principia: An International Journal of Epistemology*, *6*(1), 5–50.

Bedau, M. A. (2008). Is weak emergence just in the mind? *Minds and Machines*, *18*(4), 443–459.

Berger, B., & Leighton, T. (1998). Protein folding in the hydrophobic-hydrophilic (HP) model is NP-complete. *Journal of Computational Biology*, *5*(1), 27–40.

Boschetti, F., & Gray, R. (2013). A Turing test for Emergence. In *Advances in Applied Self-Organizing Systems* (pp. 401–416). London: Springer.

Bueno, O., & Colyvan, M. (2011). An inferential conception of the application of mathematics. *Noûs*, *45*(2), 345–374.

Bueno, O., & French, S. (2012). Can mathematics explain physical phenomena? *The British Journal for the Philosophy of Science*, *63*(1), 85–113.

Chaitin, G. J. (1969). On the length of programs for computing finite binary sequences: Statistical considerations. *Journal of the ACM (JACM)*, *16*(1), 145–159.

Conway, J. (1970). The game of life. *Scientific American*, *223*(4), 4.

Crescenzi, P., Goldman, D., Papadimitriou, C., Piccolboni, A., & Yannakakis, M. (1998). On the complexity of protein folding. *Journal of Computational Biology*, *5*(3), 423–465.

Davies, P. C. W. (2004). Emergent biological principles and the computational properties of the universe: Explaining it or explaining it away. *Complexity*, *10*(2), 11–15.

Davies, P. C. W. (2007). The implications of a cosmological information bound for complexity, quantum information and the nature of physical law. *Fluctuation and Noise Letters*, 7(4), C37–C50.

Davies, P. C. W. (2010). The implications of a holographic universe for quantum information science and the nature of physical law. www.ctnsstars.org/conferences/papers/Holographic%20universe%20and%20infor mation.pdf
Fasman, G. D. (1989) *Prediction of Protein Structure and the Principles of Protein Conformation*. New York: Plenum.
Field, H. H. (1989). *Realism, Mathematics, and Modality*. Oxford: Blackwell.
Fraenkel, A. S. (1993). Complexity of protein folding. *Bulletin of Mathematical Biology*, *55*(6), 1199–1210.
Fredkin, E. (2003). An introduction to digital philosophy. *International Journal of Theoretical Physics*, *42*(2), 189–247.
Frege, G. (1953). *Foundations of Arithmetic*. Oxford: Blackwell. Transl. by J. L. Austin.
Gough, M. P. (2013). Holographic dark information energy: Predicted dark energy measurement. *Entropy*, *15*(3), 1135–1151.
Grünwald, P., & Vitányi, P. (2008). Shannon information and Kolmogorov complexity. https://arxiv.org/abs/cs/0410002
Huneman, P. (2008). Combinatorial vs. computational views of emergence: Emergence made ontological? *Philosophy of Science*, *75*, 595–607.
Huneman, P. (2012). Determinism, predictability and open-ended evolution: Lessons from computational emergence. *Synthese*, *185*(2), 195–214.
Kolmogorov, A. N. (1965). Three approaches to the quantitative definition of information. *Problems of Information Transmission*, *1*(1), 1–7.
Landauer, R. (1996). The physical nature of information. *Physics Letters A*, *217*(4), 188–193.
Leng, M. (2002). What's wrong with indispensability? (Or the case for recreational mathematics), *Synthese*, *131*, 395–417.
Lloyd, S. (2002). Computational capacity of the universe. *Physical Review Letters*, *88*(23), 237901.
Luisi, P. L. (2002). Emergence in chemistry: Chemistry as the embodiment of emergence. *Foundations of Chemistry*, *4*(3), 183–200.
McAllister, J. W. (2013). Empirical evidence that the world is not a computer. In *Imagine Math 2* (pp. 127–135). Milan: Springer.
Ming, L., & Vitányi, P. (1997). *An Introduction to Kolmogorov Complexity and Its Applications*. Heidelberg: Springer.
Newman, D. V. (1996). Emergence and strange attractors. *Philosophy of Science*, 245–261.
Pexton, M. (2014). Can asymptotic models be explanatory? *European Journal for Philosophy of Science*, *4*(2), 233–252.
Pexton, M. (2015). *An Information Theoretic Critique of the Completeness of Physics*. forthcoming.
Pincock, C. (2004). A revealing flaw in Colyvan's indispensability argument. *Philosophy of Science*, *71*, 61–79.
Pincock, C. (2007). A role for mathematics in the physical sciences. *Noûs*, *41*(2), 253–275.
Schmidhuber, J. (1997). A computer scientist's view of life, the universe, and everything. In *Foundations of Computer Science* (pp. 201–208). Berlin and Heidelberg: Springer.
Solomonoff, R. J. (1964). A formal theory of inductive inference, parts 1 and 2. *Information and Control*, 7, 1–22, 224–254.
Turing, A. M. (1937). Computability and λ-definability. *The Journal of Symbolic Logic*, *2*(4), 153–163.
Twardy, C., Gardner, S., & Dowe, D. L. (2005). Empirical data sets are algorithmically compressible: Reply to McAllister? *Studies in History and Philosophy of Science Part A*, *36*(2), 391–402.
Von Neumann, J., & Burks, A. W. (1966). Theory of self-reproducing automata. *IEEE Transactions on Neural Networks*, *5*(1), 3–14.
Walker, S. I., Cisneros, L., & Davies, P. C. W. (2012). Evolutionary transitions and top-down causation. *Proceedings of Artificial Life XIII*, 283–290.
Welsh, D. J. (1993). The complexity of knots. *Annals of Discrete Mathematics*, *55*, 159–171.
Wolfram, S. (2002). *A New Kind of Science* (Vol. 5). Champaign: Wolfram Media.
Zuse, K. (1969). Rechnender Raum (calculating space). http://philpapers.org/rec/ZUSRR