

Incomputability, Emergence and the Turing Universe

S. Barry Cooper*
University of Leeds
Leeds LS2 9JT, U.K.
pmt6sbc@leeds.ac.uk

Abstract

Amongst the huge literature concerning emergence, reductionism and mechanism, there is a role for analysis of the underlying mathematical constraints. Much of the speculation, confusion, controversy and descriptive verbiage might be clarified via suitable modelling and theory. The key ingredients we bring to this project are the mathematical notions of definability and invariance, a computability theoretic framework in a real-world context, and within that, the modelling of basic causal environments via Turing's 1939 notion of interactive computation over a structure described in terms of reals. Useful outcomes are: a refinement of what one understands to be a causal relationship, including non-mechanistic, irreversible causal relationships; an appreciation of how the mathematically simple origins of incomputability in definable hierarchies are materialised in the real world; and an understanding of the powerful explanatory role of current computability theoretic developments.

The theme of this article concerns the way in which mathematics can structure everyday discussions around a range of important issues — and can also reinforce intuitions about theoretical links between different aspects of the real world. This fits with the widespread sense of excitement and expectation felt in many fields — and of a corresponding confusion — and of a tension characteristic of a Kuhnian paradigm shift. What we have below can be seen as tentative steps towards the sort of mathematical modelling needed for such a shift to be completed.

In section 1, we outline the decisive role mathematics played in the birth of modern science; and how, more recently, it has helped us towards a better understanding of the nature and limitations of the scientific enterprise. In section 2, we review how the mathematics brings out inherent contradictions in the Laplacian model of scientific activity. And we look at some of the approaches to dealing

*Research supported by a Royal Society International Joint Project Grant, and by EPSRC Research Grant No. EP/G000212, *The computational structure of partial information: Definability in the local structure of the enumeration degrees.*

with these contradictions. All this leads us back in Section 3 to a closer examination of those aspects of the real world which most obviously test the Laplacian model. In particular, we take a close look at the phenomenon of emergence, and learn from attempts to extract the mathematical content of emergent phenomena. Most important here is the exploration of the close relationship between emergence, definability, and invariance.

Section 4 involves a step back from placing too much explanatory burden on emergence and its mathematics. The need for this becomes particularly clear from an excursion into the philosophy of mind, and from some complementary input from neuroscience. In Section 5, we finally introduce and exercise our mathematical model, and in Section 6, give it what we call a ‘physics road test’.

1 The Laplacian Model Becomes More of a Model

Newton’s successful prediction of planetary motions assembled the important ingredients that have become the features of scientific achievement over more than three-hundred years. To the powerful combination of theoretical speculation and real-world observational data, he added the computational facilitation of mathematics. As Michael White describes [56, p. 93] in *Isaac Newton - The Last Sorcerer*:

If the mathematics had not been developed during the 1660s, Newton’s intuitive grasp of the nature of planetary motion would have remained little more than a good idea. Without his in-depth knowledge of alchemy (which he practised during the 1670s and ’80s), he would almost certainly never have expanded the limited notion of planetary motion as he saw it in 1665/6 into the grand concepts of universal gravitation, of attraction and repulsion, and of action at a distance. Finally, if the experimental evidence had not been gathered, then Newton’s theories, even if substantiated by mathematics, would not have carried the weight they did in his *Principia*, nor would they have so readily inspired the practical application of mechanics and the laws of motion which led, a century later, to the Industrial Revolution.

And the essential underlying product of this coming together was the emergence into the light of day — the conscious recognition — of the computational content of the world and its amenability to capture in mathematical predictions.

Looking more closely at the nature of Newton’s scientific revolution, one sees how computable prediction became part of the subsequent scientific benchmark. Going back to Aristotle and before, observation and speculation had had a close relationship. But the modern empiricism associated with Bacon, and Galileo before him, further emphasised the role of data, and of measurement, with its mathematical focus on real numbers. While Bacon’s view of the inductive establishment of form in nature tied theory and observation even closer: Quoting from Francis Bacon’s *Novum Organum* [3, p. 50]:

There are and can be only two ways of searching into and discovering truth. The one flies from the senses and particulars to the most general

axioms, and from these principles, the truth of which it takes for settled and immovable, proceeds to judgment and to the discovery of middle axioms. And this way is now in fashion. The other derives axioms from the senses and particulars, rising by a gradual and unbroken ascent, so that it arrives at the most general axioms at last. This is the true way, but as yet untried.

It was in this context that Newton's work laid the basis for a model of scientific practice and theory which was to fit well with the Baconian agenda, and set constraints on science which, in retrospect, would be impossible to respect in the longer term. Twentieth century science would both expose a glaring philosophical gap in the Newtonian picture — it is a *background dependent* theory, which gives no explanation of the structure of space-time it incorporated — and demand new kinds of theory from which computable predictions would be harder to extract and verify. Relativity and quantum theory are successful theories even by Newtonian standards, allowing the extraction of computable content of a high order of predictive usefulness. But collectively these have deficiencies necessitating a bizarre range of conjectural proposals, string theoretical ones being best known (of which more later).

With the benefit of our better understanding of the nature of the relationship of theory, computation and observation, one does not need to be a philosopher to recognise the inevitability of this. There is no rigid division between theories concerned with making computable predictions, and ones which are pure metaphysics. Logical analysis of the language in which theories are framed leads us to a detailed analysis of definability in the real world, connecting with well-known hierarchies, and what is known about their computational content — more of this in section 3. The point is that one cannot be surprised that reality needs a richer language than that which delivers purely computable predictions. Or, for that matter, that some mathematics capturing so-called 'causal' relationships might not be reducible to the mechanistic models sought by the good Newtonian.

Anyway, the overarching aim of science, since the time of Galileo, Bacon and Newton, became the extraction of the computational content of the world, at whatever level this this might occur. The process of discovery might not have a simple model, but the outcomes should have computational content with predictive utility, and scientific experiments and mathematical proofs should be reproducible and communicable to fellow scientists. This is what Albert Einstein [18, p. 54] is referring to when he says:

When we say that we understand a group of natural phenomena, we mean that we have found a constructive theory which embraces them.

Why not apply this approach to science itself? Just as Quine, Hilbert, Gödel and others provided us with a model of mathematical proof, and a better understanding of the constraints on the working mathematician, can one similarly model science and its deliverables? In a sense Laplace provide scientists with an aspirational model with his 'predictive demon' [29]:

Given for one instant an intelligence which could comprehend all the forces by which nature is animated and the respective situations of the beings who compose it – an intelligence sufficiently vast to submit these data to analysis – it would embrace in the same formula the movements of the greatest bodies and those of the lightest atom; for it, nothing would be uncertain and the future, as the past, would be present to its eyes.

And over the centuries many have duly internalised this model in a relatively simple form. The aim for them was to emulate the Newtonian successes on which Laplace’s conception was based in ever broader contexts. For Europeans, the late nineteenth century expansionism, such as the ‘scramble for Africa’, gave an appropriate social backdrop to the final throes of this ‘onwards and upwards’ view of science. In mathematics, Hilbert looked for a mathematical counterpart of the predictive demon. Here is the celebrated declaration from his opening address to the Society of German Scientists and Physicians, in Königsberg, September 1930:

For the mathematician there is no Ignorabimus, and, in my opinion, not at all for natural science either ... The true reason why [no one] has succeeded in finding an unsolvable problem is, in my opinion, that there is no unsolvable problem. In contrast to the foolish Ignorabimus, our credo avers:

We must know,

We shall know.

Of course, as described in John Dawson’s biography of Gödel, on the same day in another part of the same city, Gödel was announcing his incompleteness theorem which was to do severe damage to the programme of Hilbert, and usher in a new world of unsolvability.

The so-called Laplacian model was not a model in the sense that the Hilbertian model of mathematical proof was, and so was open to various interpretations. One had to wait until the 1930s for something more mathematical and vaguely relevant.

In 1936 Turing’s machines appeared [50]. This was not the first model of computability, but the one closest to the mechanistic spirit of Newton’s science, and certainly the one which is reputed to have persuaded Gödel that it did achieve its modelling aim. In the first instance, the Turing machine gave a model of computability of functions over the natural numbers. But given the existence of simple codings, it essentially provided a model of algorithmic natural processes within structures which are countably presented. The mathematics of this needs qualifying, but the wide applicability of the model is generally recognised.

But the Turing’s coding techniques for presenting machines gave a Universal Turing Machine — and with this came via the simplest of additions to the language used to describe machines — incomputable objects. Our model of computability arrived, like Sinbad the Sailor bearing the Old Man of the Sea, with a mathematically simple avatar of incomputability on its back. The Universal Turing Machine now has a secure place in the history of the computer — see Davis’ [14] *The Universal Computer: The Road from Leibniz to Turing*. In contrast, incomputability

is an irrelevance to most people beyond the confines of mathematics, and to many of those within. Teuscher's [49] comprehensive collection *Alan Turing: Life and legacy of a great thinker* contains not one article on the mathematical theory of Turing incomputability.

2 Some Uncomfortable Consequences

Since 1936 there has grown up a rich theory of incomputability, complete with hierarchies, fine structure theory, and an analysis of incomputable objects very close to being computable. The latter include *computably enumerable* sets, which have roughly the same relationship to computable sets that computably simulable events in the real world have to ones in which can be computably predicted. There are other kinds of sets which while not being computable, have approximations with computable characteristics, such as the Δ_2^0 sets of the arithmetical hierarchy, which have computable approximations to its members in which finitely many mistakes are allowed before the approximation settles down. Such sets will be dear to the hearts of those, such as Turing, who recognised the limitations of monotonic reasoning — here is Turing talking to the London Mathematical Society on February 20, 1947 (quoted by Andrew Hodges [26, p.361]):

... if a machine is expected to be infallible, it cannot also be intelligent.
There are several theorems which say almost exactly that.

Back in the real world, there was a huge investment in the Laplacian model. And any evidence to the contrary was seen more as a discipline problem than glimpse of a new world; a challenge to be soberly put down with reductionist authority. At times this was timely, such as David Deutsch's influential 1985 Royal Society paper [16] bringing the standard model of quantum computation within the Turing fold. More generally, Martin Davis writing [15] on *The myth of hypercomputation* has argued that:

The great success of modern computers as all-purpose algorithm-executing engines embodying Turing's universal computer in physical form, makes it extremely plausible that the abstract theory of computability gives the correct answer to the question 'What is a computation?', and, by itself, makes the existence of any more general form of computation extremely doubtful.

This should be read as a response to what Davis sees as the inflated hypercomputationalist claims of Jack Copeland and others. Copeland coined the term 'hypercomputation' to describe what an oracle Turing machine might perform. In his article [12] on *Turing's O-Machines, Penrose, Searle, and the Brain*, Copeland explains what oracle machines are capable of:

Let *first-order* O-machines be those whose (only) oracle returns the values of Turing's halting function $H(x, y)$... Similarly, the second-order O-machines are those that possess an oracle which can say whether or

not any given first-order O-machine eventually halts if set in motion with such-and-such input; and so on for third-order, and in general α -order . . .

It is natural to think of the functions, or problems, that are solvable by a first-order oracle machine as being *harder* than those solvable by Turing machine, and those solvable by second-order oracle machine as being harder still, and so forth.

It is the ‘might be’ that so annoys Davis. It is only ‘natural’ in a real-world sense if one can tell us where these oracles are coming from, otherwise there is no evidence such machines have any physical existence. In his article, Copeland has already by-passed this question:

Speculation as to whether there may actually be physical processes that cannot be simulated by Turing machine stretches back over at least four decades (for example Da Costa and Doria 1991; Doyle 1982; Geroch and Hartle 1986; Komar 1964; Kreisel 1967, 1974; Penrose 1989, 1994; Pour-El 1974; Pour-El and Richards 1979, 1981; Scarpellini 1963; Stannett 1990; Vergis et al 1986). If such processes do exist then perhaps future engineers will use them to implement the non-classical part of some O-machine.

Of course, Copeland and Davis are applying the perspectives of different disciplines, and neither managing to say very much new relating to the nature of physical computation. Of course, the more speculative proposals for computational models transcending the so-called ‘Turing barrier’, some of which Davis discusses in his paper, are a mixed bag. The impression one gets from the debate is that one still needs to understand more about how the real world computes.

Despite huge advances in our computational capabilities, there persist problems of predictability in the real world – at the quantum level, in the relationship between emergence and chaos, regarding relativistic phenomena (see István Németi and Hajnal Andréka [34]), and, of course, with mental phenomena. And increasingly the computational capabilities of the physical are seen as relevant to the computing machines we build. There is renewed interest in analog and hybrid computing machines leading Jan van Leeuwen and Jiri Wiedermann [55] to consider that:

. . . the classical Turing paradigm may no longer be fully appropriate to capture all features of present-day computing.

Despite his 1985 paper [16] mentioned earlier, Deutsch did not show that quantum computation *cannot* transcend the Turing barrier, just that the current model does not do it. As Andrew Hodges remarks in *What would Alan Turing have done after 1954?* [27]:

Von Neumann’s axioms distinguished the **U** (unitary evolution) and **R** (reduction) rules of quantum mechanics. Now, quantum computing so far (in the work of Feynman, Deutsch, Shor, etc) is based on the **U** process and so computable. It has not made serious use of the **R** process:

the unpredictable element that comes in with reduction, measurement, or collapse of the wave function.

Although measurement does play a role in quantum computation, and the probabilities of a particular outcome of a measurement are computable, there are still aspects of the physics which are not used which are thought to be in some sense ‘random’. Recently, under reasonable assumptions about the basic character of quantum randomness, Calude and Svozil [6] have shown that quantum uncertainty does entail incomputability — though just *how* random quantum randomness really is is still very much open to question. It may be that despite all the assumptions of physicists, nature is full of incomputability, but does not exhibit any significant level of mathematical randomness. The challenge is to integrate quantum phenomena into a general picture of physical computation. This might not entail a useable unified theory of physics, but would hopefully present quantum uncertainty as a feature of mathematical constraints operative throughout science.

There are clearly features of the classical world which challenge the Davis disciplinary regime. As observed by Copeland, some of the earliest (and deepest) thinking on the question of physical incomputability comes from another distinguished source — back in 1970 the mathematician Georg Kreisel was proposing a collision problem related to the 3-body problem, which might result in “an analog computation of a non-recursive function”.

One can find detailed accounts of Kreisel’s thinking on extensions to the Church-Turing thesis in the section of Odifreddi’s first volume of *Classical Recursion Theory* [35], and in Odifreddi’s article on the topic in his edited volume *Around and About Georg Kreisel*.

Another challenge arises from the growth of chaos theory, dealing with the generation of informational complexity via very simple rules. Features of chaotic situations include the iteration of simple rules, nonlinearity, and the sort of sensitivity to initial conditions that Edward Lorenz [31] observed in the development of weather systems. Another feature is the *emergence* of systemic formations, such as the Lorenz attractor, or the strange attractor discovered by Robert Shaw [44, 45] in studying the ostensibly very simple chaos of a dripping tap — by varying the flow to the dripping tap, unpredictable irregularities of intervals between drips were observed, while appropriate plotting of the unpredictable data revealed an interesting 3-dimensional strange attractor.

The special interest of chaos arises not so much from its undeniable novelty of computational character — there are all sorts of explanations of the apparent indeterminacy of outcomes, not usually enlisting incomputability — but from the availability of informative mathematical analogues. It is the link between such structures in nature, and mathematical objects, such as the Mandelbrot and Julia sets, which presents an opportunity of getting closer to a mathematical characterisation of what is happening. At the same time, the mathematical interest and approachability of fractals, with their grounding in the iteration of simple rules paralleling those in nature, makes their computability-theoretic character accessible to serious investigation.

The Mandelbrot set has attracted particular attention from high-profile scientists such as Roger Penrose and Stephen Smale. Its popular appeal is matched by its mathematical interest. As Penrose puts it in *The Emperor's New Mind* [?]:

Now we witnessed . . . a certain extraordinarily complicated looking set, namely the Mandelbrot set. Although the rules which provide its definition are surprisingly simple, the set itself exhibits an endless variety of highly elaborate structures.

And it is not just the observed patterns which are hard to predict. The computability of the actual point-set is still very much an open problem (despite its incomputability in the Blum-Shub-Smale model [4] of real computation – see Hertling's review article [25]).

We saw earlier that it is just the addition of a quantifier to the language used to describe a Turing machine which opens the door to the emergence of incomputability. Looking at the definition of the Mandelbrot set in terms of limiting behaviour of applications of the polynomial rule $z \rightarrow z^2 + c$, we immediately get a two-quantifier form for the set. But a little extra work gives the complement of the Mandelbrot set using just one existential quantifier. We need to pursue further the general phenomenon of emergence observed in the above examples, and to relate it to the complexity of language needed to describe it.

3 What Is Emergence? – Definability, Nonlocality

Emergence is a much over-worked concept. For example, its perceived potential for undermining determinism makes it specially appealing to those trying to create room for religion in a scientific world. Here is Stuart Kauffman making some very grand claims in his recent book [28, p. 281] *Reinventing the Sacred: A New View of Science, Reason and Religion*:

We are beyond reductionism: life, agency, meaning, value, and even consciousness and morality almost certainly arose naturally, and the evolution of the biosphere, economy, and human culture are stunningly creative often in ways that cannot be foretold, indeed in ways that appear to be partially lawless. The latter challenge to current science is radical. It runs starkly counter to almost four hundred years of belief that natural laws will be sufficient to explain what is real anywhere in the universe, a view I have called the Galilean spell. The new view of emergence and ceaseless creativity partially beyond natural law is a truly new scientific worldview in which science itself has limits.

Without saying Kauffman is wrong — his world-view has a lot of appeal — one cannot help but be nervous at such ambitious conclusions based on such a modest grasp of what emergence really is. This is Ronald Arkin's [2, p. 105] comment:

Emergence is often invoked in an almost mystical sense regarding the capabilities of behavior-based systems. Emergent behavior implies a

holistic capability where the sum is considerably greater than its parts. It is true that what occurs in a behavior-based system is often a surprise to the system's designer, but does the surprise come because of a shortcoming of the analysis of the constituent behavioral building blocks and their coordination, or because of something else?

Ronald, Sipper and Capcarrère [43] have devised a 'Test for Convergence' which usefully clarifies what we expect of an emergent phenomenon. It follows the example of the Turing Test for intelligence machinery in being observer dependent, which solves some problems even if it is not so obviously appropriate. The three criteria they list are (slightly paraphrased):

- 1) **Design:** The system has been constructed by the designer, by describing local elementary interactions between components (e.g., artificial creatures and elements of the environment) in a language \mathcal{L}_1 .
- 2) **Observation:** The observer is fully aware of the design, but describes global behaviors and properties of the running system, over a period of time, using a language \mathcal{L}_2 .
- 3) **Surprise:** The language of design \mathcal{L}_1 and the language of observation \mathcal{L}_2 are distinct, and the causal link between the elementary interactions programmed in \mathcal{L}_1 and the behaviors observed in \mathcal{L}_2 is non-obvious to the observer — who therefore experiences surprise. In other words, there is a cognitive dissonance between the observer's mental image of the system's design stated in \mathcal{L}_1 and his contemporaneous observation of the system's behavior stated in \mathcal{L}_2 .

A useful part of the test is the bringing out of the qualitative difference between the 'design' and the observed 'global behaviours' via the distinction between the languages \mathcal{L}_1 and \mathcal{L}_2 used to describe them.

On the other hand, the parallel with the observer-based Turing Test is weak, with condition 3) of the Emergence Test lacking robustness; how do we evaluate the origin of the observer's 'surprise'? For the Turing Test, the observer's inability to discriminate between the intelligence of machine and human comes with far more weight and relevance. We need to look more closely at the computational content of emergence, with the aim of extracting a clearer "surprise" criterion.

The view we want to pursue is that emergent phenomena not only yield up descriptions, using different language to that used in describing the underlying design; they are actually determined, constrained, *captured* by that which is describable in terms of the basic causal structure.

The intuition that entities exist because of, and according to, mathematical laws, is not new, of course. One can detect it in the words of Leibniz [30] from 1714 in the *The Monadology*, section 32:

...there can be found no fact that is true or existent, or any true proposition, without there being a sufficient reason for its being so and not otherwise, although we cannot know these reasons in most cases.

So natural phenomena not only generate descriptions, but arise and derive from them. So connecting with a useful abstraction, that of mathematical definability — or, more generally, invariance (under the automorphisms of the appropriate structure). And this gives precision to our experience of emergence as a potentially non-algorithmic determinant of events. On the one hand one can attempt to frame criteria for emergence in terms of the complexity of the language used to describe it, and one can also use the known associations between informational and computational complexity to constrain the computability-theoretic character of physical phenomena.

For instance, taking this approach, one might identify the halting set of the Universal Turing Machine as an emergent phenomenon; although it does not have the visual immediacy of the Mandelbrot set, it is incomputable, and that in itself qualifies it as a sufficiently surprising global attribute.

What one would expect from this very clear connection between the underlying basic causal structure (the ‘design’) and the emergent phenomenon would be a certain level of robustness of the emergence. What one is suggesting, via the association with mathematical definability, is a direct causal relationship between ‘design’ and emergent phenomenon — and one which is unlike the usual fundamental laws of nature, in that it is more global in respect of the causes it works with — and potentially, with respect to the effects. This is not so surprising from the point of view of carefully delineated experimental contexts, such as that presented by Robert Shaw’s dripping tap. More so with the higher-order emergence being called up by Stuart Kauffman. If one goes back to Samuel Alexander’s [1] magnum opus from the nineteen-twenties (another theologically inclined writer, one of the British emergentists described by Brian McLaughlin [33]) one finds the mystery of connection an integral part of the argument.

Anyway, it is just this expected robustness that Martin Nowak identified (as Director of the Program for Evolutionary Dynamics at Harvard University) in emergent aspects of evolution. This is from the interesting collection of papers from leading scientists brought together in John Brockman’s *What We Believe But Cannot Prove* [5]:

I believe the following aspects of evolution to be true, without knowing how to turn them into (respectable) research topics.

Important steps in evolution are robust. Multicellularity evolved at least ten times. There are several independent origins of eusociality. There were a number of lineages leading from primates to humans. If our ancestors had not evolved language, somebody else would have.

4 Is That All There Is? – Turing and the Human Brain

We have kept back our third challenge to Davis Discipline until we were clearer on what we wanted to summon up from it; we are now ready for the complexities of the human mind as case study. It comes with a number of strengths:

- The human mind is very *familiar*, at least to the more self-aware. Experience

of its workings is easily got through solving everyday problems, and observing others.

- And the *mechanics of the brain are well-documented*.
- The mind *does not feel, or appear to compute, like a Turing machine* — given the role of creativity, consciousness, intuition.
- The case study is *relevant* — given the importance of copying how humans think for achieving AI etc . . . and the intuition that a physical brain reflects processes in the wider universe, so can help with the modelling new aspects of physical computation.

So how do the mind and emergence match up? The surprise criterion is certainly there. Here is a well-known example from Jacques Hadamard's celebrated 1945 study [24] of *The Psychology of Invention in the Mathematical Field*, based on conversation with Henri Poincaré:

At first Poincaré attacked [a problem] vainly for a fortnight, attempting to prove there could not be any such function . . . [quoting Poincaré]:
 “Having reached Coutances, we entered an omnibus to go some place or other. At the moment when I put my foot on the step, the idea came to me, without anything in my former thoughts seeming to have paved the way for it . . . I did not verify the idea . . . I went on with a conversation already commenced, but I felt a perfect certainty. On my return to Caen, for conscience sake, I verified the result at my leisure.”

Apart from the surprise element, the unexpected arrival of a crucial idea by some unconscious process, there is another important aspect of this story — the *robustness* of the surprise solution to the problem that Poincaré had been stuck on. He could feel enough confidence in his ability to recreate the solution at some later time to be able to carry on a completely unrelated conversation. The idea, it appears, had a memetic quality consistent with the existence of a *representation* of the solution, such as one might expect from an association of emergence with definability.

So much for part 3) of the Emergence Test. But what about the design? One needs to bridge the gap between higher mental functionality and . . . what algorithmic context? One might hope to derive this from existing models of neural functionality. But this is more difficult than one might expect. According to Rodney Brooks in *Nature* in 2001:

. . . neither AI nor Alife has produced artifacts that could be confused with a living organism for more than an instant.

Another creative participant in the field of AI, Daniel Hillis, Chief Technology Officer of Applied Minds, Inc. (and ex-Vice President, Research and Development at Walt Disney Imagineering), was quoted in April 2001 as doubting whether design was even sufficient for the building of intelligent machines. Perhaps getting intelligent machines themselves would be via emergence:

I used to think we'd do it by engineering. Now I believe we'll evolve them. We're likely to make thinking machines before we understand how the mind works, which is kind of backwards.

This is not to say that paradigm-stretching features of connectionist models of computation are lacking. As Paul Smolensky (recipient of the 2005 David E. Rumelhart Prize) wrote [46] in 1988:

There is a reasonable chance that connectionist models will lead to the development of new somewhat-general-purpose self-programming, massively parallel analog computers, and a new theory of analog parallel computation: they may possibly even challenge the strong construal of Church's Thesis as the claim that the class of well-defined computations is exhausted by those of Turing machines.

And it is certainly true that connectionist models have come a long way since Turing's [1948] discussion of 'unorganised machines', and McCulloch and Pitts' [1943] early paper on neural nets.

But for Steven Pinker [36] "...neural networks alone cannot do the job". And focussing on our elusive higher functionality, he points to a "kind of mental fecundity called recursion":

We humans can take an entire proposition and give it a role in some larger proposition. Then we can take the larger proposition and embed it in a still-larger one. Not only did the baby eat the slug, but the father saw the baby eat the slug, and I wonder whether the father saw the baby eat the slug, the father knows that I wonder whether he saw the baby eat the slug, and I can guess that the father knows that I wonder whether he saw the baby eat the slug, and so on.

Less amusingly, but bringing out even more clearly the role of *recycled* emergence, the neuroscientist Antonio Damasio makes a similar point. Here is his nice description of the hierarchical development of a particular instance of consciousness within the brain (or 'organism'), interacting with some external 'object' [13]:

...both organism and object are mapped as neural patterns, in first-order maps; all of these neural patterns can become images ... The sensorimotor maps pertaining to the object cause changes in the maps pertaining to the organism ... [These] changes ... can be re-represented in yet other maps (second-order maps) which thus represent the relationship of object and organism ... The neural patterns transiently formed in second-order maps can become mental images, no less so than the neural patterns in first-order maps.

The picture is one of *re-representation* of neural patterns formed across some region of the brain, in such a way that they can have a *computational relevance in forming new patterns*. There is a key conception of computational loops incorporating, in a controlled way, these 'second-order' aspects of the computation itself. The exact

mechanism for the creation and recycling of emergent outputs is not completely clear. But the actuality of this is substantiated via our mathematical model of the definability of emergent phenomena, whereby new entities are created and defined along with a role in the original structure. It is worth noting in this context that the basic logic underlying natural language, upon which descriptions/definitions are based, does not have an irreducible, and mysterious special status in our scientific ontology; it arises from the most basic of material algorithms, ones which appear unavoidable in any viable causal context, and derive their position in human discourse via the close relationship (for us) between matter and data.

We are now ready to try and make more explicit our basic computational model. We have talked a lot about the roles of definability and invariance, without placing these notions in a specific setting. Key ingredients to be sought in such a model are those we have been talking about: imaging, parallelism, interconnectivity, and a counterpart to the second-order recursions pointed to above. And the computational content familiar from the material universe should appear explicitly in the model.

Connectionist models are strong on parallelism, interconnectivity, imaging, and can even accommodate recursions — but not in re-integrating the sort of recursions Pinker is describing into the computational process. And echoing Danny Hillis' comment above about the role of design, one may have to look for a model of the fundamental computational structure of the world, without being able to fully model the functionality. Such a model may not provide the design of an artificial brain, but it may help us understand the obstacles to doing that.

5 The Extended Turing Model

The theme of computation versus description runs through most of Alan Turing's work, and never more explicitly than in his long, hard-to-read, and immensely influential 1939 article [51]. An important thread, begun in this paper and running through much of the subsequent history of computability theory, concerns how the computational content of descriptions can be captured hierarchically — but in unpredictable ways.

Turing's approach is largely proof-theoretic, growing out of his interest in Gödel's incompleteness theorem, and what it tells us about the extent of the boundaries of the computable world. Turing shows that despite Gödel's proof [21] that no consistent first-order theory captures arithmetic, we can hierarchically transcend this barrier, in a quite constructive way — one just iterates the Gödel argument, computably generating new unprovable theorems which are then used to enlarge the theory. One uses computable ordinal notations to iterate this process into the transfinite in a constructive way, thus giving the appearance of computably transcending Gödel's theorem. But a little thought reveals the snag — identifying the route to a new theorem involves using an incomputable oracle, so we avoid the reductionist paradox.

This is how Turing explains what he had done:

Mathematical reasoning may be regarded ... as the exercise of a combination of ... intuition and ingenuity ... In pre-Gödel times it was thought by some that all the intuitive judgements of mathematics could be replaced by a finite number of ... rules. The necessity for intuition would then be entirely eliminated. In our discussions, however, we have gone to the opposite extreme and eliminated not intuition but ingenuity, and this in spite of the fact that our aim has been in much the same direction.

So here we have an explanation of why written proofs do not tell us how the proof was discovered. The ‘intuition’ involved was needed to identify the path to a proof – in the way Poincaré needed it – but having done that by some incomputable process, one immediately has a purely algorithmic demonstration (that is the proof) of why the theorem is true. The result of this process is that one delivers an emergent result into a developing body of mathematics which has a deceptively algorithmic structural appearance.

Having tried unsuccessfully to ‘compute the incomputable’, Turing introduced a model of natural causality between real data, which could be incomputable. The model – now called an *oracle Turing machine* – was essentially just a Turing machine which could ask questions of an external ‘oracle’ (usually a set of natural numbers). The number of questions during a particular computation was finite, of course. The result was that instead of getting computable real numbers via the collating of computational outputs of a machine, one now got real numbers computable *relative* to an oracle. Considering the oracles to be inputs, a given machine might capture a particular computable function over the reals, notated as a *Turing functional* from reals to reals. Given the natural form of this quite general notion, it turns out to be sufficient to capture most of the functions one extracts from basic laws of science. For instance, one can easily represent the progress of two given point masses (whose relative states at a given time are represented as a real) according to Newtonian dynamics via a Turing functional. This is not surprising, since such simple basic transformations are routinely captured via functions over the reals which can be computed up to any practicable level of approximation by a real-world computer. Given more point masses, one can still describe the motion *in terms* of that functional, but this does not allow one to extract a new Turing functional to completely express the new causal relationship. Here we have again basic computability leading very quickly via descriptions to a situation with computational content, but not necessarily computable.

But the bottom line is that in 1939 Turing’s oracle machines appeared, and that these provided a model of computable content of structures, based on *partial computable* (p.c.) functionals over the reals. This model – the *Turing universe* – was capable of capturing basic computable causal structure in the real world, with the expectation, based on experience, that any incomputable causality would be definable in some natural way from this basic structure.

This extended model of Turing’s had a very interesting history. Some of this is described in *The Incomputable Alan Turing* [9]. Around 1948 Emil Post [37] tidied

up the model by gathering together computably equivalent reals into equivalence classes called *degrees of unsolvability*, with an ordering induced by that of relative Turing computability. This gave a classification of reals in terms of their relative computability, so giving an informational landscape with a rich structure.

Back in the real world again, we know that we can often describe global relations in terms of well-understood local structure – so capturing the emergence of large-scale formations. We can now formalise this mathematically in terms of definability over structure based on Turing functionals, insofar as we understand the basic causal structure. Again, if one is concerned about the language dependency of the notion of definability – language is a human construct, and not obviously applicable to the way the universe ‘defines’ its large scale structure and laws – then one can express things in terms of invariance under automorphisms.

This brings us to *Hartley Rogers’ Programme* which (see [41]) addresses the:

Fundamental problem: *Characterise the Turing invariant relations.*

The intuition is that these relations are key to pinning down how basic laws and entities emerge as mathematical constraints on causal structure.

At one time, it was thought that the structural pathology exhibited by the Turing universe, and the disproportionate technical difficulty of proofs in the area, was evidence of mathematical ugliness, disqualifying the field from serious attention of non-specialists. It is now understood that the richness of Turing structure discovered so far provides the raw material for non-trivially defining a multitude of relations. And that the complexity and pathology of the structure is only what one would expect of something aiming to model global aspects of the real world.

6 And a Physics Road Test

The Turing model has considerable explanatory power. In [10] we apply this to the problem of clarifying the connection between the mental and the physical. Here, we focus on very different problems affecting the standard model of particle physics. Concern about the current state of physics is comes from a number of sources. Peter Woit [57], in the introduction to his 2006 book *Not Even Wrong - The Failure of String Theory and the Continuing Challenge to Unify the Laws of Physics*, describes the situation so:

By 1973, physicists had in place what was to become a fantastically successful theory of fundamental particles and their interactions, a theory that was soon to acquire the name of the ‘standard model’. Since that time, the overwhelming triumph of the standard model has been matched by a similarly overwhelming failure to find any way to make further progress on fundamental questions.

The success he refers to is in terms of practical prediction. The failure in relation to fundamental questions relates to lack of recent progress – the problems themselves have been around in some form or other for a long time. Einstein himself says in his *Autobiographical Notes* [?, p. 63]:

...I would like to state a theorem which at present can not be based upon anything more than upon a faith in the simplicity, i.e. intelligibility, of nature ... nature is so constituted that it is possible logically to lay down such strongly determined laws that within these laws only rationally completely determined constants occur (not constants, therefore, whose numerical value could be changed without destroying the theory) ...

These may not be quite the same undetermined constants that Peter Woit is pointing to (there are more of them now):

One way of thinking about what is unsatisfactory about the standard model is that it leaves seventeen non-trivial numbers still to be explained, ...

But the substance of the complaint is the same; one should not need to adjust elements of the standard model in a seemingly arbitrary way just to get the right answers delivered. The theory should give a complete explanation of the values of constants etc.

This is what it was hoped string theory would do. In a sense string theory was a departure from the Baconian paradigm, which Einstein himself had initiated, and demonstrated the power of. But things have not worked out well, and as the family of string theories and their offshoots expands, along with the arbitrary choices needed, the argument is that string theory is “the only game in town”. One-time string theorist Daniel Friedan [19] is dismissive:

The longstanding crisis of string theory is its complete failure to explain or predict any large distance physics ... String theory is incapable of determining the dimension, geometry, particle spectrum and coupling constants of macroscopic spacetime ... The reliability of string theory cannot be evaluated, much less established. String theory has no credibility as a candidate theory of physics.

Lee Smolin’s 2006 book [47] on *The Trouble With Physics* is another source of dissent. In it he lists “Five Great Problems in Theoretical Physics”. What is relevant for us is that each one can be framed as a *problem of definability*:

1. Combine general relativity and quantum theory into a single theory that can claim to be the complete theory of nature.
2. Resolve the problems in the foundations of quantum mechanics
3. The unification of particles and forces problem: Determine whether or not the various particles and forces can be unified in a theory that explains them all as manifestations of a single, fundamental entity.
4. Explain how the values of the free constants in the standard model of physics are chosen in nature.

5. Explain dark matter and dark energy. Or, if they don't exist, determine how and why gravity is modified on large scales.

An indication of the widespread concern about such problems was the 2005 statement from no less than David Gross (co-discoverer of the asymptotic freedom affecting the strong nuclear force), quoted in the Dec. 10, 2005, *New Scientist*, under the heading *Nobel Laureate Admits String Theory Is In Trouble*:

The state of physics today is like it was when we were mystified by radioactivity ... They were missing something absolutely fundamental. We are missing perhaps something as profound as they were back then.

So what is it that is 'absolutely fundamental' that is missing? It is worth noting that Smolin's thinking is consistent with our own emphasis on the modelling of basic causal structure. He proclaims that "causality is fundamental". And while pointing to early champions of the role of causality, such as Roger Penrose, Rafael Sorkin, Fay Dowker, and Fotini Markopoulou, he says [47, p. 241]:

It is not only the case that the spacetime geometry determines what the causal relations are. This can be turned around: Causal relations can determine the spacetime geometry ... Its easy to talk about space or spacetime emerging from something more fundamental, but those who have tried to develop the idea have found it difficult to realize in practice ... We now believe they failed because they ignored the role that causality plays in spacetime. These days, many of us working on quantum gravity believe that causality itself is fundamental – and is thus meaningful even at a level where the notion of space has disappeared.

And we even detect here an implicit searching for a structure in which the definable set of relations on it is rich enough to take in something corresponding to the spacetime geometry we observe.

Of course, from the point of view of Smolin's Great Problem number 2, one might also benefit from a *failure of definability* corresponding to the quantum ambiguity we encounter, and which disappears with the collapse of the wave function during a measurement. Earlier, having noted that quantum uncertainty presented a particularly strong challenge to Davis' reductionist programme, we went on to focus almost entirely on emergence. It is now time to bring quantum phenomena back into the picture. According to our picture, emergence coincides with an assertion of definability in some underlying causal structure. The complexity of the definition gives rise to a related level of surprise and unpredictability.

What we have at the quantum level is something rather different. What is being defined (or not being defined, as the case may be) is *attributes of the basic design*. Following Leibniz, lacking a definition of aspects of a given quantum state, the state has to exhibit whatever it is allowed to. But an intervention involving a measurement or whatever may enrich the context sufficiently to remove these various possibilities, and leave us with a well-defined classical reality. And the

process involves a mathematically enforced non-locality, quite in keeping with what is observed. Anyway, the classical level may not so be so much of a surprise to those of us who spend all our time there, but it is nevertheless emergent. What is surprising to us is that there is a level at which not all is unambiguously defined, and the transition between the two. One would also notice that this is a realistic interpretation, achieved without anthropic principles, many-worlds interpretations, or any other level of Max Tegmark's multiverse hierarchy.

Smolin's Great Problem number 1 also raises interesting features. Notice that when we are presented with emergent entities, described in a different language to the underlying design, they may well determine a whole new level of behaviour, complete with their own emergent causal relations. This is a picture familiar which was familiar to the British Emergentists, dealt with in Brian McLaughlin's book [33] mentioned earlier. They used it to explain the irreducibility of the 'special sciences', postulating a hierarchy with physics at the bottom, followed by chemistry, biology, social science etc. The emergence, as our model confirms, is irreversible, imposing the irreducibility of say biology to quantum theory – although the British emergentists experienced their heyday before the great quantum discoveries of the late 1920s, and as described in [33], this was in a sense their undoing.

Now, what would we think of someone who asked for a unified theory of chemistry and biology? It may be that it is equally senseless to be looking for a unified theory of quantum and relativity theory. On the other hand, with the example of the British emergentists who held that the coming together of hydrogen and oxygen to form water was an example of emergence, one can never be quite sure about the extent of application of useful models.

Smolin's Great Problem number 3 is perhaps a little too specific to be obviously within the scope of such a schematic model as we are applying. It may be that there is something basic about the automorphism group of the Turing universe and its corresponding invariant relations which tell us something very relevant about the fundamental structure of the entities making up the universe; we conjectured something of the sort in Haifa back in 1995 (see [7]). On the other hand, the answer may depend on much more specific considerations arising from physics.

Problem 4 is obviously a question of definability. And so may Problem 5 be, involving levels of failure of definability beyond our observational reach.

What we would look for is solutions to a range of fundamental problems, within a radically deconstructed universe:

- Described in terms of reals . . .
- With emergent natural laws based on algorithmic relations between reals
- With emergence described in terms of definability/invariance
- . . . with failures of definability modelling quantum ambiguity
- . . . which gives rise to new levels of algorithmic structure
- . . . and a fragmented scientific enterprise.

What the mathematics can deliver is a causality which is different in nature from that which Newton gave us back at the beginning of the modern scientific era. Alan Guth (the inventor of cosmic inflation) asks in his book [23] *The Inflationary Universe - The Quest for a New Theory of Cosmic Origins*:

If the creation of the universe can be described as a quantum process, we would be left with one deep mystery of existence: What is it that determined the laws of physics?

It is important to bring such questions firmly into the scientific domain.

Bibliography

- [1] S. Alexander. *Space, Time, and Deity*, 1927.
- [2] R.C. Arkin. *Behaviour-Based Robotics*. MIT Press, 1998.
- [3] F. Bacon. *The Works* (J. Spedding, R. L. Ellis, and D. D. Heath, eds.), Vol. IV, London, 1901.
- [4] L. Blum, F. Cucker, M. Shub and S. Smale. *Complexity and Real Computation*. Springer, 1997.
- [5] J. Brockman (ed.). *What We Believe but Cannot Prove: Today's Leading Thinkers on Science in the Age of Certainty*, Harper Perennial, New York, 2006.
- [6] C. S. Calude and K. Svozil. *Quantum randomness and value indefiniteness*. *Advanced Science Letters*, 1 (2008), 165–168.
- [7] S. B. Cooper. *Beyond Gödel's Theorem: Turing nonrigidity revisited*. In *Logic Colloquium '95* (J. A. Makowsky and E. V. Ravve, eds.), Proceedings of the Annual European Summer Meeting of the Association of Symbolic Logic, held in Haifa, Israel, August 9-18, 1995, *Lecture Notes in Logic*, vol. 11, Springer, Berlin, New York, Tokyo, 1998, pp. 44–50.
- [8] S. B. Cooper. *Computability Theory*. Chapman & Hall/CRC, Boca Raton, London, New York, Washington, D.C., 2004.
- [9] S. B. Cooper. *The Incomputable Alan Turing*. In the *Proceedings of Turing 2004: A celebration of his life and achievements*, electronically published (2008) by the British Computer Society:
<http://www.bcs.org/server.php?show=nav.9917>
- [10] S. B. Cooper. *From Descartes to Turing: The Computational Content of Supervenience*. To appear.
- [11] S. B. Cooper and P. Odifreddi. Incomputability in Nature. In S.B. Cooper and S.S. Goncharov, *Computability and Models*. Kluwer Academic/Plenum, New York, Boston, Dordrecht, London, Moscow, 2003, pages 137–160.

- [12] J. Copeland. *Turing's O-Machines, Penrose, Searle, and the Brain*. *Analysis*, **58** (1998), 128–138.
- [13] A.R. Damasio. *The Feeling Of What Happens: Body and Emotion in the Making of Consciousness*, Harcourt Brace, 1999.
- [14] M. Davis. *The Universal Computer: The Road from Leibniz to Turing*. W.W. Norton, New York, London, 2000.
- [15] M. Davis. *The myth of hypercomputation*. In *Alan Turing: Life and legacy of a great thinker* (C. Teuscher, ed.), Springer-Verlag, Berlin, Heidelberg, 2004, pp. 195–212.
- [16] D. Deutsch. *Quantum theory, the Church Turing principle, and the universal quantum computer*. *Proc. Roy. Soc., A* 400 (1985), 97–117.
- [17] A. Einstein. *Out of My Later Years*. Philosophical Library, New York, 1950.
- [18] A. Einstein. *Autobiographical Notes*. In *Albert Einstein: Philosopher-Scientist* (P. Schilpp, ed.), Open Court Publishing, 1969.
- [19] D. Friedan. *A Tentative Theory of Large Distance Physics*. *J. High Energy Phys.* JHEP10(2003)063.
- [20] Yu. L Ershov, S.S. Goncharov, A. Nerode, J.B. Remmel (Editors). *Handbook of Recursive Mathematics*, Volumes 1 and 2. Elsevier, Amsterdam, New York, Oxford, Tokyo, 1998.
- [21] K. Gödel. *Über formalunentscheidbare Sätze der Principia Mathematica und verwandter Systeme I*. *Monatsh. Math. Phys.* **38** (1931), 173–198; translated in *The Undecidable. Basic Papers on Undecidable Propositions, Unsolvability Problems, and Computable Functions* (M. Davis, ed.), Raven Press, Hewlett, New York, 1965.
- [22] E. R. Griffor (Editor). *Handbook of Computability Theory*. Elsevier, Amsterdam, New York, Oxford, Tokyo, 1999.
- [23] Guth A. H. (1997) *The Inflationary Universe – The Quest for a New Theory of Cosmic Origins*. Addison-Wesley, New York, Harlow, England, Tokyo, Paris, Milan.
- [24] J. Hadamard. *The Psychology of Invention in the Mathematical Field*. Princeton Univ. Press, Princeton, 1945.
- [25] P. Hertling. *Is the Mandelbrot set computable?* *Math. Logic Quarterly*, **51** (2005), 5–18.
- [26] A. Hodges. *Alan Turing: The Enigma*. Vintage, London, Melbourne, Johannesburg, 1992.

- [27] A. Hodges. *What would Alan Turing have done after 1954?* In C. Teuscher(ed.), *Alan Turing: Life and legacy of a great thinker*, Springer-Verlag, Berlin, Heidelberg, 2004.
- [28] S. A. Kauffman. *Reinventing the Sacred: A New View of Science, Reason and Religion*. Basic Books, 2008.
- [29] P. S. de Laplace. *Essai Philosophique sur les Probabilités*. English trans. by F. W. Truscott and F. L. Emory, Dover, New York, 1951.
- [30] G. W. Leibniz. *La Monadologie* (1714). English translation by G.M. Ross, 1999.
- [31] E. Lorenz. *Deterministic nonperiodic flow*. *Journal of Atmospheric Sciences*. **20** (1963), 130–141.
- [32] Y. Matiyasevich. *Hilbert's Tenth Problem*. MIT Press, Cambridge, Mass., London, 1993.
- [33] B.P. McLaughlin. *The Rise and Fall of British Emergentism*. In *Emergence or Reduction? – Essays on the Prospects of Nonreductive Physicalism* (A. Beckermann, H. Flohr, J. Kim, eds.), de Gruyter, Berlin, 1992, pp. 49–93. Reprinted in *Emergence: Contemporary Readings in Philosophy and Science* (M. A. Bedau and P. Humphreys, eds.), MIT Press, Cambridge, Mass., London, 2008, pp. 19–59.
- [34] I. Németi, H. Andréka. *Can General Relativistic Computers Break the Turing Barrier?* In *Logical Approaches to Computational Barriers, Second Conference on Computability in Europe, CiE 2006, Swansea, UK, June 30-July 5, 2006, Proceedings* (A. Beckmann, U. Berger, B. Löwe, J. V. Tucker, Eds.). *Lecture Notes in Computer Science 3988*, Springer, 2006, pp. 398–412.
- [35] P. Odifreddi. *Classical Recursion Theory*, Volumes I and II. North-Holland/Elsevier, Amsterdam, New York, Oxford, Tokyo, 1989 and 1999.
- [36] S. Pinker. *How the Mind Works*. W.W. Norton, New York, 1997.
- [37] E. L. Post. *Degrees of recursive unsolvability: preliminary report*. (abstract). *Bull. Amer. Math. Soc.* **54** (1948), 641–642.
- [38] E. L. Post. *Solvability, Provability, Definability: The Collected Works of Emil L. Post* (Martin Davis, Editor). Birkhäuser, Boston, Basel, Berlin, 1994.
- [39] R. Penrose. *The Emperor's New Mind: Concerning Computers, Minds, and the Laws of Physics*. Oxford University Press, Oxford, New York, Melbourne, 1994.
- [40] C. Reid. *Hilbert*. Springer-Verlag, Berlin, Heidelberg, New York, London, Paris, Tokyo, 1986 (paperback edition).

- [41] H. Rogers, Jr. *Some problems of definability in recursive function theory*. In *Sets, Models and Recursion Theory* (J.N. Crossley, ed.), Proceedings of the Summer School in Mathematical Logic and Tenth Logic Colloquium, Leicester, August–September, 1965, North-Holland, Amsterdam, pp. 183–201.
- [42] H. Rogers, Jr. *Theory of Recursive Functions and Effective Computability*. MIT Press, Cambridge, Mass., London, 1967.
- [43] E.M.A. Ronald, M. Sipper and M.S. Capcarrère. *Design, observation, surprise! A test of emergence*. *Artificial Life*, **5** (1999), 225–239. Reprinted in *Emergence: Contemporary Readings in Philosophy and Science* (M. A. Bedau and P. Humphreys, eds.), MIT Press, Cambridge, Mass., London, 2008, pp. 287–304.
- [44] R. Shaw. *Strange attractors, chaotic behaviour, and information flow*. *Z. Naturforsch*, **36A** (1981), 80–112.
- [45] R. Shaw, *The dripping faucet as a model chaotic system*, The Science Frontier Express Series, Aerial Press, Santa Cruz, CA, 1984.
- [46] P. Smolensky. *On the proper treatment of connectionism*. *Behavioral and Brain Sciences* **11** (1988), 1–74.
- [47] L. Smolin. *The Trouble With Physics: The Rise of String Theory, the Fall of a Science and What Comes Next*. Houghton Mifflin, 2006.
- [48] R. I. Soare. *Recursively Enumerable Sets and Degrees*. Springer-Verlag, Berlin, Heidelberg, New York, London, Paris, Tokyo, 1987.
- [49] C. Teuscher(editor). *Alan Turing: Life and legacy of a great thinker*. Springer-Verlag, Berlin, Heidelberg, 2004.
- [50] Turing A. M. (1936) *On computable numbers, with an application to the Entscheidungsproblem*. *Proc. London Math. Soc.* (2) **42** (1936–7), pp. 230–265. Reprinted in A. M. Turing, *Collected Works: Mathematical Logic*, pp. 18–53.
- [51] A.M. Turing. *Systems of logic based on ordinals*. *Proc. London Math. Soc.* (2) **45** (1939), pp. 161–228. Reprinted in A. M. Turing, *Collected Works: Mathematical Logic*, pp. 81–148.
- [52] A. M. Turing. *Intelligent machinery*. National Physical Laboratory Report (1948). In *Machine Intelligence 5* (B. Meltzer and D. Michie, eds.), Edinburgh University Press, Edinburgh, 1969, pp. 3–23. Reprinted in A. M. Turing, *Collected Works: Mechanical Intelligence* (D. C. Ince, ed.), North-Holland, Amsterdam, New York, Oxford, Tokyo, 1992, pp. 107–127.
- [53] A.M. Turing. *Computing machinery and intelligence*. *Mind* **59** (1950), 433–460. Reprinted in A. M. Turing, *Collected Works: Mechanical Intelligence* (D. C. Ince, ed.), North-Holland, Amsterdam, New York, Oxford, Tokyo, 1992.

- [54] A. M. Turing. *Collected Works: Mathematical Logic* (R.O. Gandy and C.E.M. Yates, Editors). Elsevier, Amsterdam, New York, Oxford, Tokyo, 2001.
- [55] J. van Leeuwen and J. Wiedermann. *The Turing Machine Paradigm in Contemporary Computing*. In *Mathematics Unlimited – 2001 and Beyond* (eds. B. Enquist and W. Schmidt), LNCS, Springer, 2000.
- [56] M. White. *Isaac Newton – The Last Sorcerer*. Fourth Estate, London, 1997.
- [57] P. Woit. *Not Even Wrong – The Failure of String Theory and the Continuing Challenge to Unify the Laws of Physics*, Jonathan Cape, 2006.