

Embodying Computation at Higher Types

S. Barry Cooper¹

School of Mathematics, University of Leeds, Leeds LS2 9JT, UK

Abstract: Since Turing's early speculations, progress with the building of intelligent machines has been slow, with obstacles tending to be dealt with in ad hoc rather than theoretical ways. This view of history is reflected in Marvin Minsky's comment (Boston University, 2003) that "AI has been brain-dead since the 1970s".

The optimism of Turing's reported "I am building a brain" (Hanslope Park, 1944) has withered before the reality of the chimera that is embodied intelligence. There is a growing sense that the universal Turing machine has given rise to a paradigm which, as well as being powerful and successful, is misleading in important respects: Firstly, the disembodiment implicit in the Turing machine was unreal even for the early construction of a digital computer, and this has become more apparent with the modern respect given to natural computing in its various guises. And more abstractly, this has made us more aware of the extent to which theory has failed to give due prominence to the role and structure of information in the computational processes that science seeks to characterize.

The theme of this article is the extent to which embodiment, information, and its type structure are relevant to the material and mental world, and their computational processes. In particular, it is argued that 'computation' in the real world is more general — and consequently harder to control — than is usually imagined.

1. The Mathematician's Bias

In October 2010 the ACM, "the world's largest educational and scientific computing society", launched an online Symposium entitled "What Is Computation?". The Symposium was hosted by the web-based magazine *Ubiquity* under the umbrella of the *Communications of the ACM*. *Ubiquity* editor-in-chief Peter Denning's introduction to the Symposium explained:

What is computation? This has always been the most fundamental question of our field. In the 1930s, as the field was starting, the answer was that computation was the action of people who operated calculator machines. By

¹ This article is based on an invited talk at *Turing 100: A Celebration of the Life and Legacy of Alan Turing*, at Boston University, Monday, November 12, 2012. Supported by the John Templeton Foundation research grant: *Mind, Mechanism and Mathematics*, July 2012 - June 2015.

the late 1940s, the answer was that computation was steps carried out by automated computers to produce definite outputs. That definition did very well: it remained the standard for nearly fifty years. But it is now being challenged. People in many fields have accepted that computational thinking is a way of approaching science and engineering. The Internet is full of servers that provide nonstop computation endlessly. Researchers in biology and physics have claimed the discovery of natural computational processes that have nothing to do with computers. How must our definition evolve to answer the challenges of brains computing, algorithms never terminating by design, computation as a natural occurrence, and computation without computers?

Lance Fortnow would have probably represented a majority of readers when he responded in December:

Alan Turing , in his seminal 1936 paper On computable numbers, with an application to the Entscheidungsproblem ..., directly answers this question by describing the now classic Turing machine model. The Church-Turing thesis is simply stated.

Everything computable is computable by a Turing machine.

The Church-Turing thesis has stood the test of time, capturing computation models Turing could not have conceived of, including digital computation, probabilistic, parallel and quantum computers and the Internet. The thesis has become accepted doctrine in computer science and the ACM has named its highest honor after Turing. Many now view computation as a fundamental part of nature, like atoms or the integers.

So why are we having a series now asking a question that was settled in the 1930s?

A few computer scientists nevertheless try to argue that the thesis fails to capture some aspects of computation. Some of these have been published in prestigious venues such as Science ..., the Communications of the ACM ... and now as a whole series of papers in ACM Ubiquity. Some people outside of computer science might think that there is a serious debate about the nature of computation. There isn't.

Before the 1990s, public dissent from this view would have been a bold step, risking ridicule and damage to academic standing. Things have changed, and the ACM Symposium is indeed a sign that the ground is moving under our feet. The revolution in progress has yet to reach Copernican scale, though promising a change in our relationship with causality — with our very conception of causality — as far-reaching as any previous shift in world view. A key contribution to the Symposium was Dennis Frailey's *Computation Is Process* — including the words:

The concept of computation is arguably the most dramatic advance in mathematical thinking of the past century. ... Church, Gödel, and Turing defined it in terms of mathematical functions ... They were inclined to the view that only the algorithmic functions constituted computation. I'll call this the "mathematician's bias" because I believe it limits our thinking and prevent us from fully appreciating the power of computation.

The view we give voice to below — increasingly widely held — is that the discussion is about more than the nature of computation. It is part of a growing appreciation of the fact that information is not flat, and that the accompanying computational modeling must be able to cope with the underlying causal confusions widely affecting science and the humanities. Mathematicians, even those aware of the wider challenges, have largely kept to their technical and conceptual comfort zones. No wonder there is widespread disillusionment regarding the role of mathematics in this informationally extended world. For Nassim Taleb, with his ‘ludic fallacy’, what mathematicians can deal with is close to his *definition* of what we should not be interested in. Those that deal with the mathematics of information beyond the reach of computation, and of language even, are largely invisible.

2. Causality Under Scrutiny



The ability to mentally image, process, and give meaning to emergent form in nature involves a computational process which the human mind manages quite impressively. The above termite ‘cathedral’ is produced by relatively primitive organisms which one might simulate on ones computer with some success, including a degree of interactivity: but neither termite nor computer will simulate the appreciation of emergent form, of higher order information, of which the human brain is observed to be capable of. Still less will the termite represent that higher order information in a form suitable for accessing by further computational routines. We will need to look more carefully at the causal character of

emergence in the light of a closer engagement with the underlying information, via various contexts.

The mathematical character of the challenge presented by causality is specially clear from the physics. As Lee Smolin elaborates on at length in *The Trouble with Physics* (p.241) “causality itself is fundamental”. Causality is key to higher order concepts, while being itself in need of clarification. For instance, regarding the tendency to relate questions about determinism to ones concerning causality, John Earman says (*A Primer On Determinism*, D. Reidel/Kluwer, 1986, p.5):

... the most venerable of all the philosophical definitions [of determinism] holds that the world is deterministic just in case every event has a cause. The most immediate objection to this approach is that it seeks to explain a vague concept — determinism — in terms of a truly obscure one — causation.

From our perspective, part of the confusion arises from the attempt to regard causality as being a simpler concept than that of determinism. While one might be happy with higher order entities being determined, the concept of causality is generally attached to more local phenomena — events, component individuals. According to the contemporary focus on process, emergence, chaos, turbulence, connectionist models of computation and global features of social and economic contexts, it is this leads to philosophical problems with causality, and a parallel dissatisfaction with the scope of the classical model of computation. Here is Albert Einstein raising (*Autobiographical Notes*, in *Albert Einstein: Philosopher-Scientist*, P. Schilpp, ed., Open Court Publishing, 1969, p.63) a particularly fundamental issue about the standard model of physics — one which is still with us. Implicitly the complaint is that, according to what we know, causality in physics is *incomplete*:

... 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) ...

Over the years Einstein’s concern has been echoed by many others. For instance Peter Woit, in his book *Not Even Wrong — The Failure of String Theory and the Continuing Challenge to Unify the Laws of Physics*, Jonathan Cape, 2006, says:

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

In recent months, CERN watchers have looked in vain for confirmation of supersymmetry, an important route to clarifying a number of issues — such as dark matter — muddying our view of the universe. Looking for causality even further up the informational scale, no less than Alan Guth (author of *The Inflationary Universe — The Quest for a New Theory of Cosmic Origins*, Addison-Wesley, 1997), asks about the determination of natural laws:

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?

If our incomplete grasp of the causal character of the physical universe is unsatisfactory, even more so is the ‘causality abrogated’ that the uncertainty phenomenon in quantum mechanics pushes us towards. Taking its cue from the probability-bound collapse of the wave function consequent on a measurement, and Hugh Everett III’s many-worlds setting for entanglement, we now have a byzantine landscape of multiverse proposals, structured by Max Tegmark into his Multiverse Levels I-IV (see his May 2003 *Scientific American* article on *Parallel Universes*, pp.40-51). The abrogation finds its quintessential expression in the *anthropic principle*, delivering our universe of conveniently appropriate fundamental constants and natural laws by some sort of cosmic accident. David Deutsch in his 1997 *The Fabric of Reality* (Penguin) sees the existence of quantum computers as pointing to the existence of the multiverse. Scott Aaronson (*Quantum Computing since Democritus*, Cambridge University Press, 2013, p.149) comments:

... a quantum computer is not a device that could “try every possible solution in parallel” and then instantly pick the correct one. If we insist on seeing things in terms of parallel universes, then those universes all have to “collaborate” — more than that, have to meld into one another — to create an interference pattern that will lead to the correct answer being observed with high probability.

3. Causality Fragmented, Supervenient ... Simulated?

So causality is basic — but a ‘truly obscure concept’, incomplete, abrogated — and, unavoidably fragmented in the context of the human brain. There are many examples of philosophers (particularly) rejecting Cartesian dualism in uncompromising style — echoing Susan Blackmore’s “grand illusion of consciousness”. Here is a quote from Professor Blackmore, the analysis in terms receptive to an information-theoretic interpretation (page 220 of *The Meme Machine*, Oxford University Press, 1999):

Dualism is tempting but false. For a start no such separate [thinking] stuff can be found. If it could be found it would become part of the physical world and so not be a separate stuff at all. On the other hand if it cannot, in principle, be found by any physical measures then it is impossible to see how it could do its job of controlling the brain. How would immaterial mind and material body interact? Like Descartes’ ‘thinking stuff’, souls, spirits and other self-like entities seem powerless to do what is demanded of them.

Such issues are visited in depth by the philosopher Jaegwon Kim, echoing such questions as: How can mentality have a computational role in a world that is fundamentally physical? And what about ‘overdetermination’, the problem of phenomena having both mental and physical causes? In *Physicalism, or Something Near Enough* (Princeton, 2005) Kim puts it:

... the problem of mental causation is solvable only if mentality is physically reducible; however, phenomenal consciousness resists physical reduction, putting its causal efficacy in peril.

The persistent problems in reconciling mentality with its physical host are arranged by philosophers who care about such things under the heading of *supervenience*, which (quoting Jaegwon Kim: *Mind in a Physical World*, MIT Press, 1998, pp.14-15):

... represents the idea that mentality is at bottom physically based, and that there is no free-floating mentality unanchored in the physical nature of objects and events in which it is manifested.

Put more mathematically (*Stanford Encyclopedia of Philosophy*): ‘A set of properties A supervenes upon another set B just in case no two things can differ with respect to A-properties without also differing with respect to their B-properties.’

The widespread consensus around supervenience hosts a huge variety of ideas regarding exactly *how* mentality is anchored in the physical. Interestingly, even a fairly committed reductionist like Daniel Dennett favours an analysis with different orders of information, even if it does not directly relate to the supervening. The term ‘subconscious’ may be a misleading one — with consciousness supervening on the physical, but unconscious mental activity — even that associated with reflex sensorimotor activity — being mediated by a developing consciousness in a quite nonlinear and selective fashion. In his 1991 *Consciousness Explained* (Little, Brown & Co.) Dennett can be found plausibly referring to unconscious ‘thoughts’ as a higher order phenomena (p. 308):

Unconscious thoughts are, for instance, unconscious perceptual events, or episodic activations of beliefs, that occur naturally — that must occur — in the course of normal behavior control. Suppose you tip over your coffee cup on your desk. In a flash, you jump up from the chair, narrowly avoiding the coffee that drips over the edge. You were not conscious of thinking that the desk top would not absorb the coffee, or that coffee, a liquid obeying the law of gravity, would spill over the edge, but such unconscious thoughts must have occurred — for had the cup contained table salt, or the desk being covered with a towel, you would not have leaped up. Of all your beliefs — about coffee, about democracy, about baseball, about the price of tea in China — these and a few others were immediately relevant to your circumstances. If we were to cite them in an explanation of why you leaped up, they must have been momentarily accessed or activated or in some way tapped for a contribution to your behavior, but of course this happened unconsciously.

There are more observable signs of fragmented causality at work: with the task being to provide the wherewithal of a successful partnership between different levels. For Alan Turing back in the 1940s it was a steep learning curve. As Solomon Feferman describes (pp. 131-2) in his magical *Turing in the Land of O(Z)* — in *The Universal Turing Machine: A Half-Century Survey* (ed. R. Herken), Oxford University Press, 1988:

Turing, as is well known, had a mechanistic conception of mind, and that conviction led him to have faith in the possibility of machines exhibiting intelligent behavior.

Since that time, scaling the gap between computer and mentality has presented investigators with a challenging phase-transition — an obstacle theoretically based in the ‘causality’ of the underlying informational structure, we will argue. For those working at the practical level, the theoretical approach — via the logic and mathematical modeling — has been singularly unrewarding. For Rodney Brooks, the experimental robotic route has made more sense. For him (quoted in *Nature* in 2001):

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

On the other hand, such ad hoc investigations leave artificial intelligence theoretician and pioneer Marvin Minsky unimpressed — commenting in characteristically trenchant terms at Boston University in May 2003:

AI has been brain-dead since the 1970s.

In his final years, Turing was already anticipating difficulties ahead, with thinking moulded by a focus on the human dimension. For instance, the mistake-prone young writer of the seminal ‘computable numbers’ article comes to mind when we read (talk to the London Mathematical Society, February 20, 1947, quoted by Andrew Hodges in *Alan Turing: the Enigma*, p. 361):

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

And, in Turing’s popular piece on *Solvable and Unsolvable Problems* (*Penguin Science News* 31, 1954) we read in his final paragraph p. 23:


The results which have been described in this article are mainly of a negative character, setting certain bounds to what we can hope to achieve purely by reasoning. These, and some other results of mathematical logic may be regarded as going some way towards a demonstration, within mathematics itself, of the inadequacy of ‘reason’ unsupported by common sense.

The human brain itself uncannily reflects elements of the Brooks-Minsky dialectic, and of the Turing preoccupation with cooperation between machine and human. In his 2009 Yale University Press book *The Master and his Emissary: The Divided Brain and the Making of the Western World*, Iain McGilchrist describes how:

The world of the left hemisphere, dependent on denotative language and abstraction, yields clarity and power to manipulate things that are known, fixed, static, isolated, decontextualised, explicit, disembodied, general in nature, but ultimately lifeless. The right hemisphere by contrast, yields a world of individual, changing, evolving, interconnected, implicit, incarnate, living beings within the context of the lived world, but in the nature of things never fully graspable, always imperfectly known — and to this world it exists in a relationship of care. The knowledge that is mediated by the left hemisphere is knowledge within a closed system. It has the advantage of perfection, but such perfection is bought ultimately at the price of emptiness, of self-reference. It can mediate knowledge only in terms of a mechanical rearrangement of other things already known. It can never really ‘break out’ to know anything new, because its knowledge is of its own representations only. Where the thing itself is present to the right hemisphere, it is only ‘re-presented’ by the left hemisphere, now become an idea of a thing. Where the right hemisphere is conscious of the Other, whatever it may be, the left hemisphere’s consciousness is of itself.

Why should there be different kinds of thinking, and why should the brain architecture so separate their physical hosts? Is this a purely ad hoc, even accidental, adjustment within the evolutionary process, or is there something more fundamental at work? The enduring role of natural selection, in the context of the growth of complexity theory, leads us to expect something ‘more fundamental’. As Steven Pinker attempts to persuade us (pp. 161-162 of his *How the Mind Works*, W W Norton, 1997), with a degree of success:

The “complexity” that so impresses biologists is not just any old order or stability. Organisms are not just cohesive blobs or pretty spirals or orderly grids. They are machines, and their “complexity” is functional, adaptive design: complexity in the service of accomplishing some interesting outcome. ... No set of equations applicable to everything from galaxies to Bosnia can explain why teeth are found in the mouth rather than in the ear. And since organisms are collections of digestive tracts, eyes, and other systems organized to attain goals, general laws of complex systems will not suffice. Matter simply does not have an innate tendency to organize itself into broccoli, wombats, and ladybugs.



There is certainly an observed level of robustness of the brain architecture across a very wide spectrum of living organisms. Though there will be doubts about the degree to which the ‘design’ is well enough defined, in terms of mechanism, to be properly termed ‘design’. A key aspect of the brain architecture of placental mammals, such as humans, is the corpus callosum, connecting and mediating the functionality of the separate hemispheres. McGilchrist comments (pp. 18-19):

... the evidence is that the primary effect of callosal transmission is to produce functional inhibition.

... it turns out that the evolution both of brain size and of hemisphere asymmetry went hand in hand with a reduction in interhemispheric connectivity. And, in the ultimate case of the modern human brain, its twin hemispheres have been characterised as two autonomous systems.

So is there actually some purpose in the division of neuronal, and therefore, mental processes? If so, what could that be?

We might further ask: Given this division, to be found in animals generally, what is the benefit of the moderated reconnection via the corpus callosum? And what is the explanation of the placental mammal connection? Relevant to our more general concerns with the computational content of causality in a very general context, we might reframe these questions:

What are the computational parallels and distinctions between the different modes of dealing with information? And between the ways in which the computational modes can effectively combine? How do these modes relate to the classical Turing model of computation? And is there a suitable adaptation of classical computability theory capable of accommodating them in an informative way? And crucially for life in our computationally hybrid world — what is the appropriate balance, and level of interaction and autonomy, to be granted these computationally different frameworks? And what are the mechanisms, or absence of mechanisms, to ensure such a balance?

The hope is that we are seeing a new level of exchange of experiences and analytical tools with the potential to give us not just enhanced human functionality, but a new consciousness of the special relationship between different ways of thinking — based in a suitably updated and more mature appreciation of the nature of causality and its computational framework.

A first step is to look more closely at the classical model.

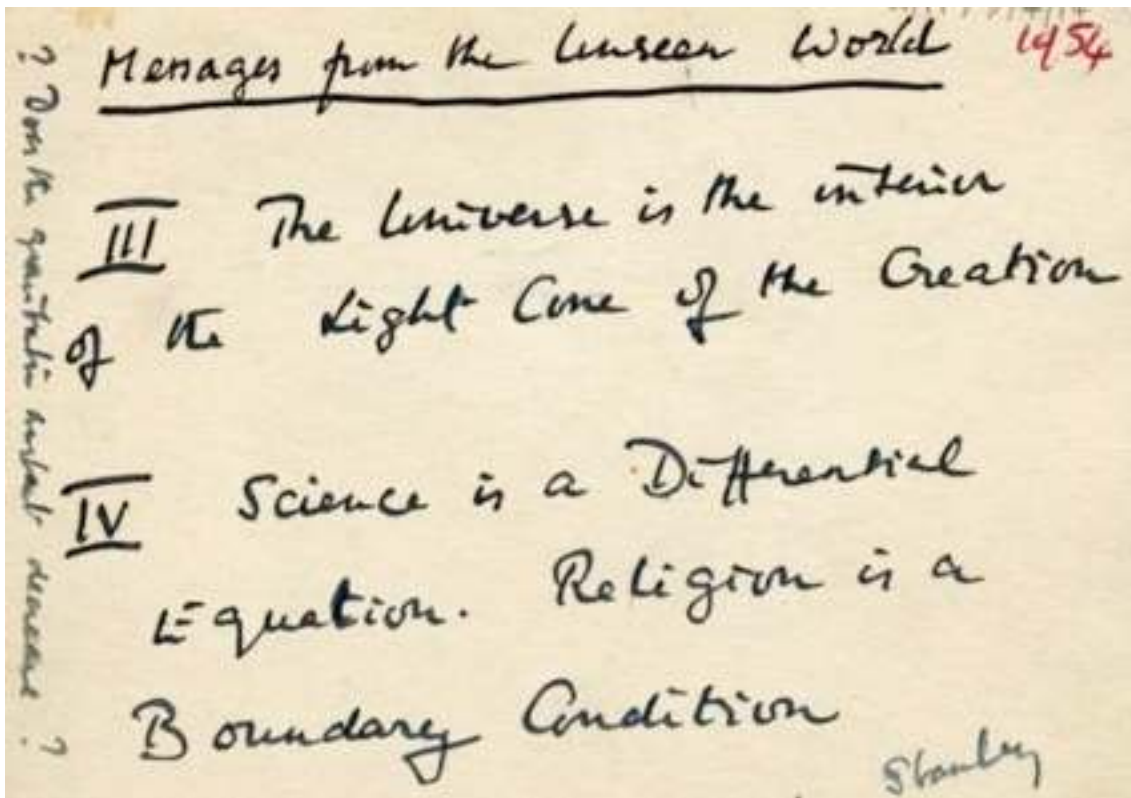
4. Causality as Computation - “I am building a brain”

Despite the distrust of those with a more spiritual, artistic or humanistic view of their world, the scientific world we inherited from Isaac Newton and his contemporaries has served us well. Einstein’s comment (p.54 of *Out of My Later Years*, 1950) that:

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

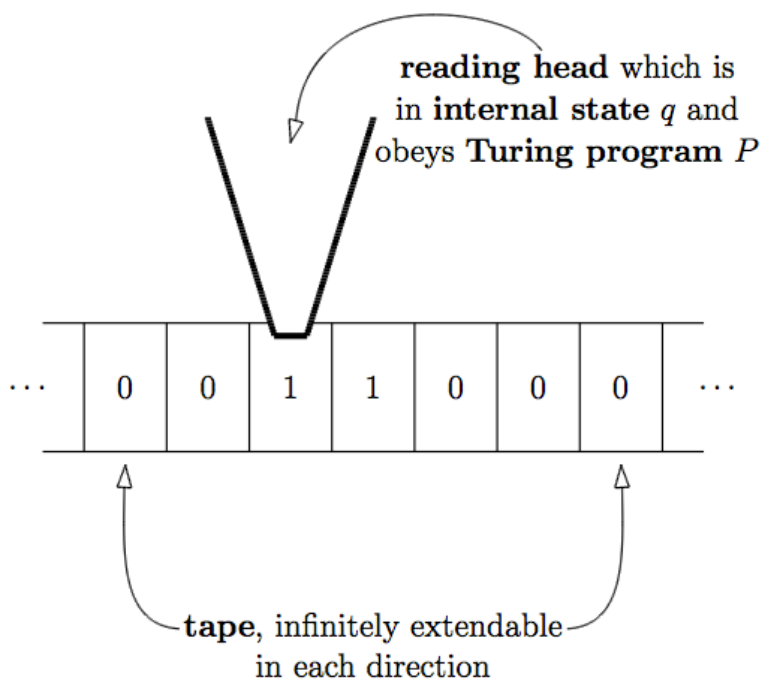
underpins a hard-nosed approach to knowledge of the world, which expects epistemology to provide off-the-peg knowledge for mass consumption, knowledge machine reproducible via proofs and computer calculation. The mechanics of the mathematical model underlying so much of the science was originally extracted from a well-defined physical context. Turing based them on the actions of human ‘computers’ (usually women in pre-1936 days) calculating in a given language, with pencil and paper, according to specific instructions.

The close relationship of the model to the physical world reflected Max Newman’s description of Turing as “... at heart more of an applied than a pure mathematician” — although the model became relevant in many other contexts: To computing machines, physics, and biology, for instance. Turing had a lifelong interest in basic elements of science and engineering. But this was not part of what was depicted in the 1936 paper on computable numbers. As a schoolboy, Turing had studied Einstein’s theory of relativity;



and in Cambridge had attended lectures of Paul Dirac on quantum mechanics. Robin Gandy, Turing's only doctoral student, actually started out as an applied mathematician, and was employed as such as a lecturer at Leeds University in the late 1950s. The postcards Turing sent to Gandy before he died in 1954 (see the example above) seem to us now to show the influence of Dirac and the keen interest in questions of causality and computation raised by the new science.

The Turing machine *disembodied* a rather specific part of the real world, and served to disembody our vision of computation much more widely — and, arguably, held in tow a deal more disembodiment than was appropriate.



Having erased the human ‘computers’ from the picture, the disembodied hardware appeared trivial. It consisted of no more than an extendable tape subdivided into cells in which to record symbols, and a single reading/writing head - see above diagram. While the basic actions performed by the reading head were chosen to generate all the more complicated actions one might envisage, ensuring that the potential was there to compute any real number computable by an algorithm. Essentially, with the right usage, the Turing machine hardware was ‘Turing complete’, in that (modulo some less than trivial re-embodiment) it could compute anything the human ‘computer’ could compute. And various people duly verified that all the other abstract models of such computation which they encountered — some apparently quite different — had exactly the same power. Of course, different models could be applied to different languages, with various computing conventions, but this was immaterial in relation to overall computing potential.

So the hardware was as simple as could be, and all the computing power was in the program it was presented with. And — something which underlies the sometimes fractious contacts between theoretical computer scientists and engineers — the convention arose whereby the theorists identified the machine with its program. Having disembodied the computing machines of earlier centuries — the machines which had to be physically built pre-Turing to do different kinds of computational chores — the theorists muddied the waters by giving the program, something which was not a machine at all, the name ‘Turing machine’. In the course of time, the distinctions between machine and information, and between program and data have receded in relation to both the engineering and theoretical worlds. And the seeds of this were laid by Turing’s very special *universal machine*, which he used to derive a computational counterpart to Kurt Gödel’s 1931 Incompleteness Theorem.

To understand the significance and mathematical sleight of hand encased in the universal Turing machine, one needs a little more detail than is usually given. The important property the universal machine (that is, its program) had to have, was crucially: That all one had to do was to give the universal machine a piece of data which said “Please, compute with input x just like that other Turing machine over there (which you never saw before) would”, and it would do it! The trick was, that instead of a lumbering great contraption, like Charles Babbage’s Analytical Engine would have been if built, which would be hard to code up as a single piece of input data! — a Turing machine was entirely determined by a finite set of instructions, a set which held within it all the logical structure needed for a particular computational chore: it could hence be coded up as a single piece of input data — say a number, or a binary string, or whatever the machine was used to.

So let’s assume we have coded up all the Turing machines as natural numbers e , say, via a listing of all the Turing machine programs. We can then program our universal Turing machine \mathbf{U} , say, so that when it is given a pair (e,x) of natural numbers, it decodes the machine \mathbf{T} coded by e , and then proceeds to compute exactly like \mathbf{T} would with input x . You can actually write down the program for such a \mathbf{U} , and this has been done in various forms. We write $\varphi_{\mathbf{T}}$ (φ_e) for the number theoretic function of one argument computed by \mathbf{T} .

The huge significance of the coding trick is that we have reduced some complicated information (a suitably set up Babbage machine, say) to a simple piece of data. You might not be too impressed. You might say “Ah, but all I need is a set of plans of the machine, and I can build it — and then a code for the plans is essentially the machine — what’s the big deal?” The ‘big deal’ is that you are using Turing’s work with hind-sight. If I gave you a

bucket of slime mould (see, for instance, Andrew Adamatzky et al, *On creativity of slime mould*, Int. J. of General Systems, vol. 42, 2013, pp. 441-457), the coding capturing the computing medium might be a little trickier ... i.e., non-trivial, and just perhaps an indication of a more complicated level of information that our Turing machine cannot handle. Later, we will look at the type structure of information, and look on the reduction of machine to code by Turing as a route to less simple computational *type reduction*. And following Turing one must then view the machine itself, regardless of the informational type a description it demands, as an input into the computational process.

In the more limited context of computer history, it needs to be noticed that the actual computing machines built involved a high degree of *re-embodiment*. It was not just the discovery of basic elements of the Turing machine hardware, such as material hosts for the memory, which were required. For very important operational reasons, basic *logical* structure was embodied in the hardware. In the more iconoclastic modern re-assessment on disembodied computation — which this article is intended as a contribution to — various writers (usually from outside mathematics) have immersed themselves in the minutiae of particular computer solutions, and discovered that Turing's universal machine was fairly unrecognizable therein. See for example Thomas Haigh's conference article on '*Stored Program Concept*' *Considered Harmful: History and Historiography* (in *The Nature of Computation. Logic, Algorithms, Applications*, ed. Paola Bonizzoni, Vasco Brattka, Benedikt Löwe, Springer, 2013). Though from a mathematical perspective, the big idea of universality can be swarmed all over with Lilliputian thoroughness, but its powerful vitality is hard to keep down.

It is the universal machine, and its strangely fortuitous genesis in type reduction, that underpins the functionalist view of the informational universe and its computational structure — and then brings it into doubt as a comprehensive computing paradigm. The functionalist view — which stresses what a computer *does* as something realizable via *different computing platforms* — was first developed in relation to philosophy of mind and AI. A seminal contribution was Hilary Putnam's 1960 writing on *Minds and Machines*. The idea can also be found in muted form in the notion of *virtual machine*, emerging in the computing world in the late 1960s, for instance in the IBM time-sharing operating system CP/CMS (Control Program/Cambridge Monitor System).

The post-1936 developments were dramatic, and must have been very exciting for those in at the start of the computer age. One can excuse Turing for getting carried away, at one point quoted by Donald Bayley as saying, while working on the 'Delilah' speech encryption project at Hanslope Park, that he was "building a brain".

In recent times, Putnam has drawn back somewhat from the 1960s functionalist simplicity, retaining some elements and disposing of others. Here he is (p.89 of *Is the Causal Structure of the Physical Itself Something Physical?* in *Realism with a Human Face*, Harvard University Press, Cambridge, MA, 1990 [35]), qualifying his earlier view:

... if the physical universe itself is an automaton ... , then it is unintelligible how any particular structure can be singled out as "the" causal structure of the universe. Of course, the universe fulfills structural descriptions - in some way or other it fulfills every structural description that does not call for too high a cardinality on the part of the system being modeled; ... the problem is how any one structure can be singled out as "the" structure of the system.

This does not derail the power of the paradigm. There are many successful reductions of ‘natural’ examples of computation to the Turing model. A particularly striking example is David Deutsch’s 1986 use of the quantum Turing machine model to bring the ‘standard model’ of quantum computation within the computationally classical fold. But there is a giant step from this to Deutsch’s assertion (Question and Answers with David Deutsch, at New.Scientist.com News Service, December 2006) that:

I am sure we will have [conscious computers], I expect they will be purely classical, and I expect that it will be a long time in the future.

5. Complexity, Emergence, and ‘Levels of Abstraction’

Turing’s last published paper is not only his most cited, but also the most dramatically prescient following the 1936 career-launching ‘computable numbers’. Today, at a time when computational models from nature and bioinformatics frame the professional lives of literally thousands of scientists, it is instructive to revisit the computational underpinnings of Turing’s patterns in nature. In 1936 the young Alan Turing had broken a computational context into basic ingredients which possessed a clear computationally causal character. In his 1952 paper *The chemical basis of morphogenesis* (in *Philosophical Transactions of the Royal Society of London. Series B*, **237** (1952), 37-72) Turing attempts something similar for a selection of pattern forming processes in nature. The differential equations he formulated from the proposed reaction-diffusion models gave solutions suggesting that the analysis was broadly correct. Research since 1952 has partly confirmed Turing’s approach, and partly led to modifications and new directions. What is important for us is that Turing had effectively given descriptions of the morphogenesis in computational terms, which pointed the way to notions of computation, embodied computation, with the potential to bring together the descriptive and computational approaches to epistemology. It opened out a route to a better understanding of ‘how the world computes’. For more explanation and history of Turing’s work and influence in this area, good sources can be found in *Alan Turing - His Work and Impact* — particularly informative are the articles by James Murray and Philip Maini.



The solutions to Turing’s differential equations were point-wise computable, and Turing was able to simulate them in Manchester using the new computer technology — in effect, placing these surprising outcomes firmly within the classical framework. However, the descriptive form of the solutions had the power, potentially, to resist this kind of — what we will call — *type reduction*. Notice that what is interesting about the simulations — and morphogenesis in nature — are the patterns. Whereas the computer’s focus is point-wise. It occupies the world of the termite, computing without being able to take full ownership of its creation. Turing’s vision brings us closer, not just to understanding morphogenesis as a

The solutions to Turing’s differential equations were point-wise computable, and Turing was able to simulate them in Manchester using the new computer technology — in effect, placing these surprising outcomes firmly within the classical framework. However, the descriptive form of the solutions had the power, potentially, to resist this kind of — what we will call — *type reduction*. Notice that what is interesting about the simulations — and morphogenesis in nature — are the patterns. Whereas the computer’s focus is point-wise. It occupies the world of the termite, computing without being able to take full ownership of its creation. Turing’s vision brings us closer, not just to understanding morphogenesis as a

phenomenon, but to a better understanding of levels of computation and how they are achieved.

Concerning the problem of pattern recognition and extraction of meaning, Judson C. Webb (*Mechanism, Mentalism, and Metamathematics - An Essay on Finitism*, D. Reidel Publishing Co., 1980, pp. 246-247) writes:

Our discussion of abstract Turing machines has, of course, taken the problem of 'pattern recognition' largely for granted, but it is by no means obvious that a physical realization of a Turing machine can always be programmed to learn to recognize the global properties of arbitrarily degraded symbols. ... Weizenbaum ... points out that the general problem of computer vision "is in many respects fundamentally the same as that of machine understanding of natural language" The historical roots of this problem go back ... to the issue of blind mechanisms doing geometry that arose in Poincare's critique of Hilbert's 'purely logical' approach to geometry. ... I believe that the ultimate test of artificial intelligence — at least in the case of mathematics — will come in this geometric arena which requires the coordinated use of both the perceptual and cognitive abilities of machines, as opposed to the relatively barren arithmetical arena in which machines need only think.


Poincare's observation turns out to be a key one, backed up by subsequent developments.

Pattern recognition is a first step towards extracting semantic content from patterns. Let us illustrate this with an example from mathematics, where both computational content and the human appreciation of higher type information is familiar to us. The Mandelbrot set provides a simple example of a computationally based structure which carries with it higher order information of great interest to many leading researchers — see for example, Lenore Blum, Felipe Cucker, Michael Shub and Stephen Smale, *Complexity and Real Computation*, Springer, 1997. The fascinating form of the Mandelbrot set emerges via an approximation to its representation in the complex plane, embodied on the observer's computer monitor. Despite the book of Blum, Smale et al, there are questions concerning the point-wise computability of the set. The members of the set are the complex numbers c for which the limiting behaviour of the sequence $z_0=0, z_1, z_2, \dots$, governed by the recurrence relation

$$z_{n+1}=z_n^2+c,$$

is bounded.

The logician will easily spot that the description of such numbers c involves the addition of a couple of alternating quantifiers, a source of alarm as we shall see that quantification was what takes the observed behaviour of Turing's universal machine \mathbf{U} out of the purview of itself. The mathematics tells us enough about the bound involved to reduce the



description to one universal quantifier, giving the description of the Mandelbrot set the flavour of Turing's 1936 example of an *incomputable* set.

For the core computable analysis community the computability of the Mandelbrot set is still correspondingly in question. Interesting as this is, it is not what impresses us out in the real world. Our mental image of the form of the set is unrestricted by the point-wise computations, and we delve deeper and deeper via our computer approximations, excavating surprise upon surprise, aesthetic pleasure upon aesthetic pleasure, as the infinite variety of shape within is unfolded. It is not the facility of the brain for receiving and internally representing the shapes which is so remarkable — it is the sense of context which comes into play, the resource-based expectations, the higher order comparison with previously stored form and the mental repercussions which feed into our subsequent thinking. Echoing Judson Webb, it is 'by no means obvious' that our computer can be programmed to share this experience, still less to iterate the shared experience as the human brain might.

The problem here parallels the challenge of 'big data' in today's wider world. In this context great strides are being taken. An appreciation of the mathematical context of the statistical tools investigators are driven to here may enable us to better understand the potential convergence of real-world models, its scheduling, and the future relationship of the available models — human and human constructed.

Underlying this scenario (see the next section) is the role of descriptions, their relationship to Turing's classical computation, and the way in which this relationship creates informational structure relevant to our residence in the real world. There are different entry points to this landscape. There is the 'levelism' of the rich repertoire of mathematical hierarchies, mainly the domain of the logician — in general, mathematicians like to be given a specific informational context to solve their puzzles within. And there is the descriptive natural language based approach of the mass of humanity, at its most analytical reaches represented by the philosophically minded — and trained. Of course, there are increasingly people who can traverse with a degree of expertise the boundaries between different approaches and disciplines.

A recent, and timely, intervention from the philosophical side is Luciano Floridi's book on *The Philosophy of Information* (Oxford University Press, 2011). *Levels of Abstraction* play a key role in his advocacy of ISR (Informational Structural Realism) as an alternative to the informational and computational flatness of the dominant digital ontology. According to Floridi:

A level of abstraction (LoA) is a finite but non-empty set of observables. ... The introduction of LoAs is often an important step prior to mathematical modelling of the phenomenon under consideration ... Use of LoAs is effective ... where a typed theory would be effective ... [but] analysis ... may be conducted at different levels of epistemological levelism.

Floridi is keen to emphasize the full diversity of LoAs, without being too specific about their basis. The more appropriate mathematical analysis is to a given context, such as in physics or other physical sciences, the more prepared he is to admit the relevance of informational typing. Of course, the mathematical analysis needs to respond to the way in which the 'levelism' can be transcended by relationships and reductions between levels — duly captured conceptually, if not in specific detail, by corresponding aspects of the

mathematics; while the semantics need to recognize relationships between the constructive levels to hierarchical levels of information, and the physical processes and phase transitions which impact on the nature and frontiers of the domains identifiable semantically as LoAs.

This impact is as significant in the ontological and epistemological domains as it is in the more abstract logical and mathematical contexts. The typing arose from Bertrand Russell's work on setting rules for the epistemology which purged the science of paradoxical description of sets. Kurt Gödel gave a simple description of the basic typing in his *Russell's mathematical logic* in 1944:

By the theory of simple types I mean the doctrine which says that the objects of thought ... are divided into types, namely: individuals, properties of individuals, relations between individuals, properties of such relations, etc. ... , and that sentences of the form: " a has the property ϕ ", " b bears the relation R to c ", etc. are meaningless, if a, b, c, R, ϕ are not of types fitting together. Mixed types (such as classes containing individuals and classes as elements) and therefore also transfinite types (such as the class of all classes of finite types) are excluded. That the theory of simple types suffices for avoiding also the epistemological paradoxes is shown by a closer analysis of these.

We can clearly see the typing at work in the transition from individual type-1 members of the Mandelbrot set to its conflated type-2 whole. And the commonly encountered descriptive progression from underlying computational causality to emergent form can be more generally related to the type-theoretic framework.

The improved clarity concerning LoAs and their relationships pays off. It makes better sense of the way in which descriptions leading to new scientific (epistemological) levels substantiate and are clarified by what we know about the computational structure accompanying the typing, and the finer mathematical infrastructure carried with it. We need to carry this forward. Without it we are only a step ahead of the termites in our understanding of the emergent forms and the observed LoAs. What is emergence, and why is it difficult to identify when it is taking us beyond the classical model? As Ronald C. Arkin points out (*Behavior-Based Robotics*, MIT Press, 1998, p.105):

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?

Appearances without analysis can be misleading, as the so-called 'British Emergentists of the inter-war years discovered (See Brian 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). This despite the fact that the conceptual grasp was impressively modern, anticipating the today's widespread use of the notion of emergence, in very different contexts. Philosopher Charlie Broad, at Cambridge at the time as Turing, characteristically wrote (C. D. Broad, *The Mind and Its Place In Nature*, Kegan-Paul, London, 1923, p.623):

... the mental properties of those events which do have mental properties are completely determined by the material properties which these events also have ... it is certainly not ... a form of Reductive Materialism; it is a form of the theory ... of Emergent Materialism.

6. Higher Type Incomputable Information & Randomness

Turing's interest in morphogenesis, and the discovery of *descriptions* of emergent patterns in nature based in *computationally characterized* causality, gave us a better understanding of how another level of information could arise — *computed* (in some sense to be clarified) by nature — with new properties of its own. Emergence takes us to a new world of higher type information computed in a higher type fashion. And in the new world of higher type computation, we must get used to things being rather different, and paradigm breaking. We must face challenges to universality, closed computational context, and an extended Church-Turing thesis.

On the positive side, we may build a bridge from the digital ontology of the scientist and classical computational model, to the world of language, common sense, imagination and intuition, still incorrigibly inhabited by most of us: a world which haunted Alan Turing throughout his life, as we observe in his 'popular' writings and broadcasts after the war, and even glimpse in occasional informal comments in his more technically arcane writings, such as in his 1939 'ordinal logic' paper. One can think of this bridge between intellectual heuristics as a form of virtual corpus callosum: a kind of abstract correlate of the complex connectivity between two kinds of thinking, a connectivity special to placental mammals, and most highly developed in humans.

Of course, the main aim of Turing's 1936 paper was not just to model basic computational practice, but to also show that the model was mathematically transcended. *Incomputability* was the *point* of the paper — not the laying of the theoretical foundations of computer science. And in a mathematically precise way he mapped out the steps up a level of the type structure, a Jack and the Beanstalk scaling of the heights leading to a new world of dangers and opportunities. The ascent would be easier than the way back, a familiar part of the computational terrain. But with ingenuity one might contend successfully with the 'big data' and deliver - like Churchill's workers at Bletchley Park - the golden egg of decrypted information content.

Taking a universal Turing machine **U** and a natural number x , and one can ask whether there is a computer program to tell us (in general) whether **U** given input x ever stops computing. In 1936 Turing (essentially) proved the unsolvability of this so-called *Halting Problem* for **U**:

*No computer can tell us, for each given input x , whether **U** will ever complete its computation with input x .*

Actually, what Turing did was have his machines compute binary real numbers. And a computable real $r \in (0,1)$ would be one whose non-terminating decimal representation was computed by a Turing machine **T**, say, writing $r = r_T$. Then Turing gave us an effective version of Cantor's theorem concerning the uncountability of the open interval $(0,1)$. In the spirit of Gödel's Incompleteness Theorem, Turing showed that **U** could not have computational knowledge of its own *collated* outputs (the jump in type): there were Turing

machines T for which U could not tell whether the decimal representation r_T was a fully computed real or not. Otherwise, one could computably list all computable reals, computably diagonalize the list, and get a new real r_{T^*} , say, not on the list.

The more usual incomputable object associated with U is its Halting Set

$$H_U = \{x \in \mathbb{N} \mid \varphi_U(x) \text{ is computed}\}.$$

The power of first-order Peano Arithmetic (PA) to ‘semi-represent’ the Halting Set gave undecidability of PA — and most strikingly, the negative solution to David Hilbert’s 1928 *Entscheidungsproblem*, in the form of ‘Church’s Theorem’ saying, in simple terms:

No computer can tell us, for each given sentence of first-order logic, whether it is logically valid or not.

Church’s Theorem can seem somewhat counter-intuitive at a first encounter, though our frustrating experiences with embodied computers prepares us for the unpredictability of computational outcomes!

Out in the real world, unpredictability is omnipresent, without shaking the confidence of many in the basic validity of the classical model, and its comprehensive extendability. To others, it seems more likely that we can blame faulty epistemology for an imperfect understanding of the structure of embodied information than for the unpredictability of the phenomena it encloses. Just as we had to get used to the reality of wave-particle duality at the quantum level, we may yet have to accept the crushingly simple observation that both quantifiers and limits provide instruments for the ascending of the informational type structure. And that machine, as embodied higher type information will not always come with parallel devices (statistical or otherwise) for type reduction.

Given that the observer will meet difficulties in observing beyond her operative LoA, particularly in traversing phase transitions between basic causal context and emergent form encasing it (in either direction), it will not be so easy to validate speculations concerning the computational character of these informational frontiers. Some of the most theoretically intractable problems in science and the humanities can be located around such boundaries — between quantum and classical reality, between brain and mentality, and around the balance between the human individual and the emergent social and economic outcomes we seem unable to control.

We conclude this section with a few remarks on how the incomputability accords with what we observe in the outside world, in particular, in regard to the tendency to replace the robustly defined notion of incomputability in the real context with the more arbitrary notion of randomness.

If the precise character of probabilities governing the outcome of a quantum measurement are determined by higher order rules emergent from subatomic structure, what is this subatomic structure? We know from work with networks and emergent biological form that the particularities are often key to the global outcome. The invoking of an absolute randomness to relieve us of the task is philosophically unsatisfying. The mathematician knows that there is no such thing as a robustly defined absolute randomness. There are different approaches to clarifying randomness as a meaningful concept, with some degree of convergence at the very basic levels. But randomness is about *avoidance* of obstacles to our intuitive vision of the concept. And it is in the pinning down of the obstacles that we discover that there is a whole inexhaustible hierarchy (even hierarchies) of randomness.

Of course, quantum randomness is a familiar experimental and theoretical phenomenon for the physicist. As computer scientist Cristian Calude says:

It passes all reasonable statistical properties of randomness.

So the question arises: Just how random is quantum randomness? And given that mere incomputability can look like randomness — maybe quantum randomness is not random at all? Such questions are hard to get a grip on. In a joint paper on *Quantum Randomness and Value Indefiniteness* (*Advanced Science Letters* 1 (2), 2008, pp.165-168), Calude and physicist Karl Svozil formulate some basic assumptions concerning the quantum physics, and on the basis of these manage to demonstrate the incomputability, but not as yet any proof of any level of randomness.

The economics presents very different problems. There are certainly basic computational mechanisms at work in the economic context. But the economics is subject to forces which originate in diverse and quite different contexts, in particular, within the minds of individuals. There is a corresponding interplay of the global and local, with the human aptitude for absorbing and representing observed macro-economic activity for further economic intervention, presenting the economist with complexities crying out for the identification of emergent pattern. This identification must needs be based on observation, as must any resulting computational structure imposed on what one takes as significantly emergent. If one attends conferences of economists, one finds different approaches to pinning down the patterns, sometimes taking the cue from more theoretical areas such as physics, with econophysics one of the more interesting and active areas. But here is Nassim Taleb again, decrying the all-too-often failure of the computational and mathematical approaches:



I have spent my entire life studying randomness, practicing randomness, hating randomness. The more that time passes, the worse things seem to me, the more scared I get, the more disgusted I am with Mother Nature The more I think about my subject, the more I see evidence that the world we have in our minds is different from the one playing outside.

Again, in what sense is it randomness? Is it incomputability? Or even some complexity-theoretic toy version of classical incomputability? If one wants to distinguish between paradigms, one needs to know a lot about both the reality and the theory one is seeking to impose on it.

K. Vela Velupillai is one of the select few economists who have a good understanding of the logic, and in particular, the computability theory. His view of the economic uncertainties that Taleb ponders, by which we are all beset, is nicely captured in *The unreasonable ineffectiveness of mathematics in economics* (Cambridge Journal of Economics, 29, 2005, pp.849-872; reprinted in *Computable Economics*, pp.723-746):

Through [parametric Diophantine equations], we enter the weird and wonderful world of undecidabilities and uncomputabilities, which is why the economic world, in its macroeconomic and microeconomic incarnations, is full of financial ambiguities. To pretend that the economic world can exist without such ambiguities and, hence, occasional and systematic exploitation of logical loopholes by unethical and immoral practitioners of financial wizardry is dangerous. On the other hand, recognising the intrinsic inevitability of such ambiguities may mean that we might happily, in a quantitative economics of the future, return to the Linnean fold, to classify and systematise particular intractable accounting schemes in increasingly and morally transparent ways.

It is in Velupillai's spirit of seeking out the source of the balance between unpredictability and form in computationally complex environments that we focus more closely on the mathematics underlying the apparent randomness.

Before doing that, it is worth saying that computational complexity does not always arrive in the material world red in tooth and claw. Georg Kreisel, an early thinker on stronger versions of the Church-Turing thesis, speculates on the potential for incomputability in the apparent simplicities of Newtonian dynamics, via the underlying infinitary mathematics (*Church's Thesis: a kind of reducibility axiom for constructive mathematics*, in *Intuitionism and proof theory: Proceedings of the Summer Conference at Buffalo N.Y. 1968*, North-Holland, 1970, pp.121–150), proposing a collision problem related to the 3-body problem, which might result in "an analog computation of a non-recursive function".

7. Meaning What We Say - and Computing What We Mean

When Ludwig Wittgenstein says (*Tractatus Logico-Philosophicus*, Proposition 1):

The world is everything that is the case.

he is setting out on a path consistent with longstanding intuitions concerning the role of language and symbolic representation. It is an intuition which predates the scientific era, one dignifying our representations of the world via language and picture, granting them meaning relevant to how things really are. Having developed this view, his concluding main Proposition 7:

Whereof one cannot speak, thereof one must be silent.

brings with it a granting of the world an epistemological, maybe even ontological, elusiveness evidenced by the failure of our representations — of the rule of 'facts'.

In *Alan Turing - His Work and Impact*, in her commentary *Turing, Wittgenstein and Types: Philosophical Aspects of Turing's 'The Reform of Mathematical Notation and Phraseology'* (1944–5), Juliet Floyd points (pp.250-253) to important parallels in Turing and Wittgenstein's respect for and understanding of the function of language:

Wittgenstein and Turing are often regarded, in a misleading caricature, as philosophical opponents. Wittgenstein is taken to be a humanistic philosopher

of meaning and 'forms of life', hostile to mathematical logic and the very idea of a Turing machine; Turing is taken to be a mechanistic or behaviouristic theorist of the mind, intent on reducing the concept of meaning to that of information. Neither picture is correct ...

Wittgenstein and Turing shared ... a particular sort of anti-reductionist attitude toward logical and conceptual analysis. On their view, it is the everyday, purposeful uses we humans make of language that crucially animate and frame the notions of meaning and information.

What Alan Turing had done in 1952 for the emergence of form in nature is capture particular examples via descriptions using the limiting mechanisms of the differential and integral calculus. He had also shown in 1936 how relatively modest extensions of the language used to describe Turing machines and their computations might lead (for example, via the Halting Set for a universal Turing machine), to levels of incomputability captured for us via the descriptions.

The intuition is that natural phenomena not only generate descriptions, but arise and derive from them. The potential for reconciling on the one hand the pre-Newtonian hybrid heuristic yielding a world captured informally via symbolic representation, with — on the other hand — the more modern scientific vision limited by modeling under controlled computational conditions, is striking. Out of such a reconciliation emerges a more coherent world, bringing together the historic experience of the power of descriptions for building a picture of reality, with an appropriate analysis of a logically parallel language and its computability theoretic infrastructure.

In mathematics, the experience of very real and novel features of structures arising via descriptions is a familiar one. And classically, one has no problem accepting such structural features as important and interesting characteristics, even if the logical complexity of the descriptions is beyond that acceptable within the computational domain. The formal notion, with which few outside of logic are very familiar, is that of *mathematical definability*. And yes, as well as the real world of ad hoc definitions, one which is important in mathematics and beyond, the logicians have developed a framework within one can discuss what, out there in the real world, one might hope to describe — or what is theoretically beyond description. There is also a working out of relationships between computation and description! This could be useful if more people knew what the logicians, a very few of them working in this area, were up to.

The 1930s work of the logician Alfred Tarski in formalizing a notion of definability of truth in a structure was a seminal development, helping build the framework, not just for a better understanding of basic algebraic and number theoretic structures, but for the better understanding via hierarchies of the relationship between definability and computability theoretic structure. See Tarski's 1931 paper *Sur les ensembles définissables de nombres réels*, I, in *Fundamenta Mathematicae*.


Hans Reichenbach's 1924 work (*Axiomatik der relativistischen Raum-Zeit-Lehre*) on axiomatizing special relativity was an early indication that 20th century science would test post-Newtonian strictures defining what a good scientific theory should deliver, bringing *definability* in science back into focus. The axiomatization of general relativity is a current project, carried forward by the Budapest group of Istvan Nemeti and Hajnal Andreka.

For the the ontological realism one might find a mathematical approach via invariance under comprehensive enough automorphisms a more appropriate. There is much

discussion of multiverses and alternative realities in relation to the physics. To properly constrain the possible alternatives one does need to filter in a comprehensive underlying causal structure, and a proper mathematical analysis of the possible representations in terms of an analysis of the potential automorphism groups. Anyone who has worked with this kind of mathematical challenge across a range of structures will know that 'many worlds' are nowhere as easy to manipulate into existence as one might expect. But more on this later. The basic question is how one models the basic underlying causal structure. For invariance a partial modeling may suffice. For failure of invariance, for ontological ambiguity, such as that encountered at the quantum level, one needs to be comprehensive. Note, that the sort of determination of natural characteristics — laws, constants, the geometry of space-time — that come to mind in this context, is unlikely to be algorithmic, or mechanical in any recognizable sense. And this challenges us to model the higher order mode of natural 'computation' via suitably generalized computability theoretic notions. The aim is to provide computational approaches to this higher world of real information, a mathematics with properties well enough devised and understood for us to substantiate its link with what we observe.

Of course, a definition in terms of basic causality can be viewed as a corresponding computation, and descriptions grounded in a logic appropriate to the context will deliver invariance of what they define. In this sense, a definition pinning down an invariant feature in nature does persuasively 'compute' that feature. We have seen this in terms of emergent form in nature, where the termite cathedral we saw earlier is there to be appreciated, an embodied example of a type-2 computation.

But unlike the Turing patterns, we do not have a description of the termite cathedral. Turing found his descriptions by paying attention to the specific chemical iterations underlying the pattern formation. And this is not an uncommon phenomenon in the real world. One often needs to take a more explicitly computational route, carefully adapted to the observed process, to provide a computational framework with an explanatory and



efficiently computational role. And it is this which the various notions of higher finite type are designed to tackle. For a history and review of notions, see John Longley's article *Notions of Computability at Higher Types I*, in *Logic Colloquium 2000*, Lecture Notes in Logic 19, ASL (2005), pp.32-142. Early work grew out of the classical case, with key contributors being Stephen Kleene and Gerald Sacks.

As Longley describes though, the notions have multiplied, and the computational frameworks confirms much of what we observe regarding the uncertainties attendant on the bringing of the mathematics of turbulence, emergence or big data into our classical computational comfort zone. Regarding the conceptual robustness we have come to expect, Longley writes:

It is ... clear that very many approaches to defining higher type computability are possible, but it is not obvious a priori whether some approaches are more sensible than others, or which approaches lead to equivalent notions of computability. In short, it is unclear in advance whether at higher types there is really just one natural notion of computability (as in ordinary recursion theory), or several, or no really natural notions at all.

The definitive book on the subject promises to be Longley and Normann's book.

8. A Computationally Host for Definability

Our underlying theme has been the extent to which mathematics can give substance and informative structure to emergent intuitions about how the world works. The background has been Alan Turing's corpus callosal joining together of shoe and star gazing. Above we commented on how Turing's discovery (along with Alonzo Church) of incomputability connected with Nassim Taleb's insights concerning economics and 'randomness'. In his 1939 paper on *Systems of Logic Based on Ordinals* (Proc. London Math. Soc. (2) 45 (1939), 161-228) we see how Turing grapples with incompleteness and incomputability, using computable ordinals during the deconstruction of mathematical phase transitions, while delivering observations on intuition and ingenuity — echoing by those of Jacques Hadamard (1945), based on Henri Poincaré's *Société de Psychologie* lectures on creative thinking and free will. Alan Turing's 1952 examples of mathematical descriptions of morphogenic emergence anticipated later developments across a whole spectrum of disciplines, and gave us a key insight into the role of definability in structuring information, while helping us understand its elusive computational character.

There exists a mathematical basis for the contemporary content of these anticipations of the role of definitions, definability and their bearings on what we can compute. This is informed by the buzz of current interest in big data, emergence and the familiar fuzzy thinking, based in statistical approximation, that makes us human. We keep in mind John Longley's questioning of the likelihood of us being able to establish a Church-Turing thesis above the type-one level.

Douglas Hofstadter and Emmanuel Sander have written a whole book (*Surfaces and Essences — Analogy as the Fuel and Fire of Thinking*, from 2013) about definability — craftily concealed within a detailed focus on the role analogies play in our thinking. They refer (p.184) to “the recognition of a correspondence between a newly-minted mental structure and an older one — in short, the construction of an analogical bridge.” In

emphasizing — what we recognize as — a computational difference between an analogy and a categorization, they describe:

— *the seemingly trivial case of the recognition of a cup as a cup. Suppose you are at a friend's house and want to fix yourself a cup of tea. You go into the kitchen, open a couple of cupboards, and at some point you think, "Aha, here's a cup." Have you just made an analogy? If, like most people, you're inclined to answer, "Obviously not — this was a categorization, not an analogy!", we would understand the intuition, but we would propose another point of view. Indeed, there is an equally compelling "analogy" scenario, in which you would have just constructed inside your head a mental entity that represents the object seen in your friend's cupboard. In this scenario, you would have created a mental link between that mental representation and a pre-existing mental structure in your head — namely, your concept named "cup". In short, you would have created a bridge linking two mental entities inside your head.*

What is it underpins this linking of previous and current mental entities? The authors clearly see these 'entities' as being far from simple. There are contextual of relationships at play, ones which are familiar and arising from experience. It is the functionality of the cup in a somewhat elaborated mental picture which supports the recognition. Something which is not simple is brought into focus by detailed relationships over which we exercise a degree of mental ownership. How can start to model such relationships and what they deliver. How can we carry out the sort of structural comparisons between diverse contexts that support such analogy forming. We suspect a corresponding hierarchical construct, involving increasingly complex descriptions — and based in relationships that we have a basic grip of. What we do have a secure grip of is the level at which *categorization*, supported by *computational* criteria, comes into play. And the analogy which comes into play *for us*, is that of the emergent higher type (mathematically type-2) patterns in nature, which Turing deconstructed in terms of computable reaction-diffusion elements, and hence found descriptions of via appropriate differential equations. Of course, networks and their pattern-forming are extensively investigated nowadays according to all their fascinating specificities, and high-impact journals trace their surprising intricacies — as one does as one encounters real emergence. But surely there is more to do than experience the botanical garden, and make interesting observations. What are the overarching informational structures which host these wonders?

Faced with the incomputability of important relations, Turing introduced the notion of an oracle Turing machine, in order to compare degrees of incomputability. At a stroke, this gave us a structuring of information according to its relative computational complexity. The oracle machine computes just like the standard machine described earlier, but has an extra facility for occasionally asking a question about a finite part of a given real number provided as the oracle — the oracle being a repository of auxiliary information to be used during any given computation.

This idea gave us a model of how we compute using data given to us from unknown sources. It not only gave us a way of comparing computability of reals. It could also be used to model real world causality, where the mathematical law for computing applications of the law from one state to another was computable. This was a framework within which Newtonian computable laws fitted nicely, at the basic level. Looked at globally, one


obtained a model of computable content of structures, based on partial computable (p.c.) functionals over the real numbers. So any physical context, characterised by computable relationships between pairs of reals representing particular physical states, could be presented as an instantiation of part of the resulting mathematical structure — often termed the Turing universe. Strangely, despite Turing's later interest in interactive computation, he never returned to the oracle machine.

In 1948, Emil Post mathematically prepared the Turing universe for further investigation by grouping reals together into equivalence classes (called degrees of unsolvability) of inter-computable reals, the degrees ordered by the induced ordering got from relative computation. The resulting structure, over the years, was discovered to have a very rich infrastructure. Mathematically, what was seen as a high degree of pathology provided the means for defining a rich array of relations and individuals. The complexity, of course, was just what one might expect from a mathematical structure capable of reflecting causal structure from different physical contexts. The definable relations formed a mathematical counterpart of the richness of emergent phenomena in the real world.

By the 1960s, the theory of this basic structure — now termed the Turing degrees — was becoming a very active area of pure research, one of formidable technical complexity. And this leads us to what has become known as *Hartley Rogers' Programme*, originating with a paper of his on *Some problems of definability in recursive function theory*, based on an invited talk at the 1965 Logic Colloquium in Leicester. The fundamental problem implicit in the paper and talk, was that of characterizing the Turing definable/invariant relations over the structure of the Turing degrees. The intuition was that these definable relations are key to pinning down how higher order relations on the real world can appear to be computed. An important observation being that the breakdown of definability potentially underpins ontological and epistemological ambiguities in the real world — such as confused mental states, or uncertainty at the quantum level in physics.

This aspect of physical computation — the Turing universe and emergent relations — is treated in more detail elsewhere (for example, in *The Mathematician's Bias, and the Return to Embodied Computation*, in H. Zenil, *A Computable Universe: Understanding and Exploring Nature as Computation*, World Scientific, 2012). The key message is that in the real world one may describe global relations in terms of local structure, so capturing the computational basis of large-scale phenomena. And that mathematically this can be formalized as definability, or more generally invariance under automorphisms, over structure based on relative computation.

This is just one important ingredient in liberating us from the computational flatness of real world information. Its role is particularly appropriate to the scientific context, with clearly observed counterparts to the computability of basic relations, representable at accessible type levels. Here, the correspondence



between the levels of abstraction,

over which Hofstadter and Sander's analogies play out, fits well with the mathematical type structure. Scientifically, one may need to focus on local instantiations of context, but the relationship to the mathematics is both useful and informative.

Over the sort of information the human brain deals with, one has difficulty isolating closed systems, and the selection of pertinent relations is essentially heuristic and shifting in particulars. Analogies weaken to metaphors, and the grasp of computers becomes at best statistical, made fallible by the patchy reducibility of informational type. What we are left with is a better understanding of the underlying patterns of knowing and not knowing, and appreciation of the appropriate relationship between the algorithmic and more characteristically *human* thought. Most importantly, we properly situate digital computation within the wider computational context, and rediscover the role of language in navigating our increasingly real sense of the typing of embodied information.

References

1. Samson Abramsky, *Two Puzzles About Computation*, In: S.B. Cooper and J. van Leeuwen, 2013, pp. 53-57.
2. Cristian S. Calude, Karl Svozil, *Quantum Randomness and Value Indefiniteness*, *Advanced Science Letters* **1** (2), 2008, 165-168.
3. S. B. Cooper, *Computability Theory*, Chapman & Hall/CRC, Boca Raton, London, New York, Washington, D.C., 2004.
4. S. B. Cooper, *From Descartes to Turing: The Computational Content of Supervenience*, In: *Information and Computation* (eds. M. Burgin, G. Dodig-Crnkovic), World Scientific Publishing Co., 2011, 107-148.
5. S. Barry Cooper, Jan van Leeuwen (eds.), *Alan Turing: His Work and Impact*, Elsevier, Amsterdam, London, New York, 2013.
6. Antonio Damasio, *The Feeling Of What Happens: Body and Emotion in the Making of Consciousness*, Harcourt Brace, London, 1999.
7. Martin Davis, *The Universal Computer: The Road from Leibniz to Turing*, W.W. Norton, New York, London, 2000.
8. Daniel Dennett, *Turing's "Strange Inversion of Reasoning"*, in S.B. Cooper and J. van Leeuwen, 2013, pp. 569-573.
9. David Deutsch, *Quantum theory, the Church-Turing principle and the universal quantum computer*, *Proceedings of the Royal Society of London, Series A*, **400**, 97-117.
10. Rodney G. Downey, Denis R. Hirschfeldt, *Algorithmic Randomness and Complexity*, Springer, Heidelberg, London, New York, 2010.
11. Albert Einstein, *Out of My Later Years*, Philosophical Library, New York, 1950.
12. Luciano Floridi, *The Philosophy of Information*, Oxford University Press, Oxford, 2011.
13. Kurt Gödel, *Russell's Mathematical Logic* (1944), in Schilpp (ed.), *The Philosophy of Bertrand Russell*, 3rd ed., Tudor, New York, 1951, 123-153. Repr. in Benacerraf and Putnam (eds), *Philosophy of Mathematics*, 2nd ed., Cambridge University Press, Cambridge, 1983, 447-469; & in Pears (ed.), *Bertrand Russell: A Collection of Critical Essays*, Anchor Books, Garden City, New York, 1972, 192-226.
14. Jacques Hadamard, *The Psychology of Invention in the Mathematical Field*, Princeton Univ. Press, Princeton, 1945.

15. Andrew Hodges, *Alan Turing: The Enigma*, Vintage, London, Melbourne, Johannesburg, 1992.
16. Jaegwon Kim, *Mind in a Physical World*, MIT Press, 1998.
17. Jaegwon Kim, *Physicalism, or Something Near Enough*, Princeton University Press, Princeton, Oxford, 2005.
18. Georg Kreisel, *Church's Thesis: a kind of reducibility axiom for constructive mathematics*. In: *Intuitionism and proof theory: Proceedings of the Summer Conference at Buffalo N.Y. 1968* (A. Kino, J. Myhill, R. E. Vesley, eds.), North-Holland, Amsterdam, London, 1970, pp. 121–150.
19. Thomas S. Kuhn, *The Structure of Scientific Revolutions*, Third edn., University of Chicago Press, Chicago, London, 1996.
20. P.S. de Laplace. *Essai Philosophique sur les Probabilités*. English trans. by F.W. Truscott and F.L. Emory, Dover, New York, 1951.
21. G.W. Leibniz, *La Monadologie* (1714). English translation by G.M. Ross, 1999.
22. Axel Leijonhufvud, *The nonlinear path of macroeconomic dynamics*, in *Computable, Constructive and Behavioural Economic Dynamics: Essays in honour of Kumaraswamy (Vela) Velupillai* (ed. S. Zambelli), Routledge, Abingdon, New York, 2010, pp. 62-69.
23. John Longley, *Notions of computability at higher types I*, in: Cori, Razborov, Todorovic, Wood (eds.) *Logic Colloquium 2000*, A K Peters, Wesley, MA, 2005.
24. W. McCulloch and W. Pitts, *A logical calculus of the ideas immanent in nervous activity*, Bull. Math. Biophys. **5** (1943), 115-133.
25. Iain McGilchrist, *The Master and his Emissary: The Divided Brain and the Making of the Western World*, Yale University Press, New Haven, London, 2009.
26. Kanti V. Mardia, S. Barry Cooper, *Alan Turing and Enigmatic Statistics*, Bull. Brazilian Section Int. Soc. for Bayesian Analysis, **5** (2) (2012), 2-7.
27. John Stuart Mill, *A System of Logic* (1843), Longmans, Green, Reader & Dyer, London, 8th ed., 1872.
28. Roger Penrose, *Quantum physics and conscious thought*, in *Quantum Implications: Essays in honour of David Bohm* (B. J. Hiley and F. D. Peat, eds) , Routledge & Kegan Paul, London, New York, 1987, pp.105–120.
29. Steven Pinker, *How the Mind Works*, W.W. Norton, New York, 1997.
30. Emil Post, *Degrees of recursive unsolvability: preliminary report* (abstract). Bull. Amer. Math. Soc. , **54** (1948), 641–642; reprinted in *Solvability, Provability, Definability: The Collected Works of Emil L Post* (ed. M. Davis), Contemporary Mathematicians, Birkhäuser Boston Inc., Boston, MA, 1994.
31. Hilary Putnam, *Minds and Machines*, reprinted in Putnam (1975).
32. Hilary Putnam, *The Nature of Mental States*, reprinted in Putnam (1975).
33. Hilary Putnam, *Mind, Language, and Reality*, Cambridge University Press, Cambridge, 1975.
34. Hilary Putnam, *Representation and Reality*, The MIT Press, Cambridge, MA, 1988.
35. Hilary Putnam, *Why Functionalism Didn't Work*. In: Putnam, *Representation and Reality*, pp. 73-89.
36. Hilary Putnam, *Is the Causal Structure of the Physical Itself Something Physical?* in *Realism with a Human Face*, Harvard Univ. Press, Cambridge, MA, 1990, p.89.
37. Hans Reichenbach, *Axiomatik der relativistischen Raum-Zeit-Lehre*, 1924. English translation: *Axiomatization of the theory of relativity*, University of California Press, 1969.

38. Hartley Rogers, Jr., *Some problems of definability in recursive function theory*. In *Sets, Models and Recursion Theory* (Proc. Summer School Math. Logic and Tenth Logic Colloq., Leicester, 1965), North-Holland, Amsterdam, 1967, pp. 183–201.
39. Bertrand Russell, *The Principles of Mathematics*, Cambridge University Press, Cambridge, 1903.
40. P. Smolensky, *On the proper treatment of connectionism*, Behavioral and Brain Sciences **11** (1988), 1-74.
41. Robert I. Soare, *Recursively Enumerable Sets and Degrees*, Springer-Verlag, Berlin, Heidelberg, New York, London, Paris, Tokyo, 1987.
42. Nassim Taleb, *The Black Swan: the Impact of the Highly Improbable*, Allen Lane, London, 2007.
43. Alfred Tarski, *Sur les ensembles définissables de nombres réels*, I, Fundamenta Mathematicae, **17** (1931), 210-239.
44. Alfred Tarski, *On Definable Sets of Real Numbers*, translation of Tarski (1931) by J.H. Woodger. In A. Tarski, *Logic, Semantics, Metamathematics*, second edition, ed. by J. Corcoran. Hackett, Indianapolis, 1983, pp. 110-142.
45. A.M. Turing, *On computable numbers, with an application to the Entscheidungsproblem*, Proc. London Math. Soc. (2) **42** (1936), pp. 230– 265; reprinted in Cooper and van Leeuwen, pp. 16-41.
46. A.M. Turing, *Systems of logic based on ordinals*, Proc. London Math. Soc. (2) **45** (1939), pp. 161-228; reprinted in Cooper and van Leeuwen, 2013, pp. 151-197.
47. 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 Cooper and van Leeuwen, 2013, pp. 501-516.
48. A.M. Turing, *The chemical basis of morphogenesis*, Phil. Trans. R. Soc. London B **237** (1952), pp. 37-72; reprinted in Cooper and van Leeuwen, 2013, pp. 689-722.
49. A.M. Turing, *Solvable and Unsolvable Problems*, Science News **31**, 1954, 7-23; reprinted in Cooper and van Leeuwen, pp. 322-331.
50. K. Vela Velupillai, *The unreasonable ineffectiveness of mathematics in economics*, Cambridge Journal of Economics, **29** (2005), 849-872; reprinted in *Computable Economics*, pp.723-746.
51. K. Vela Velupillai, Stefano Zambelli, Stephen Kinsella (eds.), *Computable Economics*, The Int. Library of Critical Writings in Economics, Edward Elgar Publishing Ltd., Cheltenham, Northampton, MA, 2011.
52. Steven Weinberg, *Is the Universe a Computer?*, The New York Review of Books, 24 October, 2002.
53. Ludwig Wittgenstein, *Tractatus Logico-Philosophicus (TLP)*, 1922, C.K. Ogden (trans.), Routledge & Kegan Paul, London. Originally published as *Logisch-Philosophische Abhandlung*, in *Annalen der Naturphilosophische*, XIV (3/4), 1921.