

Preprint from <http://www.philosophyofinformation.net>

This is a preprint of a paper accepted for publication in

Synthese (Springer)

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

It is a publisher's requirement to display the following notice:

The documents distributed by this server have been provided by the contributing authors as a means to ensure timely dissemination of scholarly and technical work on a noncommercial basis. Copyright and all rights therein are maintained by the authors or by other copyright holders, notwithstanding that they have offered their works here electronically. It is understood that all persons copying this information will adhere to the terms and constraints invoked by each author's copyright. These works may not be reposted without the explicit permission of the copyright holder.

In the case of Springer, but not of IGI Publishing, it is the publisher's requirement that the following note be added:

"An author may self-archive an author-created version of his/her article on his/her own website and his/her institution's repository, including his/her final version; however he/she may not use the publisher's PDF version which is posted on www.springerlink.com. Furthermore, the author may only post his/her version provided acknowledgement is given to the original source of publication and a link is inserted to the published article on Springer's website. The link must be accompanied by the following text: "The original publication is available at www.springerlink.com."

Against Digital Ontology

Luciano Floridi^{1,2}

¹Research Chair in Philosophy of Information and GPI, University of Hertfordshire; ²St Cross College and IEG, University of Oxford.

Address for correspondence: School of Humanities, University of Hertfordshire, de Havilland Campus, Hatfield, Hertfordshire AL10 9AB, UK; l.floridi@herts.ac.uk

Abstract

The paper argues that *digital ontology* (the ultimate nature of reality is digital, and the universe is a computational system equivalent to a Turing Machine) should be carefully distinguished from *informational ontology* (the ultimate nature of reality is structural), in order to abandon the former and retain only the latter as a promising line of research. Digital vs. analogue is a Boolean dichotomy typical of our computational paradigm, but digital and analogue are only “modes of presentation” of Being (to paraphrase Kant), that is, ways in which reality is experienced and/or conceptualised by an epistemic agent at a given level of abstraction. A preferable alternative is provided by an informational approach to structural realism, according to which knowledge of the world is knowledge of its structures. The most reasonable ontological commitment turns out to be in favour of an interpretation of reality as the totality of structures dynamically interacting with each other. The paper is the first part (the *pars destruens*) of a two-part piece of research. The *pars construens*, entitled “A Defence of Informational Structural Realism”, is forthcoming in *Synthese*.

Keywords

Analogue; continuous; digital; digital ontology; digital physics; discrete; informational structural realism; Kant’s antinomies; structural realism.

1. Introduction

In recent years, the age-old question about the discrete vs. continuous nature of reality¹ has been recast in the more fashionable terms of digital vs. analogue ontology. As such, it has enjoyed a remarkable revival. However, as will become clear in due course, this is a typical case of old wine in new bottles. And I shall argue that Kant's conclusion, reached against the more classic, old dichotomy in the context of the "antinomies of pure reason" (Kant [1998], A 434-5/B 462-3), has lost none of its value when applied to the most recent reformulation of the same alternative.

The paper is structured into four other main sections. Section two provides a brief introduction to what is known as digital ontology. In section three, the longest, a new thought experiment is introduced in order to show that digital (discrete) vs. analogue (continuous) is a Boolean dichotomy typical of the computational paradigm of our age, but both digital and analogue are only "modes of presentation of Being" (to paraphrase Kant), that is, ways in which reality is experienced and conceptualised by an epistemic agent, at a given *level of abstraction* (LoA). They do not pick up some knowledge- or LoA-independent properties, intrinsic to the external world. Although this conclusion applies both to digital and to analogue ontologies, the paper concentrates mainly on the criticism of digital ontology (hence its title) because its constructive goal is to clear the ground for a defence of informational ontology, which is often confused with digital ontology. Section four further clarifies the argument by answering three potential objections. Section four, which concludes the paper, provides a positive note and an explanation of the more constructive rationale for the argument developed in the previous sections. The paper is the first part (the *pars destruens*) of a two-part piece of research. The *pars construens* is developed in a separate article, entitled "A Defence of Informational Structural Realism", which is forthcoming in *Synthese*.

¹ Holden [2004] provides an enlightening and insightful analysis of the modern debate, to which I'm indebted. I have also relied on the excellent article by Lesne [2007].

2. What is Digital Ontology? It from Bit

Konrad Zuse² is acknowledged by many as the father of digital ontology.³ According to him and to digital ontologists in general:

- 1) the nature of the physical universe (time, space and every entity and process in space-time) is ultimately discrete.

This thesis, with which we shall be engaged for the rest of the paper, may be accompanied by (a selection of) three other, related theses:

- 2) the physical universe can be adequately modelled by discrete values like the integers;
- 3) the evolution (state transitions) of the physical universe is computable as the output of a (presumably short) algorithm; and
- 4) the laws governing the physical universe are entirely deterministic.

Theses (1) and (2) give away the neo-Pythagorean nature of digital ontology (Steinhart [2003]): reality can be decomposed into ultimate, discrete *indivisibilia*. Philosophers still disagree on the precise definition of “digital” and “analogue”,⁴ but they accept that a necessary feature of what it means for something to be digital is that of being discrete, and this will suffice for the purposes of this paper.

Thesis (3) interprets the neo-Pythagorean ontology in computational terms: the ultimate, discrete *indivisibilia* are actually computable digits, while elegance and Ockham’s razor inclines digital ontologists to favour an algorithmic theory as simple as possible (see Feynman [1992], quoted below in § 3.3.2). Thus, a digital ontology “[...] is based on two concepts: bits, like the binary digits in a computer, correspond to the most microscopic representation of state information; and the temporal evolution of state is a digital informational process similar to what goes on in the circuitry of a computer

² Zuse is famous for having constructed the first fully operational program-controlled electromechanical binary calculating machine (the Z3) in 1941, see Zuse [1993].

³ *Digital ontology* is also known as *digital physics* or *digital metaphysics* and *digital philosophy*, and has a scientific counterpart in *digital physics*, see Steinhart [1998] for an introduction and Steinhart [2003] for a review chapter. For a recent bibliography see <http://digitalphysics.org/Publications/>

⁴ The debate between Goodman [1968] and Lewis [1971] on the actual nature of the “digital” has been recently revisited by Müller [forthcoming]. On the unnecessary distinction between discrete and discretised system see Lesne [2007]: “It now appears that there is no reason to make a fundamental distinction between discrete and discretized systems: an object seems to be intrinsically discrete, even isolated, only if we choose the proper glasses” (p. 14). The “glasses” are interpreted more formally in this paper in terms of levels of abstraction.

processor.” (Fredkin [2003b], 188). In a nutshell, “we are run by a short algorithm” (Schmidhuber [1997], p. 205).

As for thesis (4), this is presented by supporters of digital ontology⁵ as a direct consequence of theses (1)-(3) and explicitly related by them to Einstein’s reluctance to accept the conclusion that the universe might be intrinsically probabilistic. The suggestion is that the analysis of physical laws in terms of deterministic state transitions may be made compatible with the ostensibly probabilistic nature of quantum phenomena, while being sufficiently flexible to overcome other criticisms.⁶

The position that unifies most supporters of digital ontology is summarised in what is known as Zuse Thesis:

ZT) “the universe is being deterministically computed on some sort of giant but discrete computer” (Zuse [1969]).

The computer referred to in ZT could be a *cellular automaton*. This is argued by Zuse [1967], also on the basis of Von Neumann [1966], by Fredkin [2003b] and, more recently, by Wolfram [2002].⁷ Indeed, a variant of ZF, which is less general, is known as Fredkin-Zuse Thesis: “The Universe is a cellular automaton” (Petrov [2003]). Alternatively, the computer in ZT could be a *universal Turing machine*, as suggested by Schmidhuber [1997], who in turn acknowledges his intellectual debt to Zuse himself; or a *quantum computer*, as proposed more recently by Lloyd [2006]. Other well-known

⁵ Fredkin [1992] presents the point with usual clarity: “Uncertainty is at the heart of quantum mechanics. Finite Nature requires that we rule out true, locally generated randomness because such numbers would not, in this context, be considered finite. [...] The deterministic nature of finite digital processes is different in that it is unknowable determinism. From within the system an observer will never be able to know very much about the true microscopic state of that system. Every part of space is computing its future as fast possible, while information pours in from every direction. The result is the same as caused by the apparent randomness of quantum mechanical processes.”. See also 'T Hooft [2002] and 'T Hooft [2005].

⁶ For example, digital ontology seems inconsistent with Bell's theorem, but a solution to bypass the problem, known as pre-determinism ('T Hooft [2005]), is based on what Bell himself acknowledged as a possibility: if a model is completely deterministic, then even the experimenter's decision to measure some components of the spins is entirely pre-determined, so she could not have decided to measure anything else but what she did measure.

⁷ To be precise, it has been argued that is unclear whether Wolfram means to support the view that the universe is a *classical* cellular automaton. Wolfram acknowledges Zuse’s and Fredkin’s work on pp. 1026-1027, but only very briefly and states that “no literal mechanistic model can ever in the end realistically be expected to work.” I follow Edwin Clark’s interpretation in reading this as a rejection of classical cellular automata. Wolfram seems to have in mind something slightly different “[...] what must happen relies on phenomena discovered in this book – and involves the emergence of complex properties [...]”. The potential differences between Fredkin and Wolfram on this issue, however, are not significant for the discussion of the tenability of a digital ontology.

proponents of versions of ZT include David Chalmers [1996], the Nobel laureate Gerard 'T Hooft [1997], and Gregory Chaitin [2005]. The latter has explicitly interpreted digital ontology as a contemporary development of Pythagoras' metaphysics, Democritus' atomism and Leibniz's monadology. Indeed, in the *Critique of Pure Reason*, Kant called any philosopher who holds a "discrete" ontology a Monadologist.

The overall perspective, emerging from digital ontology, is one of a metaphysical monism: ultimately, the physical universe is a gigantic digital computer. It is fundamentally composed of digits, instead of matter or energy, with material objects as a complex secondary manifestation, while dynamic processes are some kind of computational states transitions. There are no digitally irreducible infinities, infinitesimals, continuities, or locally determined random variables. In short, the ultimate nature of reality is not smooth and random but grainy and deterministic.⁸

Since in the rest of this paper I shall be concerned only with the digital vs. analogue nature of reality, namely thesis (1) to be found at the beginning of this section, let me conclude this brief presentation of digital ontology with a final comment about the computational nature of the physical universe. It concerns an important distinction that needs to be kept in sight and to which I shall briefly return in the conclusion.

We have seen that digital ontologists tend to be computationally-minded and hence subscribe to some version of the pancomputational thesis, according to which the physical universe is a computational system of some kind, where the kind is irrelevant as long as the preferred models are computationally equivalent, like Turing machines, cellular automata, quantum computers or indeed even recurrent neural networks.⁹ However, digital ontology and pancomputationalism are independent positions. Famously, Wheeler supported the former but not (or at least not explicitly) the latter. As he wrote:

⁸ "A fundamental question about time, space and the inhabitants thereof is 'Are things smooth or grainy?' Some things are obviously grainy (matter, charge, angular momentum); for other things (space, time, momentum, energy) the answers are not clear. Finite Nature is the assumption that, at some scale, space and time are discrete and that the number of possible states of every finite volume of space-time is finite. In other words Finite Nature assumes that there is no thing that is smooth or continuous and that there are no infinitesimals." Fredkin [1992].

⁹ I do not know anyone supporting this position, perhaps because artificial neural networks are not usually analysed algorithmically, but it is available insofar as any algebraically computable function can be expressed as a recurrent neural network, see Hyötyniemi [1996] and Siegelmann [1998].

It from bit. Otherwise put, every ‘it’ – every particle, every field of force, even the space-time continuum itself – derives its function, its meaning, its very existence entirely – even if in some contexts indirectly – from the apparatus-elicited answers to yes-or-no questions, binary choices, *bits*. ‘It from bit’ symbolizes the idea that every item of the physical world has at bottom – a very deep bottom, in most instances – an immaterial source and explanation; that which we call reality arises in the last analysis from the posing of yes–no questions and the registering of equipment-evoked responses; in short, that all things physical are information-theoretic in origin and that this is a *participatory universe* (Wheeler [1990], 5).

Physical processes in Wheeler’s participatory universe might, but need not, be reducible to computational state transitions. On the other hand, pancomputationalists like Lloyd [2006], who describes the universe not as a Turing Machine but as a quantum computer, can still hold an analogue or hybrid¹⁰ ontology. Laplace’s demon, for example, is an analogue pancomputationalist. And informational ontologists like Sayre [1976], or myself (Floridi [2004] and Floridi [forthcoming]) do not have to embrace either a digital ontology or a pancomputationalist position as described in Zuse Thesis. The distinction between digital ontology, informational ontology and pancomputationalism is crucial in order to understand the strategy of this paper, which is to criticise digital ontology in order to make room for an informational approach to structural realism, defended in Floridi [forthcoming], while not committing it to pancomputationalism. I will return to this point in the conclusion. With this clarification in mind, we can now turn to the objections against digital ontology.

2.1. Digital Ontology: From Physical to Metaphysical Problems

When discussing digital ontology, two separate questions arise:

- a) whether the physical universe might be adequately modelled digitally and computationally, independently of whether it is actually digital and computational in itself; and
- b) whether the ultimate nature of the physical universe might be actually digital and computational in itself, independently of how it can be effectively or adequately modelled.

¹⁰ Hybrid computers comprise features of analog computers and digital computers. The digital component normally serves as the controller and provides logical operations; the analog component normally serves as a solver of differential equations. Plagiarism disclaimer: the previous definition is the source of the corresponding definition in *Wikipedia*, not vice versa.

The first is an empirico-mathematical question that, so far, remains unsettled. I shall say a bit more about it in this section, but the rest of the paper is not concerned with it. The second is a metaphysical question that, in the rest of the paper, I hope to show to be ill-posed and hence, when answered, to be mis-applied.

Answers to (a) and (b) are often found intertwined. The following passage by Edward Fredkin, one of the earnest supporters of digital ontology, provides a good example of a unified synthesis:

[digital ontology] is a totally atomistic system. Everything fundamental is assumed to be atomic or discrete; and thereby so is everything else. In physics, DP [digital philosophy, what has been called in this paper digital ontology] assumes that space and time are discrete. There are two mathematical models of such systems. The first is Diophantine analysis; the mathematics of the integers. The second is automata theory; the mathematics of digital processes. We choose the latter as it also has the property of explicitly representing a discrete temporal process, while the mathematics of the integers simply establishes a set of true theorems and thus can represent implicitly only temporal processes. Tremendous progress in the sciences followed the discovery of the calculus (and partial differential equations) as a way of mathematically representing physical-temporal relationships. But we look elsewhere with respect to the most fundamental models of physical processes. What we must demand of DP is the eventual ability to derive, from our DP models of fundamental processes, the same mathematical equations that constitute the basis of science today. Conway's Game of Life [a famous cellular automaton, my added comment] is a good example of a simple digital system and the consequent emergent properties. We arbitrarily assume that DP represents state by patterns of bits, as is done in ordinary computers. All of the fundamental transformations we can do with bits in a computer are really a subset of what mathematics can do with the integers. [...] The bits of DM [digital mechanics] exist at points in a regular digital spacetime, where each point contains one bit of information. We think of spacetime as digital since it is made up only of points located where all of the coordinates are integers. Automata theory and computer science lead us to believe that the representation of state by bits imposes no limitations beyond the fact that everything is ultimately quantized. Computers and their software are the most complex things ever made by man. However, computation is based on the simplest principles ever discovered. Our world is complex and we are looking for simple models that might be at the bottom. The principles of DP require us to find and start with the simplest possible models. Thus the unit of state is the bit, which is considerably simpler than a real number. (Fredkin [2003b], 190-191).

As the passage illustrates, the empirico-mathematical and the metaphysical position with respect to digital ontology are compatible and complementary. Consider the following way of interpreting Digital Ontology. Physical simulations or models may share the same ontology with their simulated or modelled systems. Thus, a wind tunnel, used to investigate the effects of wind-speed and flow around solid objects, is actually windy, more or less spacious, may contain several physical objects, and so forth. However, digital (computer) simulations or models *often* have a different ontology from their corresponding systems. A computational fluid dynamics simulation, used to model

and simulate the behaviour of flowing air, is neither windy nor wet in itself. However, I emphasised “often” because this is not always the case. For example, computer simulations are routinely used to test and debug other application programs, in which case simulator and simulated share the same digital ontology. Digital ontology may then be interpreted as arguing that one could have a digital model or simulation of the ultimate nature of the physical universe which ends up sharing the same digital ontology with the modelled system.

Although an answer to (a) could then pave the way to an answer to (b), it is useful to keep the two issues separate because they face different challenges.

The empirico-mathematical position seeks to answer question (a) and it is weaker, and hence more defensible, than the metaphysical position, which seeks to answer question (b), because it may avoid any ontological commitment to the ultimate nature of reality. That a system might be modelled and simulated digitally, e.g. as a cellular automaton, does not imply that the intrinsic nature of that system is digital. This lack of ontological commitment may be an advantage and, occasionally, the weaker position seems to be all that digital ontologists wish to hold. Toffoli [2003], for example, who sympathizes with digital ontology, has proposed to treat Digital Ontology as a heuristically interesting line of research:

We argue that a nonfrivolous [sic] aspect of this Digital Perspective is its heuristic capacity: to help us guess which aspects of our understanding of nature are more “universal,” more robust, more likely to survive theoretical and experimental challenges. Behaviors that are substrate-independent—that can, for instance, thrive well on a digital support, even though they are traditionally imagined as taking place in a continuum—are especially promising candidates. (p. 147).

And Fredkin [online] himself suggests that

there are computer systems (cellular automata) that may be appropriate as models for microscopic physical phenomena. Cellular automata are now being used to model varied physical phenomena normally modelled by wave equations, fluid dynamics, Ising models, etc. We hypothesize that there will be found a single cellular automaton rule that models all of microscopic physics; and models it exactly. We call this field DM, for digital mechanics.

Both passages could easily be read as addressing only question (a).

If Digital Ontology is an answer to (a) and not (also) to (b), then, even if objections against the metaphysical value of digital ontology may be correct, nothing could be inferred from them with regard to the scientific tenability of digital physics weakly interpreted. Perhaps someone may wish to argue that the latter would be

damaged by a lack of ontological support but, personally, I doubt this would be of much consequence to the digital physicist.

So, let us suppose that digital ontologists would prefer to support their position as an answer to (a) and not to (b), in order to avoid metaphysical complications. Even the weaker position is not devoid of problems. In terms of empirical ontology, the majority of physicists either ignores or is positively sceptical about the value of the approach. In the words of a supporter of digital ontology:

Several speakers in this meeting [“The Digital Perspective” workshop, organized by Edward Fredkin, my note] express their optimism concerning the possibility to describe realistic models of the Universe in terms of deterministic ‘digital’ scenarios. Most physicists, however, are acutely aware of quite severe obstacles against such views. It is important to contemplate these obstacles, even if one believes that they will eventually be removed. ’t Hooft [2003], 349.

The models proposed by digital ontologists – when they are subject to testable experiments, at least in principle – show implications that are not easily reconcilable with many postulates and results commonly obtained in physics and with our current understanding of the universe. Here is a very simple illustration: Lloyd [2002] estimates that the physical universe, understood as a computational system, could have performed 10^{120} operations on 10^{90} bits (10^{120} bits including gravitational degrees of freedom) since the Big Bang. The problem is that if this were true, the universe would “run out of memory”:

To simulate the Universe in every detail since time began, the computer would have to have 10^{90} bits – binary digits, or devices capable of storing a 1 or a 0 – and it would have to perform 10^{120} manipulations of those bits. Unfortunately there are probably only around 10^{80} elementary particles in the Universe. (Ball [2002, June 3]).

A digital reinterpretation of contemporary physics may be possible in theory. After all, discrete systems can approximate continuous systems to increasing degrees of accuracy, so it is unlikely that any experiment could rule out the possibility that the world might be digital in itself. One might even argue that a digital ontology could be coherent with the many-worlds interpretation of quantum mechanics, insofar as the former satisfies what Goodman described as the multiple-realization characteristic of digital states.¹¹

¹¹ Schmidhuber [forthcoming] maintains that “computing all evolutions of all universes is much cheaper in terms of information requirements than computing just one particular, arbitrarily chosen evolution” (see version online at <http://www.idsia.ch/~juergen/everything/node3.html>). However, if there are no

The history of the world would then be more like a (discrete) chess game, in which every move generates a different state and hence a parallel universe, rather than a (continuous) football game, in which every tiny (as tiny as one may wish it to be) little difference, e.g. in the trajectory of the ball, would generate another world. What seems to be the case, however, is that, if digital ontology seeks to advance our understanding of the physical universe by providing a “conceptual strategy of looking at physics in terms of digital processes” (Fredkin [2003a]), then its success would represent a profound change in our scientific practices and outlook. Quite a bit of our contemporary understanding of the universe is firmly based not only on discrete but also on many analogue ideas (the real numbers, continuous functions, differential equations, Fourier transforms, waves, force fields, the continuum)¹² which seem to be difficult to replace entirely. Of course, this is not a final argument against digital ontology, for our analogue understanding of the universe may turn out to be digitally reinterpretable after all. But, as acknowledged by ’t Hooft in the quotation above, the burden of showing how it can actually be replaced, and why it should, is definitely on the shoulders of the digital ontologists. A quote from Einstein well highlights the overall difficulty:

I consider it quite possible that physics cannot be based on the field concept, i. e., on continuous structures. In that case, nothing remains of my entire castle in the air gravitation theory included, [and of] the rest of modern physics.¹³

If digital ontologists are right, then the “analogue” tradition that goes from Newton to Einstein will be deeply affected, as digital ontologists are willing to acknowledge.¹⁴ This is one of the reasons why, although contemporary particle physics and astrophysics increasingly depend on e-science,¹⁵ they are currently not digital-ontology-friendly.

sufficient particles for Lloyd’s computations, it is unclear how there might be sufficient particles for Schmidhuber’s.

¹² For a very instructive analysis of the “interplay between discrete and continuous behaviors and corresponding modelings in physics” see Lesne [2007].

¹³ Einstein, 1954, in a letter to Besso, quoted in Pais [2005], p. 467.

¹⁴ At least this is the way digital ontologists see the impact of their work and how it is perceived by some of their contemporaries, see for example the list of reviews of Wolfram [2002] available at <http://www.wolframscience.com/coverage.html> and http://www.math.usf.edu/~eclark/ANKOS_reviews.html

¹⁵ The term refers to any computationally intensive scientific research that relies on distributed network environments, very powerful processing capacities and very large data sets.

Steven Weinberg, another Nobel laureate in physics, has well expressed such lack of sympathy when talking about Wolfram's version of Zuse Thesis:

Wolfram himself is a lapsed elementary particle physicist, and I suppose he can't resist trying to apply his experience with digital computer programs to the laws of nature. This has led him to the view (also considered in a 1981 article by Richard Feynman) that nature is discrete rather than continuous. He suggests that space consists of a network of isolated points, like cells in a cellular automaton, and that even time flows in discrete steps. Following an idea of Edward Fredkin, he concludes that the universe itself would then be an automaton, like a giant computer. It's possible, but I can't see any motivation for these speculations, except that this is the sort of system that Wolfram and others have become used to in their work on computers. So might a carpenter, looking at the moon, suppose that it is made of wood. Weinberg [24 October 2002].

To summarise and simplify, it would be a mistake to confuse the *predicative* use of "digital" ("digital physics" is the study of the laws of the universe helped by digital-computational instruments) with its *attributive* use ("digital physics" is the study of the intrinsically digital-computational nature of the laws of the universe). A lot of contemporary physics is digital in the predicative sense, not in the attributive sense.

So much for question (a). Regarding question (b), namely whether the ultimate nature of the physical universe might be intrinsically digital and computational, does digital ontology fare any better with contemporary metaphysics? Is the latter any more "attributive-friendly"? In the rest of the paper, I shall argue that it is not. As I wrote earlier, the argument leads to a Kantian conclusion: it is not so much that reality in itself is not digital, but rather that, in a metaphysical context, the digital vs. analogue dichotomy is not applicable.

3. The Thought Experiment

Let me first provide an overall view of the argument and then of the thought experiment through which I will expound it.

If the ultimate nature of reality in itself is digital, this implies that it is either digital or analogue,¹⁶ so the premise can be refuted by showing that the disjunctive conclusion is mistaken. This can be achieved in two steps.

¹⁶ For the sake of simplicity, I shall treat the or as the logic disjunction, and hence as equivalent to asserting that reality in itself is either digital/discrete (grainy) or continuous/analogue (smooth) or hybrid (lumpy). Nothing depends on this simplification.

The first step consists in arguing that, even assuming that reality in itself is indeed digital or analogue, an epistemic agent, confronted by what appears to be an (at least partly) analogue world of experience, could not establish whether its source (that is, reality in itself as the source of the agent's experience or knowledge) is digital (or analogue).

One could object, however, that this first, epistemological step is merely negative, for it establishes, at best, only the unknowability of the intrinsically digital (or analogue) nature of reality lying behind the world of experience, not that reality in itself is not digital (or analogue). Independently of the epistemic access enjoyed by an agent – the objection continues – logic dictates that reality must be assumed to be digital/discrete (grainy) or continuous/analogue (smooth).

So the second step is more positive and ontological. It consists in showing that the initial concession, made in the first step, can be withdrawn: the intrinsic nature of reality does not have to be digital or analogue because the dichotomy might well be misapplied. Reality is experienced, conceptualised and known as digital or analogue depending on the level of abstraction (LoA) assumed by the epistemic agent when interacting with it. Digital and analogue are features of the LoA modelling the system, not of the modelled system in itself.

The negative result of the argument is that one cannot know whether reality is intrinsically digital or analogue not only because of inherent limitations in the type of epistemic access to reality that one may enjoy, but also because reality in itself might be just the wrong sort of thing to which these categories are being applied. The positive result of the argument is that there are ontologies – in particular those supported by ontic structural realism and by informational structural realism – that treat the ultimate nature of reality as relational. And since relations are neither digital nor analogue nor a combination of the two, the negative conclusion clears the ground from a potential confusion between digital and informational ontology and makes room for the development of the latter, as we shall see in section five.

In order to develop the arguments just outlined, I shall rely on a thought experiment. This is divided into four stages. At each stage an ideal agent will be introduced in order to make the argument more intuitive and vivid.

- 1) stage one sets up a scenario in which reality in itself is assumed to be either digital or analogue.

This satisfies the initial concession, made in the first, negative, epistemological step seen above, which assumes the world to be digital or analogue. At this stage, we shall need an ideal agent capable of showing that reality in itself is either digital or analogue. Call this agent Michael.

Recall now that the argument is that an epistemic agent, who is confronted by what appears to be an (at least partly) analogue world of experience, cannot establish whether reality in itself is digital or analogue. So

- 2) stage two is where the epistemic agent – which we have not specified yet – is provided with an analogue world of experience that is based on reality in itself, which, following the previous stage, is assumed to be digital or analogue.

Stage (2) is a simplification of what one of the supporters of digital ontology has aptly defined as “the stubborn legacy of the continuum” (Margolus [2003], p. 309). Epistemic agents experience the world as (at least partly, if not mainly) analogue. So an easy way of understanding stage (2) is by interpreting it as providing an analogue interface between reality in itself – which has been assumed to be digital or analogue at stage (1) – and an epistemic agent who experiences it (see stage (3)). Call the agent responsible for translating reality in itself (whether digital or analogue) into an analogue world of experience Gabriel.

Once Gabriel (i.e. the interface) has translated Michael’s digital or analogue reality into an analogue world, we move to:

- 3) stage three, where the epistemic agent is shown to be incapable of establishing whether the source (reality in itself, assumed as either digital or analogue in stage one) of the analogue, experienced world (obtained at stage two) is intrinsically digital or analogue.

This will be argued by using the method of abstraction (more on this in section 3.3.1). The epistemic agent observes his analogue world at an endless number of levels of abstraction, can never reach an end in his informational explorations, and hence can never know whether there is a digital ontology behind the observable world. Call the third, epistemic agent Rafael.

Once Rafael (the epistemic agent) is shown to be unable to establish the digital (or analogue) nature of reality in itself, one still needs to answer the objection according to which this conclusion, even if granted, shows nothing about the actual nature of reality in itself. So, what has been called above the second, positive, ontological step, requires

- 4) stage four, where the concession made at stage one is withdrawn, and it is shown that the alternative digital vs. analogue may easily be misapplied, when talking about reality in itself, because the latter could be neither, while it might be experienced as either digital or analogue depending on how the epistemic agent is related to it.

At this last stage, a fourth agent, call it Uriel, is needed in order to show that the same reality can be observed as being digital or analogue, depending on the epistemic position of the observer (Rafael) and the level of abstraction adopted.

Stage four concludes the thought experiment against digital ontology. A final note on the four agents might still be in order before seeing the details of the argument. These agents are ideal because they are endowed with the following boundless (i.e. always sufficient yet not infinite) resources:

- i) time, to be interpreted computationally as the number of steps required to achieve a specific task;
- ii) space, to be interpreted computationally as memory; and
- iii) accuracy, to be interpreted computationally as the degree of precision of an operation, which is assumed (the accuracy) to be increasable to whatever level may be required.¹⁷

Thus the four agents resemble Turing Machines. This is important because the digital ontologist, of all people, will hardly object to their assumption as perfectly conceivable. For a more colourful and vivid representation, I have proposed to call them Michael, Gabriel, Raphael and Uriel. The reader may recognise them as the four archangels and indeed some of their iconological properties turn out to be nicely consistent with their tasks in each stage, hence the choice. But of course nothing hangs on this, and the reader

¹⁷ For example, the degree of accuracy depends on how many figures or decimal places are used in rounding off the number. The result of a calculation or measurement (such as 13.429314) might be rounded off to three (13.429) or two decimal places. The first answer is more *accurate* than the second.

who finds the analogy or the names unhelpful is invited to translate the four agents into M, G, R and U.¹⁸

Let us now consider each stage and the overall argument in more detail.

3.1 Stage 1: Reality in itself is Digital or Analogue

Suppose our ideal agent, Michael, enjoys a God's eye view of reality in itself (the sort of approach criticised by Dewey) and has access to its intrinsic nature, what Kant called the *noumenal*. Whatever the noumenal reality is in itself – call it stuff – we assume that Michael is able to manipulate it and test whether it is digital/discrete or analogue/continuous. He first shapes this stuff into an ordered set, by applying some total ordering relation (e.g. “is to the right of”). For the sake of simplicity, we can represent the result geometrically, as a line. Michael then uses his sword to obtain a Dedekind cut (Dedekind [1963]).¹⁹ Following Dedekind's geometrical description,²⁰ Michael's sword could not be sharper: it has length but no thickness. When it cuts (i.e., intersects) the line, it divides it into two disjoint, non-empty parts, left and right. If the point at which the sword cuts the line belongs to either the left or the right half, then that point corresponds to a rational number. If the point belongs to neither (if neither subset of the rationals contains it), then it corresponds to an irrational number.²¹ Now, let us make the process iterative, so that the output of each operation becomes the input of the next operation. Suppose Michael is able to Dedekind-cut the line indefinitely, without his sword ever going through any empty space. Michael takes the right part, Dedekind-cuts it again, drops the left part, picks up the new right part, Dedekind-cuts it again, and so forth, never halting and never being able to drive his sword between two points (or corresponding numbers) without hitting another point (number). In this case, Michael's

¹⁸ Philosophy has already seen its fair share of demons, think of Socrates' inquisitive demon, Descartes' malicious demon, Laplace's deterministic demon, or Maxwell's intelligent demon. Time for the angels to have a comeback.

¹⁹ Intuitively, a Dedekind cut is a partition of the set of rational numbers into two non-empty subsets, in such a way as to uniquely define a real number.

²⁰ Dedekind's *Schnitt* is a cut in the geometric sense, that is, it is the intersection of a line with another line that crosses it, see Dedekind [1963].

²¹ To be more precise, the goal of a Dedekind's cut is to *construct* each real number as a pair (L, R) of sets of rationals such that every rational is in exactly one of the sets, and every rational in L is smaller than every rational in R. This couple is called a Dedekind cut. However, the distinction between “construction” and “correspondence” is not relevant here.

noumenal reality is dense²² and continuous, like the set of the real numbers. Suppose, on the contrary, that when Michael tries to Dedekind-cut the line, the point at which the sword intersects the line belongs to either the left or the right half. Or, to put it differently, Michael's sword is so sharp that it can miss the line by going through its gaps, which correspond to missing points (imagine a line in which the irrational numbers have been knocked off). Then Michael's noumenal reality is no longer continuous but discrete, and either still dense (like the rationals, which are closed under division), in which case he will not halt; or not even dense (like the integers, which are not closed under division), in which case he will halt. If reality is continuous or analogue, the output of Michael's process may consist, physically, of waves or fields, for example; mathematically, they would be equivalent to the real numbers. On the other hand, if reality is discrete, the output of Michael's process may consist of the atoms of any Democritean physics, of the natural numbers of the Pythagoreans, of Leibniz's indivisible monads or of the bits of the digital ontologists (see Figure 1).²³

At the end of this first stage, reality in itself can be assumed to be digital or analogue, with no further ontological commitment required, depending on the output of Michael's operation.

²² An ordered set (e.g. the rationals) is dense if any two of its elements have an element in between.

²³ For a similar idea, but applied to simple divisibility, see Kant's description of "annihilating composition in thought" and what Holden [2004] defines as metaphysical divisibility. Holden characterises metaphysical divisibility as divisibility which is logically possible, and distinguishes it from physical divisibility (physically possible), formal divisibility (based on distinguishable parts in space) and intellectual divisibility (a weaker notion of logical or metaphysical divisibility, based on the possibility of being able to imagine the division).




Figure 1 First stage of the thought experiment, Michael's sword.

3.2 Stage 2: The Stubborn Legacy of the Analogue

The second stage involves Gabriel and his message (see Figure 2). Recall that Gabriel is an ideal agent working as an interface between reality in itself (Michael the source) and the world as observed by the epistemic agent (Raphael the observer). If Michael's ontology is analogue, Gabriel uses it as his input to produce an analogue reality (henceforth also called the *system*). On the other hand, if Michael's ontology is digital, Gabriel puts it through a DAC (a Digital to Analogue Converter) to produce an analogue reality. In both cases, Gabriel might need an amplifier to produce his output, that is, an analogue system (a continuous world) like the one observed by the epistemic

agent. At the end of this second stage, Gabriel's output (his message) consists of an endlessly updated flow of analogue data. This is the analogue system made available to the observer, Raphael.




Figure 2 Second stage of the thought experiment, Gabriel's message.

3.3 Stage 3: The Observer's Analysis

Gabriel's message, that is, the analogue system comparable to the world surrounding us, is now observed by Raphael, the epistemic agent, who is on the other side of the screen, as it were (see Figure 2 and Figure 3). Raphael, like us, can adopt any level of abstraction to study and analyse the analogue system and formulate his theory about it (the space within the dotted line in Figure 3). So, in order to understand what this involves, we need to pause for a moment and briefly consider what a level of abstraction is.

3.3.1 The Method of Levels of Abstraction

The method of LoA comes from modelling techniques developed in an area of Computer Science, known as *Formal Methods*, in which discrete mathematics is used to specify and analyse the behaviour of information systems. What follows is rather standard knowledge in computer science, and the interested reader is referred to the relevant literature for further details (see for example De Roeve and Engelhardt [1998], Hoare and He [1998], and Floridi and Sanders [2004]). Before introducing a quick summary of the method of levels of abstraction, an everyday example may be useful.

Suppose we wish to describe the state of a traffic light in Rome. We might decide to consider an *observable*, named *colour*, of type $\{red, amber, green\}$ that corresponds to the colour indicated by the light. This option abstracts the length of time for which the particular colour has been displayed, the brightness of the light, the height of the traffic light, and so on. So the choice of type corresponds to a decision about how the phenomenon is to be regarded. To specify such a traffic light for the purpose of construction, a more appropriate type might comprise a numerical measure of wavelength. Furthermore, if we are in Oxford, the type of colour would be a little more complex, since – in addition to red, amber and green – red and amber are displayed simultaneously for part of the cycle. So, an appropriate type would be $\{red, amber, green, red-amber\}$. What we have just seen is a basic concept of *level of abstraction* (LoA), understood as a finite but non-empty set of observables, where an *observable* is just an *interpreted typed variable*, that is, a typed variable together with a statement of what feature of the system under consideration it stands for.

The definition of observables is only the first step in studying a system at a given LoA. The second step consists in deciding what relationships hold between the observables. This, in turn, requires the introduction of the concept of system “behaviour”.

Not all values exhibited by combinations of observables in a LoA may be realised by the system being modelled. For example, if the four traffic lights in Oxford are modelled by four observables, each representing the colour of a light, the lights should not in fact all be green together (assuming they work properly). In other words, the combination in which each observable is green should not be realised in the system

being modelled, although the types chosen allow it. Some technique is therefore required to describe those combinations of observable values that are actually acceptable. The most general method is simply to describe all the allowed combinations of values. Such a description is determined by a predicate, whose allowed combinations of values we call the “system behaviours”. A *behaviour* of a system, at a given LoA, is defined to consist of a predicate whose free variables are observables at that LoA. The substitutions of values for observables that make the predicate true are called the *system behaviours*.

A *moderated LoA* is defined to consist of a LoA together with a behaviour at that LoA. For example, human height does not take arbitrary rational values, for it is always positive and has an upper limit of (say) nine feet. The variable h , representing height, is therefore constrained to reflect reality by defining its behaviour to consist of the predicate $0 < h < 9$, in which case any value of h in that interval is a “system” behaviour.

Since Newton and Leibniz, the behaviours of the analogue observables have typically been described by differential equations. A small change in one observable results in a small, quantified change in the overall system behaviour. Accordingly, it is the rates at which those continuous observables vary which is most conveniently described. The desired behaviour of the system then consists of the solution of the differential equations. However, this is a special case of a predicate: the predicate holds at just those values satisfying the differential equation. If a complex system is approximated by simpler systems, then the differential calculus provides a method for quantifying the approximation. The use of predicates to demarcate system behaviour is essential in any (nontrivial) analysis of discrete systems because in the latter no such continuity holds: the change of an observable by a single value may result in a radical and arbitrary change in system behaviour. Yet, complexity demands some kind of comprehension of the system in terms of simple approximations. When this is possible, the approximating behaviours are described exactly, by a predicate, at a given LoA, and it is the LoAs that vary, becoming more comprehensive and embracing more detailed behaviours, until the final LoA accounts for the desired behaviours. Thus, the formalism

provided by the method of abstraction can be seen as doing for discrete systems what differential calculus has traditionally done for analogue systems.

Specifying the LoA at which one is working means clarifying, from the outset, the range of questions that (a) can be meaningfully asked and (b) are answerable in principle. In standard terminology, the input of a LoA consists of the *system* under analysis, comprising a set of *data*; its output is a *model* of the system, comprising *information*. The quantity of information in a model varies with the LoA: a lower LoA, of greater resolution or finer granularity, produces a model that contains more information than a model produced at a higher, or more abstract, LoA. Thus, a given LoA provides a quantified commitment to the kind and amount of information that can be “extracted” from a system. The choice of a LoA pre-determines the type and quantity of data that can be considered and hence the information that can be contained in the model. So, knowing at which LoA the system is being analysed is indispensable, for it means knowing the scope and limits of the model being developed.

Let us now return to Raphael’s theory.

3.3.2 The Observer’s Endless Levels of Abstraction

As we have seen, Raphael can observe the analogue system, produced by Gabriel on the basis of Michael’s input, at an endless number of LoA. This is tantamount to saying that Raphael may avail himself of as many different interfaces as he wishes in order to analyse the system (metaphorically: in order to read Gabriel’s analogue message), each of which will provide further information about the system itself. Figure 3 summarises this point.

Raphael never halts. This is just a colourful way of saying that, outside our thought experiment, *analogue systems are closed under modelling at levels of abstraction*, or, which is the same thing, there is no finite number of levels of abstractions that can provide all possible models of an analogue system.

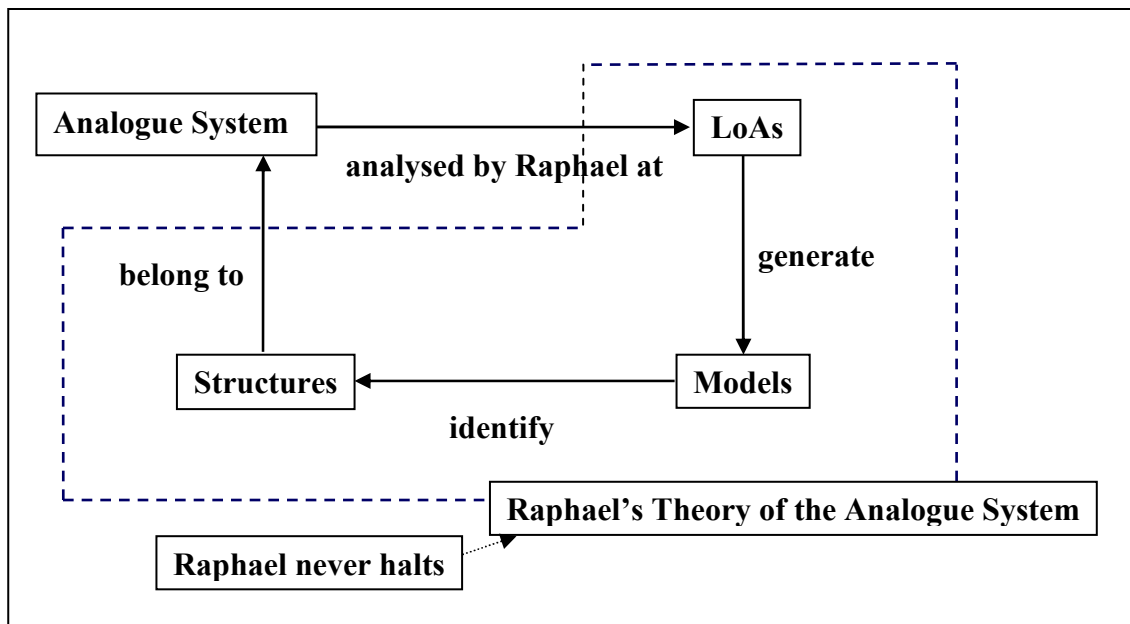


Figure 3 Third stage of the thought experiment, Raphael's LoAs.

The observer (Raphael and us) can extract an endless amount of *semantic* information²⁴ from an analogue system, so it requires an endless amount of computation to describe even the simplest system. Some people find this intellectually unpleasant. Richard Feynman was one of them and this is why we have seen (cf. the quotation from Weinberg in § 2.1) that he is some times listed among the digital ontologists. Here is how he phrased the problem:

It always bothers me that, according to the laws as we understand them today, it takes a computing machine [our Raphael] an infinite number of logical operations to figure out what goes on in no matter how tiny a region of space, and no matter how tiny a region of time [our analogue system]. How can all that be going on in that tiny space? Why should it take an infinite amount of logic to figure out what one tiny piece of space/time is going to do? So I have often made the hypothesis that ultimately physics will not require a mathematical statement, that in the

²⁴ The specification "semantic" is there to prevent one from mistaking this fact for a contradiction of Shannon's fundamental theorem according to which, if we "Let a source have entropy H (bits per symbol) and a channel have a capacity C (bits per second). Then it is possible to encode the output of the source in such a way as to transmit at the average rate of $C/H - \epsilon$ symbols per second over the channel where ϵ is arbitrarily small. It is not possible to transmit at an average rate greater than C/H . (Shannon and Weaver [1949 rep. 1998], 59). Shannon's limiting result states that if you devise a good code you can transmit symbols over a noiseless channel at an average rate as close to C/H as one may wish but, no matter how clever the coding is, that average can never exceed C/H . The limit can be extended to memory devices, which can contain only a certain amount of syntactic information.

end the machinery will be revealed, and the laws will turn out to be simple, like the chequer board with all its apparent complexities. (Feynman [1992], pp. 57-58).²⁵

However, to anyone acquainted with Kant's philosophy, the boundless informational-richness of the world comes as no surprise. The system under observation (Gabriel's message read by Raphael) works a bit like the noumenal. For the method of levels of abstraction allows one to understand that reality in itself though not epistemically inaccessible it remains an epistemically inexhaustible resource out of which knowledge is constructed.

Since Raphael never halts, he can never know whether the world is digital or analogue. On the one hand, if the world were digital, there might be a theoretical possibility of knowing it, but it probably is not, or at least Raphael's and our knowledge of the world, inescapably mediated by some LoA, seem to point in the opposite direction. It is of no avail to object that even in our thought experiment we have assumed the existence of four agents/archangels resembling four digital Turing Machines. As Turing himself remarked, even in the case of actual digital artefacts:

The digital computers [...] may be classified amongst the "discrete state machines", these are the machines which move by sudden jumps or clicks from one quite definite state to another. These states are sufficiently different for the possibility of confusion between them to be ignored. *Strictly speaking there are no such machines. Everything really moves continuously* (my emphasis). But there are many kinds of machine, which can profitably be thought of as being discrete state machines. (Turing [1950], p. 439)

By "profitably be thought of" Turing was really referring to the level of abstraction at which "many kinds of machine" are analysed as "discrete state machines". In other words, even our agents are, deep down, better understood by us as analogue agents. The world, at least as we and Raphael experience it, might well be analogue, and the digital only a convenient abstraction or the result (and technical exploitation) of some physical features in our artefacts.

On the other hand, Raphael, and ourselves with him, cannot exclude either empirically or in principle that reality in itself might actually be digital, i.e. discrete. For

²⁵ Because of similar statements, Feynman is sometimes listed among the digital ontologists but, to be fair to Feynman, the quotation continues "But this speculation is of the same nature as those other people make – 'I like it', 'I don't like it', – and it is not good to be prejudiced about these things".

the endless amount of information that Raphael can extract from Gabriel's message says nothing about the presence or absence of some ultimate *indivisibilia*.

At this point, the doubt that naturally comes to one's mind is whether the analogue/continuous vs. digital/discrete dichotomy may be sound at all, or at least whether its application to the description of the intrinsic nature of reality may not be misguided. Kant thought it was. And I agree that it is. Analogue/continuous and digital/discrete are modes of presentation of Being, i.e. ways in which a system is modelled (experienced or conceptualised) by an observer (an epistemic agent) at a given level of abstraction. The last stage in our thought experiment consists in trying to understand how this could be the case.

3.4 Digital and Analogue are Features of the Level of Abstraction

From the observer's position, which is Raphael's as well as ours, it is impossible to establish whether reality in itself (the noumenal) is analogue or digital. In the *Critique of Pure Reason*, Kant had already convincingly argued, although in very different terms, against the soundness of the dichotomy and its correct application. As is well-known, each of the four antinomies discussed by Kant comprises a thesis and an antithesis, which are supposed to be both reasonable and irreconcilable. The one which interests us here is the second. Paraphrasing Kant, it states that (A 434-5/B 462-3):

(Digital) Thesis: the world is discrete; everything in the world consists of elements that are ultimately simple and hence indivisible.

(Analogue) Antithesis: the world is continuous; nothing in the world is simple, but everything is composite and hence infinitely divisible.

As Kant argues, the conflict is not between empirical experience and logical analysis. Rather, the antinomies, ours included, are generated by an unconstrained demand for unconditioned answers to fundamental problems concerning (1) time and space, (2) complexity/granularity, (3) causality and freedom or (4) modality. Kant is right: striving for the unconditioned is equivalent to the natural, yet profoundly mistaken, attempt to analyse a system (the world in itself, for Kant) independently of any (specification of the) level of abstraction at which the analysis is being conducted, the questions are

being posed and hence the answers are being offered. In other words, trying to overstep the limits set by the adopted LoAs leads to conceptual confusions.

Kant divides the antinomies into two groups. Ours belongs to the first, in which both the thesis and the antithesis are untenable because the search for the unconditioned mistakes time and space and complexity/granularity for features of the system instead of realising that they are properties set by (or constituting) the level of abstraction at which the system is investigated and hence, as such, subject to alternative formatting. Assuming for the sake of simplicity that a LoA is comparable to an interface, it makes no sense to wonder whether the system under observation is finite in time, space and granularity in itself, independently of the LoA at which it is being analysed, since this is a feature of the interface inherited by the model, and, with some flexibility, different interfaces may be adopted depending on needs, requirements and goals, i.e. teleologically. So, although on the basis of a LoA approach, I cannot but agree with Kant: neither the thesis nor the antithesis is tenable. A good way of making sense of this conclusion in our thought experiment is by referring to a fourth and last agent, Uriel, the “sharpest sighted spirit of all in Heaven”, as Milton called him in *Paradise Lost*.

Uriel builds a wheel in which there are four nodes (see Figure 4). Each node contains either a DAC (digital to analogue converter) or an ADC (analogue to digital converter). Since it is possible to convert information from analogue into digital form and back again with as little loss of detail as needed, Uriel’s wheel generates a system – as an output from an analogue or digital ontology – which will be observed by Raphael as being either analogue or digital depending on the latter’s position with respect to the wheel. It is now obvious that it makes no sense to ask whether the system is digital or analogue in itself. We have discharged our initial assumption about reality in itself being digital or analogue.

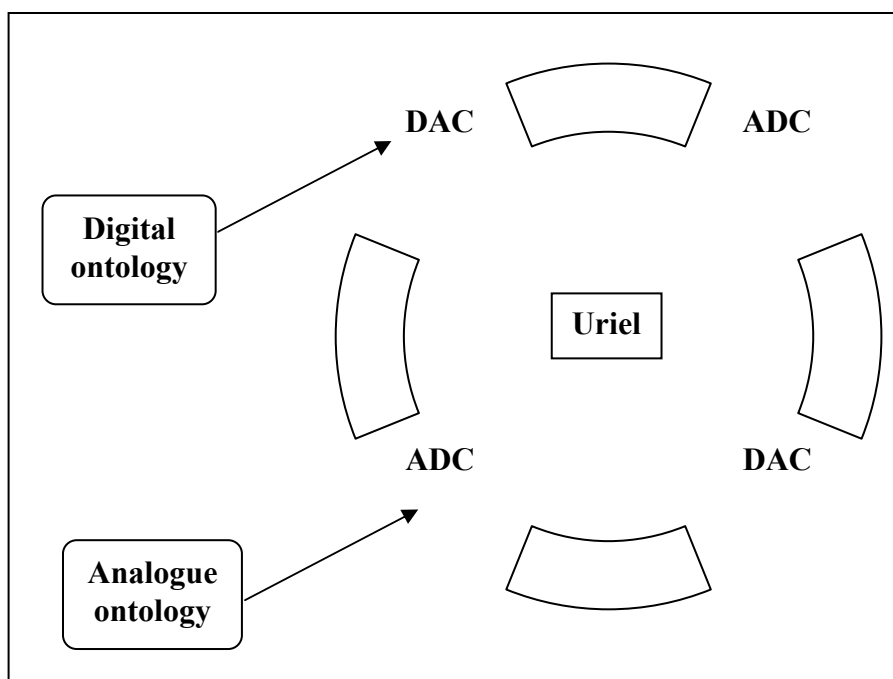


Figure 4 Fourth stage of the thought experiment, Uriel's wheel.

All this is less “philosophical” than it might seem. In quantum mechanics, we are used to seeing a similar effect, known as the wave-particle duality. The basic idea is that all objects (both micro- and macroscopic) enjoy properties associated both with continuous waves and with discrete particles, in short, that objects have a dual nature, partly wave-like and partly particle-like. Although this is experimentally detectable only on very small scales, the factual results are well-known and indisputable, if puzzling in terms of interpretation. The classic double-slit experiment (see Figure 5), for example, shows that, if only one slit is open at a time, the individual photons fired against it that manage to go through the slit and hit the screen on the other side at different times generate a pattern with a single peak, behaving like tiny bullets.²⁶ However, if both slits are open at the same time, the individual photons, still hitting the screen at different times, generate a wave-like interference pattern. In the conceptual frame presented in this paper, the slits act as a hardwired (physically implemented) level of abstraction: change the level

²⁶ To be precise, in real experiments one cannot really distinguish between two plausible explanations for the observed result (a) electrons are classical particles, i.e. they behave just like bullets fired from a gun and (b) electrons are wave trains, i.e. waves of finite length. The real difficulty occurs when both slits are open and the single electrons, sent at separate times, like as many bullets, end up generating a typical interference pattern as if they were waves (no matter of what length).

(open or close the slits) and the beam of individual photons or electrons will be (observed as) behaving “digitally” or “analogically”, like bullets or like waves, presenting a world that is discrete or continuous, not in itself, but in relation to the observer’s LoA. Similar effects can be shown with electrons and other particles. Experiments that test particle properties show that electrons behave as if they were particles, and experiments that test wave properties show that electrons behave as if they were waves. The ways in which one has epistemic access to the system affect the outcome of the analysis or experiment. As Lesne [2007] has well synthesised:

In conclusion, physics in all instances is an interplay between discrete and continuous features, mainly because any such feature actually characterizes a representation, from a given observer, of the real system and its evolution. [...] In practice, the choice between discrete and continuous models should be substantiated with the comparison between the respective scales of description, observation, variations (e.g. gradient scales, oscillation periods, or inhomogeneity sizes) and correlations [what in this paper has been called the method of abstraction, my addition]. [...] Paradoxes and inconsistencies between discrete and continuous viewpoints only appear when forgetting that our descriptions, and even physical laws, are only idealized abstractions, tangent to reality in an appropriate scale range, unavoidably bounded above and below. [...] Any physical theory is in the same way based on a representation of the system and deals only with this representation, while the reality always remains beyond and is never fully captured. Lesne [2007], pp. 35-36.

Kant might have agreed. I certainly do. Raphael observes Uriel’s world as analogue or digital depending on his position (LoA) with respect to it. What remains invariant in Uriel’s world, from Raphael’s perspective, cannot be its digital or its analogue nature, but rather the structural properties that give rise to a digital or analogue reality. These invariant, structural properties are what science is mainly interested in. So it seems reasonable to move from an ontology of things – to which it is difficult no to apply the digital/discrete vs. analogue/continuous alternative – to an ontology of structural relations, to which it is immediately obvious that the previous dichotomy is irrelevant. This is precisely the step taken, in the current debate on scientific realism, by supporters of different forms of structural realism.²⁷ The step is important, and deserves a final comment by way of conclusion.

²⁷ For an overview see Worrall [1989] and Ladyman [1998].




Figure 5 Double-slit experiment: single electrons hitting the screen build up wave-like patterns over time. Reprinted by courtesy of Dr. Akira Tonomura, Hitachi, Ltd., Japan.

4. Three Objections and Replies

The case against digital ontology can now be reinforced by clearing the ground of three potential objections.²⁸

One may contend that the argument begs the question. For it builds an analogue reality through the interface represented by Gabriel – no matter whether the original source (that is, Michael’s world) is intrinsically digital or analogue – but then it purports to show that Raphael (an agent like us), by having access to this analogue reality, cannot determine whether its original source is digital. Unsurprisingly, Raphael is trapped in a (perception of the) world as analogue, but this is merely what is already presupposed.

The objection can be answered by recalling that the goal of the argument is to establish that both a digital and an analogue ontology are untenable. The argument seeks to achieve this conclusion by making a crucial use of the boundless number of levels of

²⁸ In this section I have summarised several questions and objections that I have received during and after the conference, when the paper circulated among several colleagues. I take full responsibility for their specific formulation, although many are as close as possible to the original formats.

abstraction at which an analogue reality may be observed. So, the argument would indeed be begging the question if it were positing not Gabriel's interface, as it does, but rather Uriel's wheel, that is, if the argument were to presuppose that the world is digitally/analogously undetermined to begin with. This the argument does not, since Uriel's wheel is introduced only as a fourth step, in order to make sense of the problem disclosed by the argument. The second step in the argument relies on the analogue interface, represented by Gabriel's "translation", because it needs to make explicit the starting point that is conceded by all parties involved in the debate, namely that some aspects of the world really seem to be intrinsically and irreducibly analogue rather than digital, and that it is up to the defender of digital ontology to make a convincing case for the opposite view. This is why the section's title is a paraphrase of Margolus' comment about "the stubborn legacy of the continuum" (the reader may recall that Margolus is one of the strong supporters of digital ontology). One may retort that the argument would not work were one to assume that Gabriel translates Michael's world into a digital world. The reply is that this is correct but also not relevant here. For the question being debated is not: given that some aspects of the world appear to be digital, is a digital ontology justified? This is uncontroversial but also uninteresting. The question is rather: given that some aspects of the world appear to be analogue, is a digital ontology justified? One only needs to recall why the Pythagoreans abandoned their digital ontology to understand the pressure exercised by the question. The argument simplifies all this by generating a scenario in which Raphael (us) interacts epistemically with an analogue reality, at various levels of abstraction, and this scenario is not in need of a justification, for it is the starting difficulty under discussion.

The defender of a digital ontology might still be unconvinced because of a second object. The argument is clearly Kantian, but all it does, at best, is to explain how (an aspect of) Kant's metaphysics could be true. It does not show that Kant's metaphysics is true. So, even conceding that digital ontology may not be a promising line of research from a Kantian perspective, this fails to be convincing to anyone who does not already share such a transcendental approach.

This objection may be answered by starting from the distinction between arguing for the same conclusion and arguing in the same way for that conclusion. The

conclusion of the argument is indeed Kantian, so it is a matter of scholarship to give to Kant what is Kant's. This may have the further advantage of helping anyone familiar with Kant's philosophy to follow the argument more easily. But the argument itself is not based on, nor intends to be an exposition or a vindication of, Kant's transcendental epistemology. Rather, it seeks to provide a different and independent route to the same theoretical point, so it should be assessed on its own merits (or lack thereof). To put it visually, Kant's argument and the one presented in this paper are like two lines that intersect each other at one point, but coming from, and going in, different directions. So, the objection that the argument is unconvincing (that the new route is blocked) because it reaches a conclusion also reached by Kant (the intersecting point) is either counteractive or ineffective. Counteractive since calling its conclusion Kantian cannot weaken the argument, but it might increase its acceptability. For a reader with Kantian sympathies may find this convergence reassuring and enlightening, while anyone unwilling to follow Kant's route may still find the one presented here appealing. Ineffective because the objection, to be effective, needs to show that (i) the argument is just a version of Kant's argument and that (ii) the latter is wrong, but, even if one were to concede (ii), the fact remains that the argument supports a more constructive position that is not in itself Kantian at all, as we shall see in the answer to the next objection and in the conclusion, so (i) cannot be granted.

According to the third and last objection to be discussed here, the argument against digital ontology appears to depend on the claim that each LoA has equal right to tell us whether the ultimate nature of reality is analogue or digital (each LoA is equally valid). However, this is a claim that an opponent is likely to deny. There seems to be no reason why each LoA should be granted equal authority in answering questions about the ultimate nature of the universe. Some LoA are undoubtedly useful in some contexts and not in others, but we have no reason to think that all of them are equally good at liming the true structure of reality. In particular, someone might argue that it is reasonable that, if our best fundamental physical theory is, say, digital, then this gives us good reason to think that the fundamental nature of reality is digital. This is not deductive warrant, but it does appear to provide some degree of justification. Therefore, for such a critic, when it comes to the question of whether reality is analogue or digital,

it does not seem that each LoA should have equal say. The argument needs to do more to respond to such a critic. For it seems that it is only against the claim that one can adopt a “LoA-free” position and decide from there whether reality is digital or analogue, yet this seems a strawman. A realist about the analogue/digital divide is more likely to claim that some LoAs are better than others for telling one about the ultimate nature of reality, not to claim a miraculous ability to adopt a “LoA-free” position. Furthermore, an attack solely on the “LoA-free” position would be insufficient to support the positive claim that no LoA is a better guide than any other to the ultimate digital/analogue nature of reality.

The objection contains an important clarification, which will help to introduce the positive defence of informational structural realism, sketched in the next section. The clarification is this. The argument is not merely based on the dichotomy “LoA-free” vs. “LoA-bounded” approaches to ontology. Although important, this is only the first move. Once we accept that epistemology is LoA-based and that no ontology can be LoA-free, then a second, but equally crucial move consists in realising that digital and analogue are features of the LoAs adopted to analyse reality, not features of reality in itself, or, to put it differently, that digital and analogue features are internal features of the models made possible by various LoAs, not external features of the systems modelled by LoAs. So the argument is not just that some LoAs show reality to be digital while others show it to be analogue and that we cannot decide which LoAs are better, but, far more importantly, that some LoAs are digital and some are analogue and that, depending on which of them we adopt (because of requirements, goals etc., i.e., teleologically), reality will be modelled as digital or analogue. The case of the double-slit experiment was recalled as a clear illustration of the impossibility of determining the intrinsically digital vs. analogue nature of reality independently of how it is accessed. Now, the objection correctly points out that, even if all this is granted, LoAs are still not “born equal”, and not just as a matter of instrumental convenience. One may correctly argue, as I have done elsewhere (Floridi and Sanders [2004]), that some LoAs still fit better their systems, both in terms of taking full advantages of their affordances and in terms of respecting their constrains, and that this is very significant and worth accounting for, ontologically. So, the objection continues, perhaps one could work his

way inside out, as if were, and try to grasp what the nature of reality might be in itself, given the sort of successful LoAs that are adopted on this side of the relation. To use a terminology borrowed from computer science, one could try to reverse-engineer the output (our models of reality as obtained through our most successful LoAs) in order to obtain at least some information about the original input (Michael's world, or Kant's noumenal world). This is all correct. But we need to be careful about what can be actually inferred from this valuable suggestion. Even assuming, and this is very far from a trivial concession, that reality might be successfully modelled through a digital (or analogue, where the "or" is assumed to be exclusive of the sake of simplicity) ontology and that it is not the case that an analogue (or, if talking of an analogue ontology in the first place, a digital) ontology may not be equally satisfactory, all this would not show that reality is digital (or analogue) in itself. It would only show that our ontological commitment in favour of a digital (or analogue) ontology would be safe in the sense that we could not be proven wrong because we would be (or, our opponent would argue, we would have no way of proving that we are not) taking a feature of our LoAs for a feature of the system that they model. This is unsatisfactory. It is the impasse that I have tried to describe more intuitively by means of Uriel's wheel. The exit from this impasse, in terms of a defence of a form of realist ontology that is compatible with a non-relativistic and yet LoA-based epistemology, leads to a non-Kantian position, which seeks to reconcile digital and analogue ontology by identifying the minimal denominator shared by both. And this is the sort of position that I have defended in the constructive paper, under the label structural information realism. So, insofar as the objection is correct, it seems that an informational approach to ontology is the best way of taking advantage of its lesson.

4. Conclusion: Towards Informational Structural Realism

In this paper, I have argued against the tenability of a digital ontology, and against the soundness of the digital vs. analogue dichotomy when applied to our metaphysical understanding of reality. The criticism is motivated by the fear that informational ontology (according to which the ultimate nature of reality is informational, see Floridi [2004] and Floridi [forthcoming]) might be confused with either pancomputationalism (according to which the universe is a computational system equivalent to a Turing Machine), with a version of digital ontology (according to which the ultimate nature of reality is digital) or with a combination of the two (according to which the universe is a cellular automaton, for example). The fear is justified, for the risk is real. Take for example an influential article by Margolus [2003], in which we read that

Given more than a century of finite-state underpinnings, one might have expected that by now all of physics would be based on informational and computational concepts. That this isn't so may simply reflect the stubborn legacy of the continuum, and the recency [sic] and macroscopic character of computer science. (p. 309).

The passage conflates the informational and the digital, as if they were synonymous. Clearly, they are not, not least because information may easily be analogous or continuous. Yet even Wheeler himself fell into the same trap (see the quotation provided in § 2). Drawing no distinction between an informational and a digital ontology is a mistake. Digital ontology is very implausible and obviously behaves as a misleading distraction, when it comes to assessing the value of an informational ontology. The latter can be better appreciated once the ground has been cleared of any potential confusion. Thus, this paper has provided the preparatory *pars destruens* (digital ontology is a bad idea) of a two-stage piece of research, whose second, *pars construens* (in favour of informational structural realism) has been developed in Floridi [forthcoming]. Since the *pars construens* is really a separate project, what follows is only a brief synthesis of Floridi [forthcoming].

Informational structural realism (ISR) is a version of structural realism. As a form of realism, ISR is committed to the existence of a mind-independent reality addressed by, and constraining, our knowledge. It supports the adoption of LoAs that carry a minimal ontological commitment in favour of the structural properties of reality

and a reflective, equally minimal, ontological commitment in favour of structural objects. Unlike other versions of structural realism, ISR supports an informational interpretation of these structural objects. This second commitment, in favour of structural relations, is justified by epistemic reasons. We are allowed to commit ourselves ontologically to whatever minimal conception of objects is useful to make sense of our first commitment in favour of structures. The first commitment answers the question “what can we know?”; the second commitment answers the question “what can we justifiably assume to be there?”. We are now ready for a definition:

(ISR) Explanatorily, instrumentally and predictively successful models (especially, but not only, those propounded by scientific theories) of reality at a given LoA can be, in the best circumstances, increasingly informative about the relations that obtain between the (possibly sub-observable) informational objects that constitute the system under investigation (through the observable phenomena).

A significant consequence of ISR is that, as far as we can tell, the ultimate nature of reality is informational, that is, it makes sense to adopt LoAs that commit our theories to a view of reality as mind-independent and constituted by structural objects that are neither substantial nor material (they might well be, but we have no need to suppose them to be so) but cohering clusters of *data* (not in the alphanumeric sense of the word, but in an equally common sense of *differences de re*, i.e. mind-independent, concrete, relational points of lack of uniformity). Structural objects work epistemologically like *constraining affordances*: they allow or invite certain constructs (they are *affordances* for the information systems, like Raphael or us, who elaborate them) and resist or impede some others (they are *constraints* for the same systems), depending on the interaction with, and the nature of, the information system that processes them. They are exploitable by a theory, at a given LoA, as input of adequate queries to produce information (the model) as output. This epistemic malleability of reality as a resource seems to be what Chakravartty [2003] defines as the “dispositional nature” of structural objects²⁹ and Saunders [2003] calls their “heuristic plasticity”.

When Cassirer talked about structuralism, he had in mind, like Kant and Russell before him and Maxwell after him, a full-blooded ontology of objects as structural

²⁹ Apparently David Lewis held a similar view, see Langton [2004].

entities (Gower [2000], French [2001]). Having shown that digital ontology is not a promising line of research, we should not throw away every metaphysics vaguely resembling it, since an informational ontology and an informational structural realism remain valuable options.³⁰

³⁰ Previous versions of this paper were presented at several meetings. I am grateful to the participants for their feedback and in particular to Rainer Hammwoehner, for this invitation to give a series of seminars at the University of Regensburg, one of which was based on this paper (9 November, 2005); to Beth Lord and Timothy Chappell, for their invitation to give a talk as part of the *Research Seminar* series of the Department of Philosophy, University of Dundee (15 March, 2006); to Marcello D'Agostino, for his invitation to give a seminar at the Dipartimento di Scienze Umane, Università degli Studi di Ferrara (24 Marzo, 2006); to Amnon Eden and Raymond Turner, for their invitation to give a talk as part of the *Philosophy of Computer Science Seminar* series of the Department of Computer Science, University of Essex (21 April, 2006); to Charles Ess, Gordana Dodig-Crnkovic, Johnny Søraker, Susan Stuart and May Thorseth for their invitation to present a paper at the *Fourth International European Conference on Computing and Philosophy* (ECAP) in Trondheim, Norway (22 - 24 June, 2006); to A. Arrieta, G. Arrizabalaga, G. Boniolo, J. I. Galparsoro and G. Pin for their invitation to give a keynote speech at the *Seventh International Ontology Congress*, in San Sebastian, Spain (5 October, 2006); to Don Fallis, Michael Gill, Uriah Kriegel and Kay Mathiesen, and the Center for Consciousness Studies, the Research Group on the History and Philosophy of Information Access, the School of Information Resources and Library Science, and the International Visitors Fund for their invitation to give a series of seminars at the University of Arizona, one of which, delivered as part of the *Colloquia Series* of the Department of Philosophy, was based on this paper (9 February 2007). I am also grateful to Vincent C. Müller for having provided a preprint of his paper, to Hitachi, Ltd., Japan Akira Tonomura for their permission to reprint Figure 5; to Patrick All, Karen Mather, Sebastian Sequoiah-Grayson, Johnny Søraker and Matteo Turilli, who read the last draft and provided many useful comments; and finally to the two anonymous referees for their comments, their suggestions on how to improve the paper and above all for having spotted a very silly mistake apparently so flagrant that nobody had managed to identify it before. The usual disclaimer applies: I am the only person responsible for any remaining mistakes.

References

- 't Hooft, G. 1997, *In Search of the Ultimate Building Blocks* (Cambridge: Cambridge University Press).
- 't Hooft, G. 2002, "How Does God Play Dice? (Pre-)Determinism at the Planck Scale" in *Quantum [Un]Speakables, from Bell to Quantum Information*, edited by R.A. Bertlmann and A. Zeilinger (Berlin: Springer Verlag), 307-316.
- 't Hooft, G. 2003, "Can Quantum Mechanics Be Reconciled with Cellular Automata?" *International Journal of Theoretical Physics*, 42, 349-354.
- 't Hooft, G. 2005, "Does God Play Dice?" *Physics World*, 18(12), 21-23.
- Ball, P. 2002, June 3, "Universe Is a Computer", *Nature News*. doi:10.1038/news020527-16
- Chaitin, G. J. 2005, *Meta Math! : The Quest for Omega* (New York: Pantheon Books).
- Chakravartty, A. 2003, "The Structuralist Conception of Objects", *Philosophy of Science*, 70(5), 867 - 878.
- Chalmers, D. J. 1996, *The Conscious Mind: In Search of a Fundamental Theory* (New York: Oxford University Press).
- de Roeper, W.-P., and Engelhardt, K. 1998, *Data Refinement : Model-Oriented Proof Methods and Their Comparison* (Cambridge: Cambridge University Press).
- Dedekind, R. 1963, *Continuity and Irrational Numbers, in Essays on the Theory of Numbers* (New York: Dover). Originally published in 1872.
- Feynman, R. P. 1992, *The Character of Physical Law* (London: Penguin). Originally published: London: British Broadcasting Corporation, 1965.
- Floridi, L. 2004, "Informational Realism" in *Computers and Philosophy 2003 - Selected Papers from the Computer and Philosophy Conference (Cap 2003) Acs - Conferences in Research and Practice in Information Technology* edited by John Weckert and Yeslam Al-Saggaf (pp. 7-12).
- Floridi, L. forthcoming, "A Defence of Informational Structural Realism", *Synthese*.
- Floridi, L., and Sanders, J. W. 2004, "The Method of Abstraction" in *Yearbook of the Artificial - Nature, Culture and Technology, Models in Contemporary Sciences*, edited by Massimo Negrotti (Bern: Peter Lang), 177-220.
- Fredkin, E. 1992, "Finite Nature", *Proceedings of the XXVIIth Rencontre de Moriond, March 22-28, 1992*, Les Arcs, Savoie, France (Editions Frontieres: Gif-sur-Yvette, France).
- Fredkin, E. 2003a, "The Digital Perspective", *International Journal of Theoretical Physics*, 42(2), 145.
- Fredkin, E. 2003b, "An Introduction to Digital Philosophy", *International Journal of Theoretical Physics*, 42(2), 189-247. Also available online at <http://digitalphilosophy.org/>
- Fredkin, E. online, *Digital Mechanics - an Informational Process Based on Reversible Universal Cellular Automata*, http://www.digitalphilosophy.org/dm_paper.htm.
- French, S. 2001, "Symmetry, Structure and the Constitution of Objects", *Symmetries in Physics, New Reflections: Oxford Workshop, January 2001*, Oxford
- Goodman, N. 1968, *Languages of Art: An Approach to a Theory of Symbols* (Indianapolis: Bobbs-Merrill).

- Gower, B. 2000, "Cassirer, Schlick and 'Structural' Realism: The Philosophy of the Exact Sciences in the Background to Early Logical Empiricism", *British Journal for the History of Philosophy*, 8(1), 71–106.
- Hoare, C. A. R., and He, J. 1998, *Unifying Theories of Programming* (London: Prentice Hall).
- Holden, T. A. 2004, *The Architecture of Matter : Galileo to Kant* (Oxford: Clarendon Press).
- Hyötyniemi, H. 1996, "Turing Machines Are Recurrent Neural Networks", *STeP'96 - Genes, Nets and Symbols*, edited by J. Alander, T. Honkela, and M. Jakobsson (Publications of the Finnish Artificial Intelligence Society (FAIS), Helsinki, Finland), 13-24.
- Kant, I. 1998, *Critique of Pure Reason* repr. w. corr. (Cambridge: Cambridge University Press).
- Ladyman, J. 1998, "What Is Structural Realism?" *Studies in History and Philosophy of Science*, 29A(3), 409-424.
- Langton, R. 2004, "Elusive Knowledge of Things in Themselves", *Australasian Journal of Philosophy*, 82(1), 129-136.
- Lesne, A. 2007, "The Discrete Versus Continuous Controversy in Physics", *Mathematical Structures in Computer Science*, 17(02), 185-223.
- Lewis, D. 1971, "Analog and Digital", *Nous*, 5(3), 321-327.
- Lloyd, S. 2002, "Computational Capacity of the Universe", *Physical Review Letters*, 88(23), 237901-237904.
- Lloyd, S. 2006, *Programming the Universe : From the Big Bang to Quantum Computers* (London: Jonathan Cape).
- Margolus, N. 2003, "Looking at Nature as a Computer", *International Journal of Theoretical Physics*, 42(2), 309-327.
- Müller, V. C. forthcoming, "What Is a Digital State?"
- Pais, A. 2005, *Subtle Is the Lord: The Science and the Life of Albert Einstein* (Oxford; New York: Oxford University Press). Originally published in 1982, republished with a new foreword by Sir Roger Penrose.
- Petrov, P. 2003, "Church-Turing Thesis Is Almost Equivalent to Zuse-Fredkin Thesis (an Argument in Support of Zuse-Fredkin Thesis)", *Proceedings of the 3rd WSEAS International Conference on Systems Theory and Scientific Computation, Special Session on Cellular Automata and Applications (ISTASC'03)*, Rhodes Island, Greece.
- Saunders, S. 2003, "Structural Realism, Again", *Synthese*, 136(1), 127-133.
- Sayre, K. M. 1976, *Cybernetics and the Philosophy of Mind* (London: Routledge and Kegan Paul).
- Schmidhuber, J. 1997, "A Computer Scientist's View of Life, the Universe, and Everything ", *Lecture Notes in Computer Science*, 1337, 201-208.
- Schmidhuber, J. forthcoming, " All Computable Universes", *Spektrum der Wissenschaft (German edition of Scientific American)*.
- Shannon, C. E., and Weaver, W. 1949 rep. 1998, *The Mathematical Theory of Communication* (Urbana: University of Illinois Press). Foreword by Richard E. Blahut and Bruce Hajek.

- Siegelmann, H. T. 1998, *Neural Networks and Analog Computation : Beyond the Turing Limit* (Boston: Birkhàuser).
- Steinhart, E. 1998, "Digital Metaphysics" in *The Digital Phoenix: How Computers Are Changing Philosophy*, edited by T. Bynum and J. Moor (New York: Blackwell), 117-134.
- Steinhart, E. 2003, "The Physics of Information" in *Blackwell Guide to the Philosophy of Computing and Information*, edited by L. Floridi (Oxford: Blackwell), chapter 13.
- Toffoli, T. 2003, "A Digital Perspective and the Quest for Substrate-Universal Behaviors", *International Journal of Theoretical Physics*, 42, 147-151.
- Turing, A. M. 1950, "Computing Machinery and Intelligence", *Minds and Machines*, 59, 433-460.
- Von Neumann, J. 1966, *Theory of Self-Reproducing Automata* (Urbana; London: University of Illinois Press). Edited and completed by A. W. Burks.
- Weinberg, S. 24 October 2002, "Is the Universe a Computer?" *The New York Review of Books*.
- Wheeler, J. A. 1990, "Information, Physics, Quantum: The Search for Links" in *Complexity, Entropy, and the Physics of Information*, edited by W. H. Zureck (Redwood City, Cal.: Addison Wesley),
- Wolfram, S. 2002, *A New Kind of Science* (Champaign, Ill.: Wolfram Media).
- Worrall, J. 1989, "Structural Realism: The Best of Both Worlds?" *Dialectica*, 43, 99-124.
- Zuse, K. 1967, "Rechnender Raum", *Elektronische Datenverarbeitung*, 8, 336-344.
- Zuse, K. 1969, *Rechnender Raum* (Braunschweig: Vieweg). Eng. tr. with the title *Calculating Space*, MIT Technical Translation AZT-70-164-GEMIT, Massachusetts Institute of Technology (Project MAC), Cambridge, Mass. 1970.
- Zuse, K. 1993, *The Computer, My Life* (Berlin ; New York: Springer-Verlag).