

Varieties of Analog and Digital Representation

Whit Schonbein

Received: 26 July 2013 / Accepted: 8 April 2014 / Published online: 24 April 2014
© Springer Science+Business Media Dordrecht 2014

Abstract The ‘received view’ of the analog–digital distinction holds that analog representations are continuous while digital representations are discrete. In this paper I first provide support for the received view by showing how it (1) emerges from the theory of computation, and (2) explains engineering practices. Second, I critically assess several recently offered alternatives, arguing that to the degree they are justified they demonstrate not that the received view is incorrect, but rather that distinct senses of the terms have become entrenched specific fields, perhaps most notably in the cognitive psychology of mental imagery.

Keywords Mental representation · Analog representation · Digital representation · Computation

The received view of the distinction between analog and digital mental representation is that the former are continuous and the latter are discrete (Maley 2011). Indeed, this ‘standard interpretation’ is central to computer science, electrical engineering, and computational modeling. However, several recent analyses propose alternatives which carve the distinction orthogonally to the standard interpretation, treating some traditionally digital systems as analog, and some traditionally analog systems as digital (Katz 2008; Maley 2011). In this paper I examine how these alternatives contrast with the standard interpretation, whether they are justified, and if so, how they should be accommodated.

The structure and contributions of the paper are as follows. First, in “[Making sense of implementation](#)” and “[Applying the standard interpretation](#)” sections, I

W. Schonbein (✉)

Department of Computer Science, University of New Mexico, Mail Stop: MSC01 1130, 1
University of New Mexico, Albuquerque, NM 87131-0001, USA
e-mail: whit.schonbein@gmail.com

defend the received view by emphasizing its origins and utility in computer science and engineering, concluding that it is too important to be redefined or otherwise discarded. Second, in “[Applying the standard interpretation](#)” section, I consider the proffered alternatives in detail. Each has its roots in how the distinction manifests itself in particular contexts: The engineering of binary digital devices (Maley’s ‘digital’), literature on mental rotation (Maley’s ‘analog’), and literature on numerical cognition (Katz’s ‘analog’ and ‘digital’). Regarding Maley’s proposals: (1) I show how each contrasts with the standard interpretation, (2) illustrate how his ‘digital’ is consistent with the more general framework established by the received view, (3) question whether his ‘analog’ correctly characterizes how the term is used in the mental rotation literature, and (4) conclude that, if Maley’s alternatives are indeed entrenched in the literature he identifies, then it is preferable to explicitly flag them as distinct senses. Regarding Katz’s proposed analysis, I argue that it (1) is inconsistent with the standard interpretation, and (2) cannot be justified by appeal to numerical cognition literature from which it appears to be drawn; consequently, it should not be retained as a distinct sense.

The Received View and the Theory of Computation

Slide rules are prototypical (mechanical) analog computers, while laptop computers are prototypical (electronic) digital computers. According to the ‘received view’ or ‘standard interpretation’, the difference between the representations used by these two computational devices is that those of the former are *continuous* while those of the latter are *discrete* (Maley 2011, p. 119). The terms ‘continuous’ and ‘discrete’ are here understood mathematically: Discrete representational schemes are bijective with (a finite subset of) the natural numbers, and continuous representational schemes are bijective with the reals. Less formally, the received view holds that slide rules are analog devices precisely because they utilize representations that vary by ‘arbitrarily small degrees’, while standard electronic computers are digital because they use representations that do not. In this section I argue that the received view is one that should be endorsed because it is central to the theory of computation, i.e., the abstract, theoretical foundations of computation in general.

These foundations can be explicated in a variety of equivalent ways, but the most common is through computational automata, e.g., Turing machines. Turing machines are conceptualized as having three components: A controller, a memory store, and a means for reading from and writing to that store. A specific Turing machine is defined by specifying (among other things) the finite set of controller states for that machine, and the rules for transitioning between states given inputs and the current contents of memory.

These inputs and memory contents are the representations used by the system, and they are *thoroughly discrete*: The inputs and the contents of memory take the form of sequences of symbols taken from a finite set of primitive representations Γ (Hopcroft and Ullman 1979). Finite sets are, by definition, bijective with a subset of the natural numbers, which is to say that representations in traditional computational

automata are discrete. The same observation holds for other, equivalent formulations of computation, such as the lambda calculus or register machines.¹

There are good reasons *why* Γ is discrete. One is that the problem Turing is addressing requires it. Turing's goal is to give a formal account of the proper subset of reals that can be exhaustively represented using a base- n representational system, through the operation of a finite algorithm, up to an arbitrary digit i (Turing 1936, p. 230). If Γ contained symbols that differed by arbitrarily small degrees, Γ would replicate the structure of the reals, defeating the purpose of the Turing machine. So the discreteness of Turing machine representations is a precondition on addressing the problem the machines are supposed to address.

A second reason for making Γ discrete is that Turing views continuous representational schemes as a *problem* for describing what humans do when they calculate, and hence for what Turing machines do. For instance, when discussing why the set of controller states must be finite, he writes,

The reasons for [positing a finite number of controller states] are of the same character as those which restrict the number of symbols. If we admitted an infinity of states of mind, some of them will be “arbitrarily close” and will be confused. (Turing 1936, p. 250)

The set of representations must be discrete, according to Turing, because otherwise they could not be reliably distinguished.²

These sorts of commitments to discreteness could be multiplied,³ but for present purposes I hope this brief discussion is sufficient to dispel any doubt that representations grounded in Turing machines (and their equivalents) *are* discrete in the sense given above, and that being discrete is an essential property of such computational systems. The issue, then, is how the discreteness of Turing machine representations is related to representation in *digital* devices.

One reason to think that digital devices use discrete representational schemes is that ‘digital’ can be viewed as a colloquialism for ‘discrete’: Digital representations are representations that can be *counted* using *digits*, i.e., fingers. When we count in this way, we enumerate in whole numbers, and do so with a finite set, so it makes

¹ For instance, the expressions of lambda calculus are explicitly defined as consisting of a finite number of symbols plus an “enumerably infinite” number of variable names, of which only a finite number are used in any given formulation of a lambda-definable function (Church 1936). Similarly, Post's (1936) definition of an effectively calculable system uses one symbol, a finite set of operations, and any given problem uses a finite number of contiguous squares, each of which may contain a single instance of the symbol.

² I am not endorsing Turing's argument; I'm only noting that the fact he makes it supports the view that Turing machines (and hence the foundations of computational theory) are committed to discrete representations.

³ A particularly interesting instance of the commitment to discrete representation occurs in Turing's discussion of how symbols are instantiated in the memory of a Turing machine, a paper tape. The tape is divided into squares, each of which can contain a single symbol token. Each square is continuous along each dimension, but is closed and bounded, and hence compact (Turing 1936, p. 249, footnote). Consequently, the space in which symbols are instantiated is such that that the set of possible symbols ‘reduces’ to a finite number of types in the sense that any arbitrary subset of points will be closer to one of these finite subsets than another. Despite their being continuous, Turing's mathematical medium for symbols—points in a plane – respects the requirement that Γ be (finite and) discrete.

intuitive sense that when we speak of a “digital” representational system we mean one that is discrete.

A better reason for associating ‘discrete’ with ‘digital’ comes from the fact that digital computers are engineered in accordance with Turing’s account of computation, i.e., digital devices aspire to realize discrete representations. For example, contemporary digital computers are based on a representational scheme with two representational types, and researchers have also explored implementing devices using unary (base-1), ternary (base-3), quaternary (base-4), and other countable schemes. All of these approaches are grounded in a process of rendering discrete some underlying physical parameter, i.e., quantizing a signal. Take for instance processing a sound wave, which is continuous along both time and magnitude. It is possible to quantize one or both of these dimensions. When we quantize the temporal dimension, we sample the magnitude of the waveform at countable points in time, and when the information contained at each sample is also quantized, the result is a *digital* waveform, i.e., one that has been rendered discrete on both time and magnitude dimensions.⁴ Similarly, in everyday binary digital computers, the CPU speed is the temporal sample rate, and the voltage magnitudes are quantized into two values at each of those samples; ternary computers quantize into three values, quaternary computers quantize into four values, and so forth. Building digital computers thus involves a process of quantization that renders a (continuous) signal discrete in the sense identified by the received view.

This is not to say that digital devices have always been built on the basis of some explicitly articulated prior mathematical theory of discrete structure. For many years the design and implementation of digital devices proceeded in relative independence from computational theory, and it was not until the 1950’s that the issue was addressed in any detail, through the introduction of ‘register machines’. These automata were explicitly motivated by the goal of reconciling discrete computational theory with digital engineering. For instance, Wang (1957) begins his seminal paper by appealing to this goal:

Turing’s theory of computable functions antedated but has not much influenced the extensive actual construction of digital computers. These two aspects of theory and practice have been developed almost entirely independently of each other. ... It cannot, however, fail to strike one as rather strange that often the same concepts are expressed by very different terms in the two developments. One is even inclined to ask whether a rapprochement might not produce some good effect. (p. 63)

The received view did not spring fully-formed from the brows of Turing, Post, Kleene and Church, but rather is the result of historical contingencies involving the refinement of concepts drawn from different fields of research. The end product *is*

⁴ As Sarpeshkar (1998) notes, in signal processing and electrical engineering the terms ‘analog’ and ‘digital’ are associated with the discreteness or continuity of the signal rather than the time dimension. Consequently, there are in addition to ‘full-blooded’ analog and digital systems, asynchronous digital systems (continuous time, discrete signal) and clocked analog systems (discrete time, continuous signal). For purposes of this paper I bracket these variations, focusing only on the ‘full-blooded’ versions of each type. However, the fact the standard interpretation facilitates these distinctions should not be overlooked.

the standard interpretation, i.e., Wang's desired "rapprochement" between theory and practice, i.e., between discrete computation and digital engineering.⁵

The relevant contrast class to discrete representation is, according to the received view, *continuous* representation. Given that the theoretical core of effective calculation is a branch of discrete mathematics, this contrast is a reasonable one to make. For instance, as already noted, if we allow members of Γ to vary by arbitrarily small degrees then Turing machines no longer address the problem they were designed to address. Second, the discrete/continuous distinction is significant in that allowing for a continuous symbol space can result in automata computationally more powerful than Turing machines (Siegelmann and Sontag 1994; Blum et al. 1989; Koiran 1994; Bournez and Campagnolo 2008). Both of these observations show that the continuous/discrete boundary plays an important role in demarcating different types of computational system, and hence forms a reasonable boundary between types of representational scheme.

Furthermore, the discrete/continuous distinction is consistent with familiar philosophical debates in the philosophy of psychology. In particular, influential critiques of computational approaches to cognition (e.g., van Gelder 1995, Horgan and Tienson 1996) are reasonably interpreted as asserting that the problem lies with their reliance on a finite and discrete symbol space (Schonbein 2005). For instance, these critiques claim that cognition is too 'subtle' to be captured by traditional computational automata, and so real-valued equations must be used instead. This debate over mental representation thus tacitly acknowledges that discrete representations are to be contrasted with continuous ones.⁶ Again, this suggests the received view is correct in contrasting digital/discrete representations with continuous ones. The issue, then, is whether continuous representational schemes should be associated with analog ones.

One reason to posit a close relationship between analog and continuous representations is historical: Analog computers were the default computational mechanisms for thousands of years prior to the introduction of the digital computer, and they operate over continuous values. So, drawing on some 20th century examples, slide rules perform multiplication by adding two logarithms, and traditional artillery guidance systems compute integrals (Clymer 1993). A second reason is that the term 'analog' is used in engineering and computer science to indicate the presence of continuous values. For instance, in Bournez and Campagnolo (2008) survey of automata that make use of continuous values, analog systems are categorized as continuous, not as discrete (p. 395). Similarly, in signal processing, analog signals are those that are continuous, i.e., not quantized. In short, in the context of engineering, computer science, and mathematics, 'analog' is indicative of continuity.

To sum, the received view of the distinction between analog and digital representation holds that analog devices are so-called because they use continuous

⁵ These historical contingencies help explain why 'discrete' and 'continuous' tend to be used in mathematical contexts, while 'digital' and 'analog' are often used when talking about physical devices. So, for instance, there are textbooks on 'discrete' math, but not on 'digital' math; similarly, there are textbooks on 'analog' computers, but not 'continuous' computers. Discrete and continuous formats are the mathematical blueprints for what gets implemented by digital and analog devices, respectively.

⁶ For the purposes of this paper I ignore the fact that some proponents of continuity hold that physical states described using real-valued equations are not representational.

representational schemes, while digital devices use discrete schemes. In this section I've argued that given our theoretical understanding of computation, the received view is a reasonable position to take. First, contemporary computational theory contrasts systems with discrete representations and those with continuous ones. Second, discrete systems form the theoretical foundations for digital computational devices. Finally, analog computation is identified both historically and in contemporary literature with continuous representational schemes.

Making Sense of Implementation

In this section I explore the relationship between formal theory and practice in more detail. One of the primary reasons to accept the received view, I argue, is that it *explains characteristic engineering practices*.

To facilitate discussion, let us introduce a tripartite distinction between (1) the structure of the underlying physical medium used to implement a representational scheme, (2) the structure of the representational scheme itself, and (3) the structure of the domain being represented. Maley and Katz independently articulate the initial aspects of this distinction; Maley's version is:

Let a *representational medium* be the physical substrate in which a representation is instantiated. Let a *representational format* be the structure of the system of representation, regardless of the medium. (Maley 2011, p. 118, original emphasis)⁷

He illustrates the concept with a simple example: The use of individual grains of sand to represent some quantity (such as the number of people in a room), where the property represented is simply equal to the number of grains of sand. In this system, the medium is the sand, and the representational format is "equal to the natural numbers" (Maley 2011, p. 118). I interpret Maley to mean that the system of representation used to indicate the number of people has the structure of the natural numbers, i.e., it has a discrete *representational format*. Furthermore, we can also note that, since any given sample of sand will contain a finite number of distinct grains, the structure of the medium is itself discrete, i.e., it has a discrete *medium format*. Finally, we can further extend the example by adding in the third member of the distinction: The domain being represented (people in a room) also has the structure of the natural numbers. So, in the present example, there is a discrete medium, a discrete representational format, and a discrete representational domain. In contrast, the format of voltages serving as the medium in digital electronic computers is continuous, being properly described by Maxwell's equations; so a representation of the same domain in digital computer would have a continuous medium, a discrete representational format, and a discrete domain.⁸

⁷ Katz's version is: "Call whatever stuff from which representations are constructed the medium of representation. Call whatever structure is imposed on that medium the format of representation." (2008, p. 404).

⁸ While neither Maley nor Katz mention domain formats, they implicitly acknowledge media format, e.g., in Maley's definition of 'continuous' (2011, p. 123), and in Katz's occasional reference to the

To sum, by extending Maley's and Katz's initial discussion we can distinguish between three different formats: The *medium format* (MF) is the mathematically characterized structure of the underlying physical processes that will be used to implement a representational scheme; the *representational format* (RF) is the structure of the scheme imposed on the MF; and the *domain format* (DF) is the structure of the domain to be represented by the representational scheme.⁹

According to the received view, the process of realizing a digital device involves 'imposing' a discrete mathematical structure on the physical world. In some cases (e.g., grains of sand), this can be straightforward because the medium format is itself discrete. Other cases are more complex. For example, building a digital electronic computer requires (roughly speaking) imposing a set of two discrete symbols on a continuous medium (voltage) by setting upper and lower thresholds $i < j$ such that $f(v) = \{0 \text{ if } v < i, 1 \text{ if } v > j, \text{ and undefined otherwise}\}$.

Any successful implementation of a discrete scheme requires that each token physical state of the implementing medium be reliably typed as an instance of one and only one representational type (or discarded as non-representational). This is because discrete representational systems do not allow for meaningful states 'between' or 'within' those identified in the scheme; the physical states must be typed so that they can even be used as representations by the device. Consequently, a practical consequence of the received view is that physical realizations of discrete representational schemes—i.e., digital devices—will include mechanisms to enforce this design goal.¹⁰

The need to impose a discrete structure is most apparent in how digital devices cope with the problem of noise: Voltages fluctuate and transitions between high and low states are not instantaneous. If we are trying to implement a discrete scheme with a continuous medium, the mechanisms for imposing that scheme must cope with this noise for purposes of representational typing. The received view thus predicts that digital devices will contain mechanisms dedicated to fulfilling the goal of reliably typing physical states in the face of noise.

This situation is, of course, what we find in digital electronic computers. In these devices, the necessity of imposing the threshold function leads to significant effort

Footnote 8 continued

physical medium (rather than the representational format) being continuous or discrete (e.g., Katz 2008, p. 404).

⁹ There are, of course, theoretical and practical interactions between MFs, RFs, and DFs. For instance, on the theoretical side, one cannot use a discrete RF (e.g., binary code) to exhaustively represent a continuous DF (e.g., real-valued solutions) because the cardinality of the latter is greater than that of the former. For the same reason, one cannot use a discrete MF (e.g., grains of sand) as a basis for a continuous RF, even if the MF were infinite. On the practical side, limitations on resources (e.g., finite memory or error introduced by noise) mean that no RF, regardless of its format, can fully represent an infinite (discrete or continuous) DF.

¹⁰ A popular engineering textbook makes this point as follows: "The discipline of discretization states that we choose to deal with discrete elements or ranges and ascribe a single value to each discrete element or range. Consequently, the discretization discipline requires us to ignore the distribution of values within a discrete element [and] this ... requires that systems built on this principle operate within appropriate constraints so that the single-value assumptions hold." (Argawal and Lang 2005, p. 4). That is, implementing a discrete RF requires mapping ranges of MF states to individual representational types, and this requires that "constraints" are imposed to guarantee that this design goal is satisfied.

being expended on building mechanisms whose purpose is to guarantee that states are reliably typed so that they can be used as (discrete) representations. First, to cope with fluctuations in voltage, the thresholds i and j are never equal; rather, they are chosen so that there is a gap between them, so that wavering voltages cannot be mistyped. Instead, the region between thresholds is undefined, i.e., it has no representational function, and the system enters an error state (c.f. Haugeland 1981, 1998, p. 79, example 3).

Second, since $f(v)$ requires comparing v to i and j , there must be a way to reliably measure v with respect to those thresholds. To accomplish this a *reference voltage* is introduced into the system. This additional circuitry produces a stable voltage against which the values of v , i and j can be compared. This is a significant engineering concession, since it requires increasing the size, power consumption, and operating temperature of integrated circuits.

Third, the propagation of voltages along wires is fast but not instantaneous, and signals have to travel different distances depending on their points of origin and reception. Therefore, a logic gate must coordinate the time at which it samples its inputs so as to access them only when they representationally relevant: The signal must be temporally quantized. Part of the solution is to introduce a central clock to provide a timing signal used by the millions of logic gates for the discrete sampling of inputs. Furthermore, for distances that would be too short, timing *delays* are introduced into the system, e.g., by using delay circuits or by lengthening the distance a signal has to travel.

To sum, according to the received view, digital representations are discrete, and this has specific consequences for how digital devices are engineered: To satisfy the need to type physical states as instances of one representational type rather than another, the system must contain support mechanisms for reliably ‘ignoring’ transitory fluctuations due to state transitions or noise. This broad methodology for coping with noise is precisely what we observe in actual engineering practices, e.g., in the form of clocks, reference voltages, and delay lines. In other words, what we see in the engineering of digital devices is a consequence of the discreteness of the representational scheme to be realized.

Not surprisingly, analog devices adopt a different strategy for coping with noise, paralleling the distinction between digital and continuous RFs. To begin, note that it is impossible for a discrete MF to instantiate a continuous RF, even in theory, since the cardinalities of the two formats differ. Consequently, on the received view, all genuine analog systems utilize a continuous MF, and the design goal is to use this underlying MF to instantiate a continuous RF. For instance, in a slide rule, the relative distance between points on the two rules varies continuously, and this distance represents an input value (a real).¹¹ Because of the close relationship

¹¹ This correspondence in structure between MF, RF, and DF tends to make the distinction between MF and RF ‘transparent’ in analog systems, as the continuous states of the MF appear to be directly related to the representational domain (a situation sometimes referred to as “native” representation (Argawal and Lang 2005, p. 44)). Transparency is not unique to analog systems; a similar situation occurs when a discrete MF is used to realize a discrete RF for purposes of representing a discrete domain, as in Maley’s example of grains of sand being used to represent an integer quantity. The lesson is that care is required to avoid conflating MF and RF.

between MF and RF in analog systems, the noise affecting the MF also impacts the RF, causing fluctuations in the representations themselves. That is, since representational kinds in analog systems differ by arbitrarily small degrees, any variation in the MF is a variation in the representational type, and a difference in what is represented.

As a consequence, how we deal with the noisy physical world in implementing analog computers is very different than in the digital case. In the latter, since a state cannot function as a representation unless it is typed, our goal is to quantize the signal *prior* to treating it as a representational vehicle. In contrast, in analog devices, since the representational type will always be subject to noise, the best we can hope for is a token representation that is ‘in the vicinity’ of the ideal result, i.e., within some margin of error: In analog devices we cope with the effect of noise on the representations themselves.

The design and implementation of analog computers reflects this aspect of the received view: A major engineering challenge for building analog devices is that of minimizing the margin of error introduced to the representational scheme through physical realization. So, for instance, engineers may develop materials and machine tools that can achieve greater precision under broader operating conditions (e.g., making larger slide rules out of materials less susceptible to variations due to temperature), engage in redundant calculations, or quantify and correct for error through the development of error models (Clymer 1993).

Taking stock, if associating ‘digital’ with ‘discrete’ and ‘analog’ with ‘continuous’ is appropriate, then doing so should be consistent with implementation differences. Digital devices are to be designed with the goal of unambiguously typing the MF so that the resulting states (of the discrete RF) can reliably fulfill their representational roles. In contrast, since analog devices utilize continuous media and representational formats, every difference in physical state is a difference in representational type. Therefore, noise in the MF results in misrepresentation, i.e., likely discrepancies between what is actually tokened and the desired ideal; this deviation is measured by a margin of error.

As we have seen, these consequences are borne out by actual engineering practices. The design and implementation of digital computers includes hardware dedicated to coping with error in the process of quantization—clocks, reference voltages, delay lines, etc.—and that of analog computers involves minimizing error through choice of materials, machine tools, and error models. Since the received view explains these differences in engineering practices, we have additional reason to conclude that the received view is an indispensable account of the distinction between analog and digital representation.

I conclude this section with two brief addenda. The first is that the received view allows for the possibility that some digital devices engage in approximate representational typing; error-free typing is not a necessary condition on a device being digital, but rather an ideal incurred by pursuing the goal of implementing a discrete representational scheme. For instance, a digital computer operating at the limits of its design parameters—e.g., under intense heat—should be expected to make typing errors, but this does not imply that it is no longer digital; rather, it is a malfunctioning digital device. Furthermore, if it is known that a digital device will

be operating in situations where it will likely mistype, it is possible to design the device so that it still behaves in reasonable ways even as it fails to achieve the ideal.

The second observation concerns the relation between the current account (which is based on literature in computer science and engineering) and earlier influential philosophical work, e.g., Goodman (1968), Lewis (1971) and Haugeland (1981, 1998). Goodman suggests that the distinction between digital and analog representation is that the former is ‘differentiated’ while the latter is not, which *prima facie* corresponds to the received view as I’ve characterized it: Digital representations are such that candidate tokens belong to one type rather than another (i.e., ‘differentiated’), whereas analog representations differ by arbitrarily small degrees (‘undifferentiated’), so any difference in a candidate token is a difference in type.

Lewis argues that some analog devices have differentiated states, and hence would be incorrectly categorized as digital on Goodman’s account. For example, suppose we have an analog circuit that computes some $f(x) = kx$, where x is an input value, and k is a user-adjustable coefficient. One way to set the value k is by using variable resistors, e.g., knobs that rotate to different positions. Lewis notes that these resistors come in different varieties: Some use contacts that vary (spatially) continuously along a wire, while others have contacts that incrementally move across a series of single-ohm resistors. According to Lewis, in the latter case we have ‘differentiated’ input voltages, yet the circuit is still properly deemed analog; hence ‘analog’ does not correspond with ‘undifferentiated’.

The present account avoids Lewis’ critique because it agrees the circuit as analog. First, the underlying circuitry still operates over the reals, i.e., a continuously-varying voltage, potentially interpreted as such. If we were to replace the stepped input with a smooth one, nothing else in the circuitry would have to be changed; rather, bracketing the influence of noise, doing so would ‘unlock’ additional precision afforded by the underlying continuous representational scheme. The second reason the circuit is analog is that digital representation requires quantization of the time dimension as well as the magnitude of the signal, for, as noted above, without time quantization the ‘temporal endpoints’ of the segment of signal to be quantized cannot be determined. The present circuit does not quantize time, and hence is not digital. Therefore, Lewis’ counterexample does not apply to the current interpretation of the received view.¹²

The present account is also consistent with Haugeland (1981). In contrast to Goodman, Haugeland focuses on the kinds of mechanisms (or ‘procedures’) that each kind of device uses for typing token instances of representational types. Digital devices use “positive procedures” which, when functioning properly, type instances with complete certainty, while analog devices use “approximation procedures”, which necessarily involve a margin of error in their typing. The account offered above coheres with this approach by showing how Haugeland’s distinction between positive and approximate procedures is related to the difference in the formal

¹² My interpretation of the received view also agrees with Lewis’ second proffered counterexample to Goodman (it is analog) and with his example of a digital system (it is digital). These observations raise the issue of how to properly interpret Goodman’s version of the distinction in comparison to my own, but this project is outside the scope of the present paper.

structure of a representational scheme: Digital devices use positive procedures because they aim to reliably type token states into discrete categories, while analog devices use approximation procedures because continuous representational schemes can never be typed without error. The received view is thus complementary to Haugeland's, despite his emphasis on implementation over formal structure.

Applying the Standard Interpretation

At this point I've argued that the standard interpretation is an indispensable part of the analog–digital distinction: It is part and parcel of computational theory, and explains why the engineering of analog and digital devices differ in their treatment of error. However, as noted in the introduction, the received view conflicts with recent alternatives offered by Maley (2011) and Katz (2008). In this section I illustrate these incompatibilities and—on the assumption that the received view does not require additional defense—consider whether these departures can be justified as alternative senses, entrenched in those contexts cited by the authors as inspiration for their respective analyses, or whether they lack such support.

Maley (2011)

Maley defines a continuous representation as “one that takes on continuous values, either in the representational medium or the representational format or both” (p. 123), and a discrete representational scheme as one where “representations are distinct from other representations in the same representational scheme, (i.e.) there are gaps between the possible representations, and representations come only in wholes.” (p. 125) The former definition must be amended to avoid the potential conflict with the latter: Electronic computers use continuous media yet have discrete RFs, so if a continuous MF is sufficient for continuous representation, typical computers are simultaneously continuous and discrete. This can be avoided by reiterating the distinction between MF and RF; aside from this adjustment, Maley's definitions of ‘continuous’ and ‘discrete’ are consistent with those of the received view.

The more interesting cases in Maley's analysis concern digital and analog representation. Digital computers (and the engineering practices surrounding them) are the basis for his account of the former. Since all of these devices use a place-value mapping (usually 32 or 64-place binary), and the results of such mappings are integer contents, Maley concludes that digital representations are those that (1) represent integer values, and (2) do so via positional notation. (p. 125)

This analysis imposes restrictions on both *how representations represent* and *what they represent*, and each contrasts with the standard interpretation. For instance, consider your current laptop computer, which uses a binary place-value representational scheme. Suppose that, without adjusting any physical feature of that device, we adopt a summative rather than positional mapping from physical states to integer contents, so that, e.g., ‘101’ represents two rather than five. According to Maley's definition of ‘digital’, this shift in how states are interpreted

preserves the discreteness of your laptop's representational system while rendering it non-digital. In contrast, the received view takes any discrete representational scheme to be indicative of a digital device, regardless of how (or if) it encodes integers.

Maley's restriction that "a digital representation necessarily represents a number" (p. 125) is also a significant departure from the received view. For example, suppose we have ordered triples of tokens constructed from two symbols, 'a' and 'b', and we adopt the following semantics: In position 0, 'a' means 'Adele' and 'b' means 'Bill'; in position 1, 'a' means 'hates' and 'b' means 'loves'; and in position 2, 'a' means 'cats' and 'b' means 'dogs'. The semantic content of any given triple is the compositional outcome of the meanings of the constituents interpreted in accordance with English grammar. In this case the strings are binary but have non-integer contents. More generally, according to Maley's account, *any* discrete representational system whose (original or derived) function is to carry information about something besides integers is non-digital. The contrast between Maley's account and the received view is obvious: Just as the received view does not restrict digital devices to a certain way of representing integers, it does not limit those devices to *representing integers*.

To sum, Maley's account of 'digital' introduces a new category of device: The 'non-digital' discrete computer. Such a device (1) uses a discrete RF, yet (2) either represents integers without using a place-value interpretation, or does not represent integers at all. In other words, Maley's proposal is that 'digital' should be restricted to a proper subset of those devices identified as such by the standard interpretation, since every device that is digital on Maley's account is also digital according to the received view, but the converse is not true.

This being said, there do seem to be contexts where 'digital' is used in this restricted way. For example, Maley appeals to ASCII code in defense of his analysis: ASCII encodes characters as natural numbers, and those numbers are themselves represented using binary place-value notation. This particular argument may nonetheless leave one suspicious that there exist alternatives; after all, ASCII is an arbitrary conventional numerical code for the representation of text. Perhaps there exists a distinction between bit patterns interpreted as instructions (e.g., 'move the contents of register A to register B') and those interpreted as numbers (e.g., ASCII code). However, a brief glance at an assembly language text suggests this worry is unfounded:

Everything inside the computer is indicated as a number. It is what the number represents that determines the difference between one thing and another. Numbers may represent instructions for the computer to perform specific actions (a program), values used in calculations (data), or characters to be printed (ASCII code). (Howe 1981, p. 6)

This passage is a random sample drawn from my bookshelf, but it coheres with more recent texts (e.g., Argawal and Lang 2005). Therefore, Maley's claim that there is a sense of 'digital' corresponding to the place-value representation of natural numbers appears to be correct, provided we limit ourselves to the context of 'industry-standard' or 'popular' architectures.

Regardless, this sense is overly restrictive as a general account. First, the dominant usage of ‘digital’ in the context identified by Maley is not merely place-value, but base-2. So, if we take context of usage seriously, it seems we should restrict ‘digital’ to referring to binary code. But researchers continue to consider systems with bases of greater value (e.g., base three or five). The received view captures what these alternatives have in common with binary schemes; while it may be that binary place-value encodings are common, it seems premature to project this definition to all possible digital devices on pain of obscuring their shared features.

A similar point can be made with respect to semantics. One of the reasons philosophers are interested in the analog–digital distinction is that it may provide insight into the format of mental representation. If we restrict ‘digital’ to applying only to representational schemes with integer contents, then it’s a trivial matter to argue that minds and brains are not digital: Mental states and brain states represent things besides integers. In other words, philosophical tradition treats ‘digital’ as neutral regarding semantics (cf. Searle 1990).

I thus propose that we treat Maley’s ‘digital’ as a distinct sense, one that defines a proper subset of the more general category of digital representations. Within ‘industry-standard’ or ‘popular’ contexts of use, the word ‘digital’ means ‘a discrete representational scheme that represents integers using base- n (or binary) place-value notation’. However, the received view accommodates the more general sense underlying what all digital systems have in common, i.e., the theoretical foundations that drive their construction. For the remainder of this paper, I will use ‘digital_E’ to refer to this sense; ‘digital’ will continue to be used in the received way.¹³

Maley’s proffered account of ‘analog’ is also divergent. After describing his alternative, I (1) show how it contradicts the standard interpretation in such a way that it cannot be viewed as identifying a subset of standard analog devices, and (2) question whether it is indeed an adequate characterization of the way the terms are used in the mental rotation literature. Finally, since it seems to be an open question whether there is an agreed upon sense of ‘analog’ in the mental rotation literature, (3) I conclude that the same strategy used for digital_E—clarification and accommodation—is appropriate.

To arrive at an analysis of ‘analog’, Maley considers an influential usage of the term from cognitive science: the case of mental rotation. For example, it is well established that subjects’ response times on visual matching tasks vary

¹³ In the preceding discussion I interpret Maley as asserting his account of ‘digital’ should replace the standard interpretation. An anonymous reviewer suggests that I misinterpret Maley’s proposal: Rather than offering an account of ‘digital’ intended to supplant the received view, he is doing nothing more than what I propose, namely, identifying a distinct sense as it appears in a particular context. In response, I believe my interpretation is justified. For example, Maley writes, “I claim that the term ‘digital’ should be reserved only for representations of this type [i.e., those that represent integers using a place-value scheme], rather than discrete representations more generally.” (p. 125). In other words, Maley asserts we should limit ‘digital’ to a subset of discrete systems instead of allowing the term to apply to all discrete systems as specified in the standard interpretation. Regarding ‘analog’ and ‘continuous’, Maley is not as straightforward, writing “I have presented a distinction between analog and continuous representation, and suggested that this distinction be adopted on the grounds that it provides a more useful way to classify representations” (p. 124). However, given his treatment of ‘digital’ and ‘discrete’, I assume his proposal regarding ‘analog’ and ‘continuous’ is also one of replacing the standard view.

proportionally to the distance a probe image must be mentally rotated to be compared to a target (Shepard 1978). In his discussion of these results, Shepard contrasts ‘logical’ (i.e., ‘digital’) and ‘analog’ processes as alternative explanations. When we rotate a physical object or image from position x to position z , it travels through intermediate degrees of rotation, e.g., y . According to Shepard, digital processes do not *intrinsically* exhibit this property, e.g., a matrix transformation from states x to z has no guarantee of involving an intermediate result corresponding to state y , and hence cannot explain the behavioral results. In contrast, Shepard asserts, analog processes do have this property, so the underlying mechanisms must be analog.

Based on Shepard’s usage, Maley defines an analog representation as follows:

A representation R of a number Q is analog if and only if: (1) there is some property P of R (the representational medium) such that the quantity or amount of P determines Q ; and (2) as Q increases (or decreases) by an amount d , P increases (or decreases) as a linear function of $Q + d$ (or $Q - d$). (Maley 2011, p. 123; formatting adjusted)¹⁴

An analog representation is one where a quantitative description of a property of the medium format varies monotonically with a quantitative description of the property being represented.¹⁵ For example, in the case of mental rotation, there is some relevant physical property P of those brain states that represent the orientation Q of an object, and P is monotonically related to Q , in the sense that as the quantity of P increases (firing rate, oscillation rate, or some other property), the quantity of Q increases (degree of rotation).

Unlike digital_E, this sense of ‘analog’ is unrelated to the standard interpretation, since it implies an analog representational scheme can be simultaneously discrete. Consider a unary Turing machine that enumerates the natural numbers using a singleton alphabet, implemented using a paper tape upon which instances of the lone symbol are inscribed. Because this is an implementation of a Turing machine, it qualifies as digital according to the received view, and because it uses a discrete (base-1) place-value scheme to represent integers, it is also digital_E. Furthermore, because there is a monotonically increasing relation between the value being represented and the physical medium doing the representing (i.e., the amount of paper and ink), it is also analog on Maley’s account. In contrast, the received view does not allow a device to be simultaneously analog and digital; it can implement a digital RF using an analog MF, or implement an analog RF using an analog MF, but it cannot be simultaneously analog and digital.

¹⁴ The restriction that the relationship between P and Q be linear is relaxed to monotonicity in a footnote.

¹⁵ My proposed abbreviated version of Maley’s definition is intended to gloss over some complications for purposes of facilitating discussion. Most notably, R is said to be a ‘representation of a number Q ’, which suggests that analog representations have restricted contents (like digital_E representations). For present purposes I assume that analog representations are *not* restricted to representing numbers – they could represent colors, propositions, objects, etc. The important point is that what is represented can be described quantitatively, which, taken in tandem with a similar description of the physical parameter P , results in a (monotonic) function from the latter to the former.

Note furthermore that unary Turing machines are theoretically interesting because they define a class of systems allowing pseudo polynomial solutions to otherwise non-polynomial problems. That is, unary digital representational schemes can tractably solve certain problems that base-2 and greater schemes cannot—but only in exchange for a (proportionally massive) increase in memory (n digits to represent integer n in comparison to $\log_b n$ digits for bases $b > 1$). Consequently, the class provides a nice illustration of fundamental relations between how a domain is represented, how efficiently problems involving that domain can be solved using representations with that structure, and how much memory is required to do so. However, unlike the standard interpretation, Maley's definition of 'analog' cuts the class of unary Turing machines in two. For instance, a stack can be implemented using a representational medium whose relevant physical parameter exhibits the following pattern: $f_{START}(1) = 1$, $f_{PUSH}(n + 1) = -0.5f(n)$, and $f_{POP}(n - 1) = -2f(n)$, where n is the number of symbols in the stack. The push and pop operations are non-monotonic, and since two stacks can be used to simulate a Turing machine tape, a unary Turing machine implemented in this way will not be analog on Maley's account, despite being *computationally identical* to the same formal automata implemented using a paper tape. The point is simply that not only does Maley's proposed definition of 'analog' treat some discrete systems as analog, it also imposes distinctions between types of discrete systems that, from a computational perspective, should be grouped together; in contrast, the received view is consistent with computational theory.

This, of course, does not mean Maley's account should be discarded; instead, we must ask whether it amounts to an entrenched, alternative sense. If it is entrenched in the psychological literature on mental rotation, then it is reasonable to expect there to be agreement across researchers in that field regarding the meaning of 'analog'. Whether there is such agreement is difficult to resolve. For instance, here is an alternative interpretation of Shephard's appeal to analog representation: Consider a prototypical mechanical analog computer such as a slide rule. This device performs multiplication by adding the logs of two input values, and this is accomplished by positioning two logarithmic scales (the 'rules') in parallel and adjusting their spatial relationship with respect to each other. A side effect of the way these devices are implemented is that the distance one must move the rules is proportional to the input values—the higher the input values, the further one must move the components. This means that (1) to set the system to calculate higher values, it must travel through states representing results for lower values, and (2) since moving these components takes time, the higher the input values the longer it takes to calculate the result (assuming the rules are moved at a constant speed). In other words, slide rules exhibit the basic behaviors motivating Shephard's use of 'analog': The calculating of larger results requires going through the calculations of the smaller results, and this takes time proportional to the values being multiplied. Consequently, perhaps it is not monotonicity that is central to Shephard's appeal to analog representation, but rather that *the behavior to be explained is a direct consequence of the way the representational system is implemented*.

This seems to be the interpretation adopted by Pylyshyn (1981) in his influential discussion of mental imagery. On his account of analog representation,

the observed [input–output function expressing the relation between distance and reaction time] has that form as a consequence of the intrinsic lawful relations that hold among the particular physical properties that in fact represent distance and mean speed in the brain. ... This corresponds to what I would call the *analogue* view. (Pylyshyn 1981, p. 19)

As in the case of the slide rule, we get a relation between values of inputs and the time it takes to compute them because the device is built in a certain way, one that requires—through ‘intrinsic lawful relations’—that the computation of larger values go through the same physical routines of computing the same function for smaller values. Differences in reaction time are a direct consequence of how the system is realized. The emphasis on this interpretation is not monotonicity, but rather the nature of the implementing mechanisms, and those mechanisms reflect the fact that analog devices use continuous representational schemes.

If this interpretation is correct, then the contrast class is *not* non-monotonic systems. Instead, according to Pylyshyn, the contrast class consists of devices whose processing characteristics are not constrained by such intrinsic lawful relations; rather, they can be adjusted by tacit knowledge:

a subject makes [the relationship between distance and reaction time] come out ... because he or she has tacit knowledge of [this relationship]. In other words, regardless of the form of his or her representation, the subject knows [the relationship] holds in the world and therefore makes it be the case ... that the amount of time spent imagining the scanning will conform to this relation. (p. 19)

In contrast to analog devices, in this case the underlying implementation *could* function without respecting the proportional relationship between distance and reaction time. However, subjects *know* that this relationship holds in mind-independent objects, and this knowledge is used to tailor performance to respect this real-world phenomenon. Pylyshyn’s account of ‘analog’ thus seems to be broader than Maley’s, because it allows for the possibility of non-monotonic relationships qualifying as analog, and for monotonic relationships to not be analog: What matters is *where* the behavior comes from—intrinsic lawful relationships or tacit knowledge—not *what* that behavior is. Furthermore, the observation that analog (i.e., continuous) systems exhibit certain characteristic behaviors (in this case, a monotonic relation between degree of rotation and magnitude of representing medium) is consistent with the standard interpretation. The challenge to Maley’s account, then, is explaining away the possibility that the analog–digital distinction as it appears in the mental rotation literature is not simply a consequence of the standard interpretation and the implications it has for the engineering of analog devices.

Suppose, however, it can be demonstrated that the usage observed in the mental rotation literature cannot be boiled down to specific aspects of the standard interpretation. In that case, I believe sufficient reason has been given above to treat this alternative usage as a distinct sense rather than as cleaving the association between ‘analog’ and ‘continuous’ (cf. Maley 2011, p. 127). First, we’ve already

seen that the standard interpretation is central to computation and engineering practice. And second, as pointed out above, Maley's proposal treats paradigmatically discrete systems—unary Turing machines—as simultaneously analog and discrete, which clearly invites confusion given the ubiquity of the standard interpretation. The solution, then, is the same as for digital_E: If the usage present in the mental imagery literature cannot be interpreted as a version of the standard interpretation, allow there to be multiple, distinct senses of 'analog', one of which is the received view ('analog'), while another is unique to mental imagery research; call it 'analog_M'.

Katz (2008)

I've shown that Maley's versions of 'digital' and 'analog' conflict with the received view, and argued that to the degree they are entrenched in particular contexts, they ought to be treated as distinct senses rather than as displacing the standard interpretation. In this section I subject Katz's (2008) account to a similar treatment, except I argue that it is not justified.¹⁶

Rather than appealing explicitly to how the distinction is used in a particular context, Katz motivates his analysis through four hypothetical devices (p. 404):

- A Water is measured in increments of size x (e.g., using a shot glass), and the quantity of water as measured in those increments represents some quantity n . So, for example, if four shots of water are poured into the beaker, then the water in the beaker represents the quantity four.
- B Marbles of ten distinct colors corresponding to the values 0 through 9 are placed into n cups arranged in a row and ordered 0 to $n-1$ from right to left. The value represented is calculated according to base-10 positional notation.
- C Beakers of water from example A are arranged into an ordered n -tuple, and interpreted according to the place-value notation used for system B. So, for example, if there are three beakers, with the location 0 containing four shots of water, location 1 containing none, and location 2 containing two, then the array of beakers represents the quantity 204.
- D Large quantities of marbles (e.g., oil drums) are poured into a very large beaker (e.g., a missile silo). As in system A, the quantity represented is simply the number of units of marbles poured into the container; so, if twenty drums of marbles are poured into the silo, the marbles in the silo represent the quantity 20.

¹⁶ To be clear, whereas Maley considers specific contexts (binary digital engineering and mental rotation research) in order to uncover alternative senses, Katz does not explicitly adopt the same strategy; rather, he begins by noting the importance of the analog–digital distinction to numerical cognition literature but justifies his account by appealing to his thought experiment, not the literature. A reviewer suggests that this makes my subsequent consideration of the numerical cognition literature misplaced. Nonetheless, as I argue below, (1) Katz's account contradicts the received view, (2) his thought experiment does not represent the received view, and (3) his thought experiment is very similar to the 'mental accumulator' model of numerical reasoning as it appears in the literature he cites. Since his proposal departs from the received view, is not motivated by the received view, and is seemingly based on a model drawn from the numerical cognition literature, it is reasonable to consider whether his nonstandard alternative can be justified by appeal to that literature, even if Katz does not explicitly pursue that strategy. After all, the goal is to see if the nonstandard interpretation can be justified, and there is nowhere else to turn.

It is important to note that, according to the received view, all four of these hypothetical devices utilize a digital representational scheme. First, their representational format is that of the natural numbers, i.e., they are discrete. Second, each of the examples accommodates error in the medium format in the process of representational typing. For instance, in system *D*, we don't care about the exact number of marbles; rather, we care about the number of drums of marbles. By using drums as the measurement we have implicitly corrected for noise in the precise number of marbles (the MF) prior to semantically interpreting the state of the silo, in the same way digital electronic computers accommodate error in the voltage of a wire before classifying it as zero or one. The same observation applies to *A* and *C*. Case *B* is a limit, i.e., the mapping from MF to RF is transparent (i.e., 'native') because there is no need to accommodate noise in the MF. So all four systems are digital according to the received view. In other words, Katz is not addressing the received view with his thought experiment, because none of the devices are continuous according to that view.

Regardless, Katz's working assumption is that systems *A* and *D* are analog, while systems *B* and *C* are digital; consequently, the task of his analysis is to figure out why this is so. Bracketing the fact that all four are digital on the received view, here is an attempt at making his reasoning explicit: First, system *A* has two potentially relevant features: (1a) it uses a *continuous* MF (water), and (2a) it is *unary* (i.e., it implements a base-1 scheme with a single atomic symbol, realized as instances of quantity *n*). System *B* also has two potentially relevant features: (1b) it uses a *discrete* MF (marbles), and (2b) it has a *decimal* place-value semantic interpretation. Second, we can eliminate (1a) and (1b) (the MF) as a contender for what makes a system analog or digital, because we can implement a decimal place-value scheme using a continuous medium (as shown by system *C*) and a unary representational system using discrete media (as shown by system *D*). This means that what is relevant has something to do with the base-one/base-ten distinction. Third, Katz notes that one feature shared by the unary systems *A* and *D* is that their representational schemes do not depend on tracking individual elements of their respective physical media; in each case you could vary, within some degree of tolerance, the exact number of molecules or marbles contained in each increment without changing the representation tokened. In contrast, in the base-10 system *B*, if you swap a marble for one of a different color, the representation changes; *B* does not tolerate variation in its representational medium. The same holds for the base-10 system *C*: If you add or subtract an increment *n* of water from one of the beakers in the array, the representation changes. Therefore, Katz concludes, the relevant difference between the two pairs is that *A* and *D* tolerate variation in their representational media, while *B* and *C* do not.

Katz summarizes his conclusion in terms of the consumers—i.e., the *users*—of the respective representational schemes:

whether or not a representational system is analog or digital may turn on facts about the user of the system. For example, if the perceptive powers of a particular human being were sufficiently enhanced, then the representations of

[a system using water as a representational medium] would be readily distinguishable to that user. ... It would therefore be a digital, rather than an analog system. (p. 405)

In other words, *A* and *D* are analog because users of the scheme do not track individual increments of the representational media (molecules or marbles), while *B* and *C* are digital because users do track individual increments of the representational media (colors and increments of size *n*).

Katz's argument arguably begs the question;¹⁷ however, here I grant Katz's intuitions that *A* and *D* are analog and that *B* and *C* are digital, and ask how this relates to the standard interpretation. On Katz's view, the analog–digital distinction is tethered to contingent facts about users—namely, their ability to distinguish individual elements of the medium—and consequently a device can change from digital to analog by holding the discriminatory capacities of the user fixed while modifying the device. For example, since users cannot count individual molecules in system *C*, there is a lower limit to the size of *n* below which the user cannot reliably discriminate increments. In the original version of *C*, each beaker holds between zero and nine increments of water of size *n*, and the user is assumed to be reliable in distinguishing between a beaker containing *i* increments and one containing $i \pm 1$ increments. Suppose, then, we construct a new system, *C**, by modifying *only* the base of the scheme from decimal to some base *b* such that the user can no longer discriminate nearby quantities represented by individual beakers in the array. To do this, we reduce the size of the increments of water poured into each beaker, from *n* to $m < n$ so that each beaker is divided into $b-1$ units rather than nine. According to Katz's analysis, *C** is now analog, because the user can no longer 'readily discern' nearby values. In contrast, according to the received view, the base-*b* representation of the naturals is as digital as unary, binary, octal, decimal, hexadecimal, or any other finite base place-value encoding, because they are all discrete.

¹⁷ His justification for labeling *C* 'digital' is that, despite not being able to track individual water molecules, a human user "can readily distinguish representations of some number *n* from representations of other nearby numbers." (p. 405) So, for example, John can't tell how many molecules there are in beaker #2, but he can tell that there are two (rather than one or three) increments of water in that beaker, and hence distinguish the overall representation of 302 from 402 and 202. Now, for this claim to hold of device *C*, it must also hold for each individual beaker in the array, and since device *C* is built from instances of device *A*, it follows that the user can readily distinguish nearby representations in device *A*. Since the only difference between *A* and *D* is their medium, the fact the user cannot discriminate individual molecules in device *A* should not be an impediment to the user distinguishing representations in *D*, either – after all, it's the same situation only rendered at a different physical scale. Yet, Katz unexpectedly asserts that the user cannot discriminate nearby representations in case *D*:

Because a large number of marbles are employed in each increment, the user will likely be unaware exactly how many marbles are employed in any given representation. Because of this, the user is likely to be unable to readily discern whether a representation is of some number *n*, or whether it is of some other relatively nearby number. (p. 405)

Substituting 'molecules' for 'marbles', it seems that the premise given in this passage holds for systems *A* and *C*, and so by parallel reasoning the user of *C* should not be reliable in discriminating nearby representations if they are not for system *D*; or, if they are reliable when using *C*, they should also be reliable in using *D*. In short, somewhere in this chain of inferences Katz is begging the question that the discriminatory capacities of the user are relevant to the analog–digital distinction.

This same example illustrates how, on Katz's account, a device could be changed from analog to digital by holding the representational scheme fixed while varying the discriminatory capacities of the user. For example, if we swap out the unreliable user of C^* for one who is reliable, C^* is once again digital. Furthermore, if we keep both users on hand, C^* is simultaneously analog and digital. In contrast, the received view grants that building a robust digital device requires implementing a reliable user, but holds the lack of a reliable user to imply only an imperfect attempt at engineering a digital system, not a shift from digital to analog representation.

These examples illustrate how Katz's analysis is nonstandard; but, as we saw above, departures from the received view may be justified by the existence of contexts of use wherein the proposed alternatives have lives of their own. So we can ask whether Katz's version is so entrenched. Katz introduces his account by referencing work on the cognitive psychology of numerical reasoning, and while he does not explicitly discuss the relationship between that research and his version of the analog–digital distinction, systems A and D (the 'analog' systems) bear a significant resemblance to the *mental accumulator* model for numerical memory. This model is conceptualized as a beaker into which uniform increments of liquid are poured; when it comes time to recall the stored value, it is read off of the accumulated liquid, but, since the system is subject to noise—the water in the beaker 'sloshes' around—the recall is less than 100 % reliable. Furthermore, experimental evidence indicates that numerical memory is scalar, which is to say that it is less accurate the larger the value to be stored. This is accommodated by the model in that representing large values necessitates using smaller increments of water (because either the amount of water or the size of the beaker is fixed), in turn requiring more precise discrimination (and hence more error). Katz's argument thus resembles the mental accumulator model in that (1) it appeals to beakers filled with liquid as a motivating illustration, and (2) it claims that users (i.e., humans) are less reliable the greater the value being represented. Given these parallels, perhaps Katz's account can be contextualized: Within the field of research on the cognitive psychology of numerical reasoning, the analog–digital distinction is tied to the reliability of the user, not features of the MF or RF.

It's not clear that this strategy will work, as the sense in which 'analog' and 'digital' are used in the context of numerical reasoning is consistent with the received view. For instance, Katz refers to Gallistel et al. (2006; henceforth GGC), who ask whether humans represent integer quantities based on a prior capacity to represent reals, or instead use a scheme for representing integers to approximate reals. From the perspective of the received view, GCC's question is interpreted as follows. First, since neuronal information processing models often appeal to continuous values, let us assume that the MF provided by neuronal activity is continuous. The goal, then, is to explain how humans are capable of representing both real magnitudes and discrete integers using a continuous medium. There are two possibilities:

- (P1) A digital scheme is used to represent both integers and reals. The neural MF is first constrained so that it implements a discrete RF (e.g., analogous to the way the step function imposes a digital symbol set on the voltage of a wire), and the resulting symbol set is used to represent both integers and reals.

(P2) An analog scheme is used to represent reals, and a digital scheme is used to represent integers. The continuous properties of neurons are themselves used as native representations of real-valued magnitudes. Furthermore, a discrete structure is imposed on those same properties, and the resulting symbol set is used to represent the integers.¹⁸

The received view thus generates an intelligible interpretation of the question being asked. Furthermore, GGC's answer is consistent with this interpretation:

it is the real numbers, not the integers, that are the primitive foundation of numerical reasoning. The integers are a special case whose prominence in the cultural history of numbers derives from the discrete character of language. When a discrete system like language attempts to represent a quantity, it will find it much easier to represent a countable (discrete) quantity than an uncountable (continuous) quantity. (p. 265)

There are several aspects of this passage that deserve emphasis. First, it coheres with the interpretation just offered: The correct option is (P2). Second, insofar as the authors appeal to cardinality, the passage favors a mathematical interpretation of discreteness and continuity, a perspective shared by the received view. Furthermore, the received view also arguably underlies GGC's arguments in favor of option (P2). For example, regarding the relation between noise and representational format, they write,

When a device such as an analog computer represents numerosities by different voltage levels, noise in the voltages leads to confusions between nearby numbers. If, by contrast, a device represents countable quantity by countable (that is, discrete) symbols, as digital computers and written number systems do, then one does not expect to see [that] kind of variability (p. 252).

Their point is that digital systems are affected by noise in ways different their analog counterparts. For example, in binary representation, if noise is applied to each bit (so that each digit in a binary representation as some chance of being misconstrued), then certain patterns of error are more probable than others, e.g. three will be conflated more frequently with seven than with four. This consequence of noise on represented value is a natural feature of digital representation as articulated above (“[Making sense of implementation](#)”). Hence it appears that GGC are working with the received view of ‘digital’. A parallel observation holds for their use of ‘analog’ in this passage.¹⁹

As this brief discussion illustrates, the research on numerical reasoning cited by Katz appears to adopt the received view of the analog–digital distinction. Taken

¹⁸ Technically there are two possibilities: The first is that a full-blooded digital scheme is imposed on the underlying continuous MF, and the second is that a continuous RF is ‘notched’ (e.g., stable attractors are added) so that the system tends to settle into integer values. In this case an analog scheme is used to represent both reals and integers. In an effort to simplify the exposition, I’ve deliberately collapsed these two possibilities into one.

¹⁹ The authors also equate ‘discrete’ with ‘countable’, which is precisely how the received view conceives of discrete representation.

together with the counterintuitive consequences of Katz's proposal, additional argument is needed to establish that the senses proposed by Katz should be treated as entrenched alternative senses.²⁰

Conclusion

The received view holds that the distinction between analog and digital mental representation is properly understood as turning on two types of mathematical structure: Analog representations share the structure of the reals (they are continuous), while discrete representations share the structure of the integers (they are discrete). In this paper I've supported the standard interpretation by surveying how it (1) coheres with the formal foundations of computational theory, and (2) explains the engineering of analog and digital devices (especially in terms of how error is handled). I've also considered at length two recent alternatives; both depart significantly from the received view, but they differ in that one (perhaps) identifies distinct senses of the terms, unique to specific contexts of use (digital_E , digital_M), while the other lacks such justification.

There is a broader concern to be gleaned from this discussion. Whereas the received view draws its inspiration first and foremost from formal, mathematical concepts, the strategy pursued by both Maley and Katz emphasizes 'implementational contexts', i.e., situations where the terminology is used to describe representational devices such as brains or artifacts. For instance, Maley's account of digital_E appeals to the construction of industry-standard digital devices, his treatment of analog_M appeals to how physical magnitudes relate to representational domains, and Katz's account indirectly attempts to characterize how the terms are used in literature on numerical reasoning. On the one hand, these approaches *neglect* the theoretical foundations of the analog–digital distinction, and in doing so run into conflict with how the terms are understood by computer scientists, mathematicians and electrical engineers. On the other, we must also acknowledge there currently

²⁰ Katz also argues that his account is implied by Haugeland (1981, 1998), and hence draws support from that source. This seems to be based on a misinterpretation. As mentioned in "Making sense of implementation" section, Haugeland claims that analog devices are distinguished from digital devices in that the former use *approximation* procedures for reading and writing representational states while the latter use *positive* procedures. Approximation procedures are subject to noise – they are uncertain about the type of representation they read or write – and hence can only represent within a margin of error. In contrast, positive procedures reliably succeed in typing representational tokens. Katz interprets Haugeland as claiming that an analog device can be transformed into a digital one by making the read-write mechanisms more tolerant to noise (so that the error in typing token representational states is eliminated) and that a digital device can be transformed into an analog one by making the typing criteria more strict (so that error is introduced). This is only partially correct: The reason analog devices use approximation procedures is that they implement continuous representations, and continuous representations are always subject to noise ("Applying the standard interpretation"). This is why Haugeland notes that *no amount of technological innovation can eliminate error* in analog devices (p. 83). In contrast, digital devices use positive procedures because they must quantize the representational medium, and this process does not tolerate any error if the resulting states are going to be treated as representations. In short, Haugeland's account is properly viewed as an articulation of the engineering consequences of the received view.

exist cases where these alternatives are sufficiently entrenched to render it futile (not to mention presumptuous) to object. Consequently, we should recognize there are *varieties* of analog and digital representation.

That being said, someone might ask, ‘Is the point simply that different groups propose different definitions?’ I am not claiming that how we interpret the analog–digital distinction is merely a verbal issue; instead, the standard interpretation is correct, because it is grounded in a well-understood mathematical distinction. However, this does not mean there cannot be nonstandard interpretations that play important roles in particular contexts. In this way the analog–digital distinction is like computation: There are standard models (e.g., Turing machines) and nonstandard alternatives (e.g., asynchronous time digital devices). But the fact there exist nonstandard models does not mean we give up the original account; instead, we retain it and use it as a touchstone for assessing the benefits and drawbacks of nonstandard versions. Furthermore, as illustrated above with my discussions of Maley’s ‘digital’ and Pylyshyn’s ‘analog’, possible alternatives are nonetheless related to the received view; the standard interpretation is the trunk of a tree from which nonstandard branches emerge. So the position taken in this paper is not that disagreements over the analog–digital distinction are merely verbal. Rather, the point is that the standard interpretation is primary, so rather than treating nonstandard alternatives as contenders for a throne, a strategy of clearly articulating their relation to the standard and (possibly) accommodating them as distinct senses is a reasonable strategy.

Acknowledgments Thanks to Corey Maley and two anonymous reviewers for valuable comments on an earlier draft of this paper, and to Bill Schonbein for teaching me how to use a slide rule.

References

- Argawal, A., & Lang, J. H. (2005). *Foundations of analog and digital electronic circuits*. San Francisco: Morgan Kaufmann.
- Blum, L., Shub, M., & Smale, S. (1989). On a theory of computation and complexity over the real numbers: NP-completeness, recursive functions and universal machines. *Bulletin of the American Mathematical Society*, 21(1), 1–46.
- Bournez, O., & Campagnolo, M. L. (2008). A survey on continuous time computations. In: S.B. Cooper, B. Lowe, A. Sorbi. (Eds.), *New computational paradigms: Changing conceptions of what is computable* (pp. 383–423). New York: Springer.
- Church, A. (1936). An unsolvable problem of elementary number theory. *American Journal of Mathematics*, 58, 245–363.
- Clymer, A. B. (1993). The mechanical analog computers of Hannibal Ford and William Newell. *IEEE Annals of the History of Computing*, 15(2), 19–34.
- Gallistel, C. R., Gelman, R., & Cordes, S. (2006). The cultural and evolutionary history of the real numbers. In S. Levinson & P. Jaisson (Eds.), *Culture and evolution* (pp. 247–273). Cambridge: MIT Press.
- Goodman, N. (1968). *Languages of art*. New York: Bobbs-Merrill.
- Haugeland, J. (1981). Analog and analog. *Philosophical Topics*, 12, 213–225 [Page references are to the reprint appearing in Haugeland (1998)].
- Haugeland, J. (1998). *Having thought: Essays in the metaphysics of mind*. Cambridge, MA: Harvard University Press.
- Hopcroft, J. E., & Ullman, J. D. (1979). *Introduction to automata theory, languages, and computation*. Reading, MA: Addison-Wesley.

- Horgan, T., & Tienson, J. (1996). *Connectionism and the philosophy of psychology*. Cambridge, MA: MIT Press.
- Howe, H. S. (1981). *TRS-80 assembly language*. Englewood Cliffs, NJ: Prentice-Hall.
- Katz, M. (2008). Analog and digital representation. *Minds and Machines*, 18, 403–408.
- Koiran, P. (1994). Computing over the reals with addition and order. *Theoretical Computer Science*, 133, 35–47.
- Lewis, D. (1971). Analog and digital. *Noûs*, 5(3), 321–327.
- Maley, C. (2011). Analog and digital, continuous and discrete. *Philosophical Studies*, 155, 117–131.
- Post, E. L. (1936). Finite combinatory processes. Formulation I. *Journal of Symbolic Logic*, 1, 103–105.
- Pylyshyn, Z. (1981). The imagery debate: Analogue media versus tacit knowledge. *Psychological Review*, 88, 16–45.
- Sarpeshkar, R. (1998). Analog versus digital: Extrapolating from electronics to neurobiology. *Neural Computation*, 10, 1601–1638.
- Schonbein, W. (2005). Cognition and the power of continuous dynamical systems. *Minds and Machines*, 15(1), 57–71.
- Searle, J. R. (1990). Is the brain a digital computer? *Proceedings and Addresses of the American Philosophical Association*, 64, 21–37.
- Shepard, R. N. (1978). The mental image. *American Psychologist*, 33, 125–137.
- Siegelmann, H. T., & Sontag, E. D. (1994). Analog computation via neural networks. *Theoretical Computer Science*, 131, 331–360.
- Turing, A. M. (1936). On computable numbers, with an application to the Entscheidungsproblem. *Proceedings of the London Mathematical Society*, 42(series 2), 230–256.
- Van Gelder, T. (1995). What might cognition be, if not computation? *Journal of Philosophy*, 91(7), 345–381.

Copyright of Minds & Machines is the property of Springer Science & Business Media B.V. and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.