

Analog and Digital, Continuous and Discrete

Corey J. Maley
Princeton University

June 5, 2009

Abstract

Representation is central to contemporary theorizing about the mind/brain. But the nature of representation—both in the mind/brain and more generally—is a source of ongoing controversy. One way of categorizing representational types is to distinguish between the analog and the digital: the received view is that analog representations vary smoothly, while digital representations vary in a step-wise manner. I argue that this characterization is inadequate to account for the ways in which representation is used in cognitive science; in its place, I suggest an alternative taxonomy. I will defend and extend David Lewis’s account of analog and digital representation, distinguishing analog from continuous representation, as well as digital from discrete representation. I will argue that the distinctions available in this four-fold account accord with representational features of theoretical interest in cognitive science more usefully than the received analog/digital dichotomy.

1 Introduction

Cognitive science is committed to at least two ideas: that the mind/brain is a computer, and that the mind/brain operates on representations. Von Eckardt (1993) characterizes these as the substantive assumptions of the field. Yet representation is not an unproblematic notion: serious theoretical and empirical questions about representation drive current research in cognitive science. The problem of intentionality has commanded much attention in philosophy,

and in psychology, significant effort has gone into characterizing representations thought to underlie various psychological phenomena.

One fact about representation that has been noted but insufficiently interrogated is that the *form* of a representation determines both the kinds of things that can be represented and the ways in which that representation can be manipulated. The received analog/digital dichotomy is one way of distinguishing such forms: some representations vary smoothly, others vary step-wise. But examples from the psychological and computational literature make it clear that ‘analog’ is not, in fact, synonymous with ‘continuous’, and ‘digital’ is not synonymous with ‘discrete’; thus, the analog and digital labels must be tracking something else.

In this essay, I will argue that we can refine the received dichotomy along lines originally suggested by Lewis (1971): analog representation is distinct from continuous representation, and digital representation is distinct from discrete representation. I will first examine some of the accounts of analog and digital representation found in the philosophical literature. A proposal for a more refined characterization of representational types follows. I will then argue for the adoption of this taxonomy because it both clarifies the essential aspects of representational types that are explanatorily relevant to cognitive scientists, and it provides researchers across the constitutive disciplines of cognitive science a unified framework for discussing representation.

Before we begin, however, I will introduce two terms to ease the discussion. 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. So, for example, I might take an amount of sand as representing some quantity Q . In this case, the representational medium is sand. If I represented the quantity by the total number of grains of sand, then—assuming grains of sand do not come in fractions—the representational format would be discrete, and isomorphic to the whole numbers (i.e. $Q \in \mathbb{N}$). On the other hand, if I represented the quantity by the *weight* of the sand, then, depending on the range of numbers I wanted to represent, the representational format would be either continuous or discrete. For example, if I have several tons of sand for representing a number between 0 and 1, it might be most expedient to consider the representational format to be continuous, ranging over all real numbers between 0 and 1 (i.e. $Q \in \mathbb{R}, 0 \leq Q \leq 1$). On the other hand, if I’m representing a range of numbers with a small amount of sand, it might be better to consider the format to be discrete, broken up into, say, hundredths, with some margin of

error used for rounding (i.e. $Q \in \{x : x = \frac{n}{100}, n \in \mathbb{N}\}$). Thus, the representational format may be discrete or continuous, although the representational medium (grains of sand) is discrete.

2 Previous Accounts

Goodman (1968) was perhaps the first philosopher to analyze the so-called analog/digital distinction. Although he acknowledges possibilities to the contrary, his account is essentially the received view, in which ‘digital’ is synonymous with ‘discrete’, and ‘analog’ is synonymous with ‘continuous’:

Plainly, a digital system has nothing special to do with digits, or an analog system with analogy. . . . Since the misleading traditional terms ‘analog’ and ‘digital’ are unlikely to be discarded, perhaps the best course is to try to dissociate them from analogy and digits and a good deal of loose talk, and distinguish them in terms of density and differentiation—though these are not opposites. (Goodman, 1968, 160)

Here, the basic distinction between a digital and an analog representational scheme is that a digital scheme is differentiated or discrete, while an analog scheme is continuous or dense. Goodman was also concerned with whether analog representation was a matter of the representational medium, the representational format, or both; we will discuss these issues below. In any case, it is clear that Goodman equates ‘analog’ with ‘continuous’.

In response to Goodman, Lewis (1971) offers an alternative account, taking issue with the claim that digital systems have nothing to do with digits, and that analog systems have nothing to do with analogy. First, Lewis claims that differentiated representations can be analog in some circumstances. For example, an analog computer might represent a positive integer by the resistance in ohms along a particular part of a circuit. According to Goodman, this would only count as an analog representation if the resistance was set by a continuously-variable resistor; Lewis contends that using a series of single-ohm resistors, along with a device to bypass the unneeded resistors, would count as an analog representation. The number is still represented by the amount of resistance, but instead of varying continuously, the resistance varies in unit steps. So according to Lewis, whether the representation varies continuously or in discrete is irrelevant to its being analog; what counts is that

the representation does its work via a quantity. Thus, analog representation is not necessarily dense in Goodman’s sense. Lewis’s alternative formulation is that “*analog representation* of numbers is representation of numbers by physical magnitudes that are either primitive or almost primitive,” (Lewis, 1971, p. 163, original italics) where “primitive or almost primitive” refers to terms in an appropriate reconstruction of the language of physics (such as ohms).

Lewis then gives an account of digital representation, whereby “we can define *digital representation* of numbers as representation of numbers by differentiated multidigital magnitudes” (Lewis, 1971, p. 165, original italics). Thus, a digital representation uses more than one digit, where both the tokens of the digits and their values are differentiated. An example is an odometer with six digits: each digit can take one of ten tokens (the numerals ‘0’ through ‘9’), each of which is discrete and differentiated, and the sequence of all six digits is used to represent a number in a systematic way (i.e. the first digit on the right represents the number of units, the second digit represents the number of tens, etc., and the sum of all six of these quantities is taken to be the value of the number represented). This is in contrast to Goodman’s analysis of digital, which only requires differentiation (i.e. discreteness).

Haugeland (1981) weighs in on the analog/digital distinction from a different perspective, although his account is ultimately allied with the received view. Haugeland claims that the primary guide for distinguishing analog and digital representations is the reliability of the physical reading (or writing) procedure accessing (or producing) those representations. Specifically, Haugeland understands a digital device to be one in which the tokens of a set of specified types can be written and read reliably and with absolute certainty; analog devices, in contrast, are those in which “the procedures for the write-read cycle are approximation procedures—that is, ones which can ‘come close’ to perfect success,” (Haugeland, 1981, p. 83). Consequently, a representation is analog if and only if it is continuous (or *seems* continuous to the user of the representation; (Katz, 2008) defends this interpretation of Haugeland’s account), because real-world continuous representations only allow for approximate read-write cycles; *mutatis mutandis* for digital and discrete representation. This immediately discounts Lewis’s single-ohm resistor example, about which Haugeland states “I think it’s clearly digital—just as digital as a stack of silver dollars, even when the croupier ‘counts’ them by height” (Haugeland, 1981, p. 80).

Intuitions clearly differ among all three authors, and adjudication is ob-

viously not an empirical matter. Choosing any particular account of representational types should be a matter of deciding which is most useful for taxonomizing the representations we care about. In particular, given the crucial role that representation plays in cognitive science, and given that ideas about representation in cognitive science draw from both psychological and computational theorizing, we should choose the one that most fruitfully allows us to think about how representations in the natural computational system of the mind/brain relate to representations in the artificial computational systems we have engineered. In the next two sections, I will argue that a modified version of Lewis's account fits this bill.

3 Analog and Continuous

The term 'analog', in the sense I am concerned with, originates in computer engineering. Early computing machines employed both continuous and non-continuous representations, using many varieties of mechanical, electromechanical, and electronic media (Mindell, 2002). As digital computing machines became more prevalent, the label 'analog' came to stand for those machines that used a continuous representation of one sort or another. 'Analog electronics' now typically refers to electronics systems that are analyzed with continuous variables; hence the first part of the received analog/digital distinction.

Another sense of 'analog' representation can be found in experimental psychology. This sense is invoked in debates about whether mental imagery is analog or propositional. Of central concern in these debates is how to best interpret data from human performance on certain tasks involving spatial reasoning. Examples include remembering the location of landmarks on a map, or making similarity judgments on two objects, one of which is presented in a rotated form. To get a sense of what psychologists mean by analog representation, we will look at the latter task—known as mental rotation—in a bit of detail.

The seminal finding in the mental rotation literature is that of Shepard and Metzler (1971). Participants in this experiment were presented with two line drawings of three-dimensional objects. The portrayed objects were either identical, or one was a mirror image of the other. Additionally, one of the depicted objects was rotated relative to the other, where the amount of rotation varied between trials. Upon presentation of a pair of objects, par-




Figure 1: Shepard-Metzler objects. The top two are ‘same’; the bottom two are ‘different’.

ticipants were asked to determine whether the two objects were the same or different, where “same” means that one could rotate one of the objects to exactly match the other (top row in Figure 1), and “different” means that no such rotation is possible (bottom row in Figure 1).

Shepard and Metzler found that the time it takes for a person to make a decision of “same” or “different” is directly proportional to the amount (in degrees) that the two objects are rotated relative to one another: in short, the larger the rotation, the longer the response time. Shepard later argued that these results suggest that analog representation underlies mental rotation abilities:

By an analogical or analog process I mean just this: a process in which the intermediate internal states have a natural one-to-one correspondence to appropriate intermediate states in the external world. Thus, to imagine an object such as a complex molecule rotated into a different orientation is to perform an analog process in that half way through the process, the internal state corresponds

to the external object in an orientation half way between the initial and final orientations. And this correspondence has the very real meaning that, at this half-way point, the person carrying out the process will be especially fast in discriminatively responding to the external presentation of the corresponding external structure in exactly that spatial orientation. The intermediate states of a logical computation do not in general have this property. Thus, a digital computer may calculate the coordinates of a rotated structure by performing a matrix multiplication. But the intermediate states of this row-into-column calculation will at no point correspond to—or place the machine in readiness for—an intermediate orientation of the external object. (Shepard, 1978, 135)

It should be noted that, although Shepard mentions analog *processes* here, the point applies to analog representation as well: “We have taken this dependence of transformation time on the extent of the transformation as supportive of the notion that the mental transformation is carried out over a path that is the internal analog of the corresponding physical transformation of the external object” (Shepard, 1978, 134).

Experimental psychologists have adopted Shepard’s notion of analog representation, and argued that it provides the best explanation of the available data. Recent fMRI research supports the argument for analog representation in another way: the amount of activity in brain regions thought to be involved in mental rotation is directly proportional to the amount of mental rotation. So, the larger the rotation, the greater the response time, and the greater the brain activity in specific, rotation-related regions (Zacks, 2008).

I have mentioned these two uses of ‘analog’—the computer engineer’s and the experimental psychologist’s—not in order to show how they are different, but to suggest a way to understand analog representation that includes both. Note that nothing in Shepard’s discussion implies that, by *analog* he just means *continuous*. Rather, he refers to a relationship between the thing represented and the representation itself (in his case, the degrees of rotation required to manipulate an object, and some property of a mental representation of that object).

The proposal I offer is this: contrary to Goodman and Haugeland, analog representation does not refer to continuity in the representational medium or the representational format. Rather, analog representation is represen-

tation in which the represented quantity covaries with the representational medium, regardless of whether the representational medium is continuous or discrete. This is similar to Lewis’s definition of analog representation, but without the requirement that the medium of representation be “primitive or almost primitive.” As long as the quantity doing the representing increases (or decreases) in a systematic way with what is being represented, then the representation is analog. So, for example, we could represent the numbers of days that have passed since my birthday by a number of marbles in a jar: as the number of days increases, so do the number of marbles. The marbles are then an analog representation, even though “number of marbles” is clearly not an almost-primitive physical magnitude.

We can put this characterization more formally 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 maps onto Q ; and
2. as Q increases (or decreases) by an amount d , P increases (or decreases) as a linear function of $Q + d$ (or $Q - d$).

More succinctly, R is an analog representation of Q if and only if P is a property of R and $P = f(Q)$, where f is linear. In the case of Lewis’s resistor example, as well as more typical examples of analog representation in analog computers, P is resistance or voltage, Q is the quantity to be represented, and f is the identity function. In the case of mental rotation, P is some hypothesized (assumedly neural) property of a mental representation, Q is degrees of rotation, and based on the response time data, f is some linear function. Note that, although we know that the relationship between response time (RT) and degrees of rotation is approximately $RT = \frac{Q}{50} + 1$ (Shepard and Metzler, 1971), we do not know the precise relationship between RT and the hypothesized analog representation. Nevertheless, the relationship between RT and degrees of rotation is taken as evidence of an analog representation (to criticize this inference is to criticize a basic methodological assumption driving experimental psychology). We can thus assume that the relationship between the hypothesized analog representation and RT is also linear.¹

¹For example, it takes approximately one second for a person to make a similarity judgment about two objects that are not rotated with respect to one another at all. Although no mental rotation is necessary, there is a certain amount of “cognitive overhead,” per-

This definition captures other examples of so-called analog representation, except that continuity is not necessary—the representational medium or the representational format can be continuous *or* discrete. Thus, the analog computer is no less analog if the voltages of its components come in small units, and mental rotation is no less analog if it turns out that people can only mentally rotate objects in half-degree steps. Although many examples of analog representation use continuous media or formats, this is not what makes them interesting; rather, it is the relationship that they maintain with what they represent, as I have characterized here.

I have been discussing continuity in this essay without precisely characterizing it; doing so would be both beyond the scope—and unnecessary for the argument—of this essay. Nevertheless, it is worth mentioning a few brief points about continuity and continuous representation. First, it should be clear that a continuous representation is simply one that takes on continuous values. These values may vary across either the representational medium or the representational format. Second, pace Goodman, continuity is different than density. Density is the property of having an element between any other two: the rational numbers are a clear example. However, because there are gaps among them (e.g. π and $\sqrt{2}$), the rational numbers are not continuous.

In this section I have presented a distinction between analog and continuous representation, and argued for its adoption on the grounds that it provides a more useful way to talk about representations. Some aspects of representation concern a particular relationship between representation and thing represented (e.g. mental rotation); other aspects concern whether values can vary continuously or not. My proposal amounts to simply noticing this difference, and making the relevant distinction. In the next section, I will do the same for digital and discrete representation.

4 Digital and Discrete

Similar to ‘analog’ and ‘continuous’, the terms ‘digital’ and ‘discrete’ are synonymous under the received analog/digital dichotomy. In this section, I will argue that one can usefully distinguish between digital and discrete representation. The basic idea is again due to Lewis (1971): a digital representation

happens involved in visual orientation and response initiation. But this overhead is assumed to be constant with respect to the manipulation of the analog representation necessary for performing the task.

is one in which a number is represented via its digits, and a discrete representation is one that is not continuous. The upshot of this refinement is parallel to that of the previous section, in that this distinction is useful for characterizing the representations used among the various areas of cognitive science.

Digital representation is the scheme we normally use to represent numbers. More explicitly, it is the representation of a number consisting of: a series of digits, where *digit* just means a numeral in a specific place; and a *base*, which is used to interpret the relative value of digits. So, for example, to represent the number three hundred forty eight in base 10, we use the three-digit series ‘348’, which has ‘3’ in the hundreds place, ‘4’ in the tens place, and ‘8’ in the ones place—the value is $(3 \times 10^2) + (4 \times 10^1) + (8 \times 10^0)$. More generally, a digital representation ‘ $d_n d_{n-1} \dots d_1 d_0$ ’ in base b is interpreted as $(d_n \times b^n) + (d_{n-1} \times b^{n-1}) + \dots + (d_1 \times b^1) + (d_0 \times b^0)$.

Our familiarity with digital representation can obscure its actual complexity. Notice that ‘3’ cannot be understood as a component of a digital representation except in the context of its place. Without this context, ‘3’ could represent 300 (as in the example above), 3,000, or just 3. Worse, without knowing the base, ‘3’ could stand for virtually any number (e.g. in base four, ‘30’ = 12). Additionally, the base dictates how many numerals can be used for each digit: in base n we have exactly n numerals. With these considerations in place, we can now precisely define digital representation. To wit:

1. a series of digits, each of which is a numeral in a specific place within the series; and
2. a base, which determines the value of each digit as a function of its place, as well as the number of possible numerals that can be used for each digit.

The value of a digital representation is then determined using the formula given above.

Nothing about this characterization is controversial: it is just the way we typically represent numbers. However, I claim that the term ‘digital’ should be reserved only for representations of this type, rather than discrete representation more generally, as the received view would have it. Digital representation is a variety of discrete representation, but the converse is not true. To see this, let us examine discrete representation in more detail.

Whereas a digital representation necessarily represents a number, a discrete representation has no such restrictions. To be precise, an individual representation is not by itself discrete; rather, it is the representational scheme of which it is a part that is discrete. The essential requirement for being a discrete representational scheme (which both Goodman (1968) and Haugeland (1981) capture in their accounts of digital representation) is that representations are distinct from other representations in the same representational scheme, and there are gaps between the possible representations. For example, letters of the alphabet, written words, and poker chips are all discrete, non-digital representations. In these cases, the representation is either wholly present or it is not—half of a letter does not count as a letter. Furthermore, there is no representation between any two representations that are adjacent according to some ordering. There is no letter between ‘a’ and ‘b’, and if we sort English words by length and then alphabetically, there is no word between ‘cam’ and ‘can’.

Although a minor point, it should also be noted that—contra Goodman—discrete representations can be dense. The most obvious example is the representation of positive rational numbers by fractions. Again, this is so familiar that it is worth making explicit. The rational number one half is represented by two digital representations separated by a ‘/’: i.e. ‘1/2’. Thus, a fraction is a discrete representation, consisting of two digital representations and a special symbol. If we try to order fractions by the values they represent, we find that there is always another fraction between any two; this is precisely the property of being dense. However, if we order them as follows: ‘1/1’, ‘1/2’, ‘2/1’, ‘1/3’, ‘2/2’, ‘3/1’, ‘1/4’, ‘2/3’, . . . , then there is no representation between successive pairs. So this is both a discrete and a dense representational scheme. In fact, this property is precisely what separates dense from continuous representational schemes.

The most important reason for separating digital from discrete representation is to make clear the peculiarities of digital computers—particularly if we are to treat them as models of cognition. Digital computers are not digital simply because they are discrete, but because their representations are digital in the sense I have presented here. Space prohibits a detailed description of digital computation, but it must be noted that an essential aspect of such computation is that it uses so-called digital logic. This is the lowest level of abstraction used in designing digital computers, where components operate on strings of individual bits (at levels lower than digital logic, one is dealing with hardware and electrical engineering). These strings of bits are inter-

preted as binary representations of numbers, and operations on these strings rely on the fact that they are composed of digits in the sense articulated above. For example, an adder might determine the sum of two numbers by adding them digit-by-digit, using the same algorithm children are taught.² In the case of machine instructions, certain digits will encode a particular operation, while other digits in the instruction encode addresses or registers in memory.

Two simplified examples will help illustrate this point. In digital computers, characters in ASCII format are represented as numbers, which are in turn represented by seven binary digits. Also, all uppercase letters are 32 less than their lowercase counterparts. The letter ‘K’, for example, is represented by the number 75, which is ‘1001011’ in binary. The lowercase ‘k’ is represented by 107, or ‘1101011’. So, a digital component that converts uppercase letters to lowercase letters could simply add 32 to the uppercase letter. This would work by changing the second digit from the left from a ‘0’ to a ‘1’, because 32 is ‘0100000’ in binary. In other words, this component would turn the digit in the “32s place” to a ‘1’.

Another example is accessing memory locations. A (very) simple digital computer could have four separate devices (call them chips) for storing bytes of data, each of which has eight locations for a particular byte of memory, for a total of 32 units of memory. When the computer is given an instruction to read a particular unit—an address between 0 and 31—a natural way to access the memory is to use the most significant places of the address to determine which chip should be accessed, and then use the remaining digits to fetch the particular byte in that chip. In other words, each chip has locations 0 through 7, and the chips are numbered 0 through 3. To access location 6 in chip 2, we would use address ‘10110’: the first two digits (on the left) indicate that the address is stored in chip 2 (10 in binary); the last three indicate that it is location 6 that we want. The programmer, of course, does not need to know these details: as far as she is concerned, there are simply 32 bytes of memory. But those wishing to understand the computer at a lower level must attend to the digital nature of the representations involved. Now, modern computers do not have specialized components for turning uppercase letters into lowercase ones, and memory is usually measured in the gigabytes. Nevertheless, these toy examples illustrates the point: digital computers use

²Strictly speaking, the computer operates on binary representations of numbers, and not numbers themselves. For the sake of brevity, I have left this qualification out.

digital representations *qua digital*: their representations and operations take advantage of—and can only be understood using—the special structure of numbers represented via digits.

5 Discussion

I have argued that the received analog/digital distinction comes apart in two ways: we can distinguish between analog and continuous representation, and we can distinguish between digital and discrete representation. One obvious objection is that these are just distinctions without a useful difference: after all, we have gotten along quite well using the received view, and it may be futile to make the distinctions proposed here (recall the quotation from Goodman (1968)). Or, one could argue that the synonymies assumed within the received view *really are* synonymies, and that by using ‘digital’, one may as well have used ‘discrete’ (as Haugeland (1981) assumes). In response, it seems best to point to instances where the received view is misleading or discounts useful representational types that one can make using the four-fold view I propose. I have already mentioned how theorizing about mental rotation benefits from one half of the story (the distinction between analog and continuous representation); in the following section, I will discuss some further examples that can benefit from the proposed analysis.

Before proceeding to particulars, however, it is worth considering a more general point about the potentially misleading use of ‘analog’ as synonymous with ‘continuous’. If some phenomenon of interest is continuous, describing it as analog presupposes that the phenomenon is a representation, which may not be the case. It is an empirical question whether all of the continuous phenomena in neural systems are actually representations, even if they are continuous at some level of theoretical description. Similarly, the ‘digital’-‘discrete’ synonymy also presupposes representation, when all that may be of interest is discreteness. Again, in neural systems, it may be quite useful to characterize action potentials or ion channels as discrete phenomena, without thereby declaring them representations (as one would by calling them digital). Saddled with the received view, however, we cannot usefully make these distinctions: the terms ‘analog’ and ‘digital’, coming as they do from computer science, imply that representations are at hand. They may not be.

The best example of what I claim we should avoid comes from Eliasmith (2000), who proposed an answer to whether the brain is analog or digital. Us-

ing the received view, he argues that: the brain uses ‘digital’ representations; the brain processes information; information used by the brain cannot be infinite; ‘analog’ representations require infinite information, and thus cannot be used by the brain; therefore the brain is digital. If this is true, are the researchers studying mental rotation fundamentally misguided? And what kind of number system is the brain using? Is it binary, like contemporary digital computers, or does it use a decimal system, like some of the first computers? Further, might the neuroscientists gain some crucial insight into the functioning of a digital system like the brain by studying how other digital computer systems operate?

The problem here is obvious. Being reasonably charitable, Eliasmith probably does not mean to claim that the brain is like a digital computer (and I am not sure what he would think about mental rotation). So why use the misleading terms? By insisting on the ambiguity inherent in the received view, one must defend Eliasmith by insisting that both ‘analog’ and ‘digital’ have different senses: sometimes they mean one thing, sometimes another. I claim that it is much simpler to use ‘discrete’ and ‘continuous’ in cases like Eliasmith’s, and use ‘digital’ and ‘analog’ in cases that merit their more refined definitions given above.

Another example coming from the cognitive science literature discusses representations of quantities in rats, pigeons, and monkeys (Gallistel et al., 2006). Gallistel and his colleagues review data suggesting that, when remembering some interval of time, the variability of the error in the time remembered increases as the interval increases:

The fact that the memory for numerosity exhibits scalar variability suggests that numerosity is represented in the brains of nonverbal vertebrates like rats, pigeons, and monkeys by mental magnitudes, that is, by real numbers, rather than by discrete symbols like words or bit patterns. When a device such as an analog computer represents numerosities by different voltage levels, noise in the voltages leads to confusion 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 the kind of variability seen [earlier]. For example, the bit-pattern symbol for 15 is 01111 whereas for 16 it is 10000. Although the numbers are adjacent, the discrete binary symbols for them differ in all

five bits. Jitter in the bits (uncertainty about whether a given bit was 0 or 1) would make 14 (01110), 13 (01101), 11 (01011), and 7 (00111), all equally and maximally likely to be confused with 15 because the confusion arises in each case from the misreading of one bit. . . . Thus, the nature of the variability in a remembered target number implies that what is being remembered is a magnitude—a real number. (Gallistel et al., 2006, 252)

Gallistel, like Eliasmith, assumes that the received digital/analog dichotomy provides the only option for distinguishing representational types. One possibility consistent with these results is that numerosity is a discrete analog representation. An organism could quite plausibly represent numerical quantity by a discrete number of neural firings; noise in the number—some neurons fire too much, others too little—would slightly increase or decrease the quantity represented, just as noise introduced into the jar of marbles representing days since my birthday would only increase (or decrease) the magnitude as a function of the intensity of the noise. Obviously, how numerosity is actually represented is an empirical matter; but we should not limit the possibilities to continuously-varying quantities, as Gallistel suggests.

The last two examples I will mention comes from neuroscience. Recent work by a number of researchers (e.g. Shu et al. (2006) and Alle and Geiger (2006)) has cast doubt on the conventional view of neural firing as an “all-or-none” phenomenon. Instead, the post-synaptic potential invoked by some neurons may vary as a function of the input received. Neuroscientists describe this finding as alternately *graded*, *analog*, or *continuous*. The variability is clearly continuous; what’s not so clear is what this phenomenon has to do with representation. The relevant neural recordings are all done *in vitro*: they are not responses to stimuli given to an organism, but responses to current induced into cultured neurons by microelectrodes. As one might expect, I claim we should insist upon calling such phenomena continuous. The point may seem small, but it is critical that we not import representational capacities into the activities performed by neural systems until we know that those activities do, in fact, underlie or implement representations. Much of what the brain does is likely to be irrelevant to mentality; it is better to be clear about when such claims are being made, and when they are not.

Finally, Romo et al. (2000) discuss the representation of a vibrating tactile stimulus applied to the fingers of monkeys. Recordings of a particular neuron show that it fires at a rate varying linearly with the frequency of the

stimulus. The individual firings of this neuron are discrete, and the number of firings in a given time represents the stimulus frequency. This is a clear case of analog representation: whether one thinks of the neural firings as a discrete number, or the frequency of neural firing as a continuous variable, it is clear that the representation covaries with what is being represented in a manner consistent with analog representation as I have proposed. And importantly, we are justified in calling this a representation, as opposed to the preceding example: the quantity of neural firing is representing the quantity of a stimulus in the world, presented to an organism.

6 Conclusion

I have presented an analysis of representational types that complicates the received analog/digital distinction into a four-fold framework. The examples presented throughout this essay demonstrate the benefit of my proposal. Individual researchers or research programs in the cognitive sciences may not be confused about what they mean by the terms they use to label phenomena of interest, but a standard framework for analyzing representational types across research areas is obviously beneficial. But the benefit of the analysis goes beyond a mere legislation of terms. It is clear that the essential feature of analog representation is the relationship between the thing doing the representing and the thing represented, not just whether continuity is to be found. Similarly, digital representation is a matter of representing numbers via digits, and not just an absence of continuity. It should be noted that the analysis proposed here extends to representations outside of cognitive science: digital and analog music recordings and broadcast signals fall in place as they should. Less obvious examples are covered nicely as well. Consider dimmer switches, where turning a knob varies the intensity of a light. This is analog in the received view, because it seems continuous. However, these devices work by rapidly turning a constant voltage on and off (fast enough that humans do not perceive any flicker). Maybe these would be digital under the received view, because there are only two voltages. In any case, my proposal gets the example right: the dimmer switch is analog, because what matters is the relationship between the light output and the amount that the knob is turned.

The nature and uses of representation are serious objects of investigation for many philosophers and cognitive scientists. I have argued in this essay

for a new way of characterizing representations, which reflects important dimensions along which they vary. I have also illustrated why it matters to make the distinctions I have suggested. Given that we do not know how representation works, we would be wise to avail ourselves of as many different types of representation as is appropriate. The framework I have argued for in this essay provides a clear improvement upon the traditional analog/digital distinction.

References

- Alle, H. and Geiger, J. R. P. (2006). Combined analog and action potential coding in hippocampal mossy fibers. *Science*, 311:1290–1293.
- Eliasmith, C. (2000). Is the brain analog or digital? the solution and its consequences for cognitive science. *Cognitive Science Quarterly*, 1(2):147–170.
- Gallistel, C. R., Gelman, R., and Cordes, S. (2006). The cultural and evolutionary history of the real numbers. In Levinson, S. C. and Jaisson, P., editors, *Evolution and Culture*. MIT Press, Cambridge, MA.
- Goodman, N. (1968). *Languages of art*. Hackett Publishing, Indianapolis, IN.
- Haugeland, J. (1981). Analog and analog. *Philosophical Topics*, 12:213–222.
- Katz, M. (2008). Analog and digital representation. *Minds and Machines*, 18:403–408.
- Lewis, D. (1971). Analog and digital. *Nous*, 5:321–327.
- Mindell, D. A. (2002). *Between human and machine*. Johns Hopkins University Press, Baltimore, MD.
- Romo, R., Hernández, A., Zainos, A., Brody, C., and Salinas, E. (2000). Exploring the cortical evidence of a sensory-discrimination process. *Philosophical Transactions of the Royal Society of London B*, 357(1424):1039–1051.

- Shepard, R. N. (1978). The mental image. *American Psychologist*, pages 125–137.
- Shepard, R. N. and Metzler, J. (1971). Mental rotation of three-dimensional objects. *Science*, 171(3972):701–703.
- Shu, Y., Hasenstaub, A., Duque, A., Yu, Y., and McCormick, D. A. (2006). Modulation of intracortical synaptic potentials by presynaptic somatic membrane potential. *Nature*, 441:761–765.
- Von Eckardt, B. (1993). *What is cognitive science?* MIT Press, Cambridge, MA.
- Zacks, J. M. (2008). Neuroimaging students of mental rotation: a meta-analysis and review. *Journal of Cognitive Neuroscience*, 20(1):1–19.