Structural Representation is Analog Representation

(to appear in *Philosophy and the Mind Sciences*)

Corey J. Maley
Purdue University
cjmaley@purdue.edu

Abstract

Recent years have seen an increasing amount of attention devoted to the subject of structural representation. Is there one type of structural representation or many? How do they differ from other types of representation? Are they really a genuine type of representation in the first place? All good questions which I will not address directly. Instead, I argue that structural representations are nothing more than analog representations, and as such, some much needed clarity can be gained by understanding them as such. Typical analog representations (e.g., liquid thermometers or analog clocks) are often "one-dimensional;" extending how they are analyzed into multiple dimensions elucidates the structure of more complex analog representations, such as photographs, maps, or three-dimensional models. However, this analysis applies to structural representations without remainder. The upshot is that we can directly apply what we have learned about analog representation to our understanding of structural representation, which, if not directly answering these recent questions, greatly adds to our theoretical resources for doing so. The analog wheel has already been invented; characterizing structural representation does not need to reinvent that wheel.

1 Introduction

We know that representations come in many types, but finding the right theoretical taxonomy is an ongoing project. This project is complicated by many factors: theorists from different disciplines come to the table with different backgrounds, training, and auxiliary theoretical commitments. For example, some neuroscientists and other cognitive scientists may couch their understanding of representation in terms of Marr's three levels, while others who explicitly reject the utility of that framework understand representation differently. Cognitive scientists who endorse computationalism of some form or another may understand representation one way, while those committed to varieties of enactive, embodied, or embedded cognitive science may understand representation differently (or reject the necessity of representations in cognitive science altogether).

Having a clear taxonomy of representational types is of obvious importance. Similarly to taxonomic hierarchies in biology, seeing how particular types of representation are divided and subdivided would allow us to see how they are similar and different and what follows from these similarities and differences.

How might we best taxonomize representational types? As a first pass, we might start by considering divisions between already-established dichotomous groups: perhaps the analog/digital distinction, or the distinction between structural and non-structural representations. One might focus on semantic features of representations, like the distinction between those with original versus derived intentionality; or perhaps one might focus on artificial versus natural representations. Or maybe we should begin with more than two divisions: perhaps there are three primary types of representation that can be further subdivided. In any case, the first step in developing any taxonomy is to become clear about what kinds of representation there are in the first place. In biology, we had to discover that there are frogs, salamanders, snakes, and caecilians before we could conclude that all three are similar enough to one another that they should fall into the clade "amphibian," and that frogs and salamanders are more alike to one another than either are to caecilians, but snakes (which look superficially like caecilians) do not belong to the amphibians at all.

Structural representation has received much attention in the last decade and may be a contender for a unique and important representational type Shagrir (2012); Shea (2014); Gładziejewski and Miłkowski (2017). Recent work has challenged this contention. Facchin (2024), for example, has argued that there are at least four distinct types of representation that have been labeled as "structural," and as such very little can be said about structural representation in general: one needs to specify *which* species of structural representation is in question. Moreover, Artiga (2023) argues that no plausible characterization of structural representation can support claims about the supposedly unique utility of appealing to this type of representation in representation-using systems.[1]

The thesis for which I argue here is that structural representation and analog representation are the same type of representation, and that structural/analog representation is, in fact, its own unique and robust representational type. Although I find this interesting in its own right, this identification provides some clarity to the nature of this type of representation. Analog representation has been largely (but not exclusively) developed in the context of artificial analog computational and representational systems, while structural representation has been largely (but not exclusively) developed in the context of neuroscience and cognitive science. If it is true that they are the same thing (the contention of this very essay), then what we know about analog representation applies to structural representation, and vice versa. This is particularly useful when it comes to applying principles of analog representation to structural representation: despite some minor disagreement about the proper way to characterize analog representation, there is more unity in contemporary theorizing about analog representation than there is about structural representation. Moreover, noting this identification provides some evidence that we have identified a robust type of representation: one that spans natural and artificial systems.

The first part of what follows aims to show how analog representations and structural representations are the same. Many prototypical examples of analog representations are one-dimensional (e.g., liquid thermometers and hourglasses); however, a careful characterization of this type of analog representation can be extended to multiple dimensions to account for things like photographs, maps, and three-dimensional models (Maley, 2023a). This multi-dimensional account can then be applied seamlessly—without remainder—to structural representations. In the second part, I argue that this identification can be useful for addressing some questions about structural representation that have been noted in the literature, simply by appealing to what we know about analog representation. This is simply because, in many cases, analog representation is understood more clearly.

Let me make two preliminary notes before we get started. First: All too often in contemporary analytic philosophy, one can find symbolization and jargon where

[1] A different take on a similar point is made in Facchin (2021), although in later work (((Facchin, 2024)) states that he had conflated different types of structural representation and disavows these earlier claims.

they are completely unnecessary, where prose alone would do the best job of conveying the ideas being discussed. Unfortunately, we are not in that situation here. In order to correct some of the mistakes and clarify some of the concepts at issue, we have to use some mathematical terminology and notation.

Second: Some accounts of structural representation (or other types of representation) try to accomplish two goals: characterizing both the "structural" aspect of structural representation on one hand and the "representational" aspect on the other. In other words, it is sometimes implicit that an account of representation must explain both how and why that type of representation works, as well as what makes it representation in the first place. Lofty goals, to be sure, and goals that I would ultimately like to accomplish myself. However, in what follows, I set aside the latter concern about what makes a representation a representation, and focus instead only on what makes a given type of representation the type that it is. In a slogan: you give me a representation, and I will tell you whether it is structural/analog and why, but I am not addressing the question of what makes something a representation in the first place.

2 The Structure of Analog (and Structural) Representation

It might seem obvious that structural and analog representations cannot be the same, because paradigmatic examples of each typically differ. According to Shea (2014), for example, maps are a clear example of structural representations, while analog representations like liquid thermometers are not. This chapter will spell out why they are, in fact, the same type of representation, for principled reasons. We begin with a characterization of cases such as liquid thermometers, hourglasses, and analog watches. From there, this characterization can be extended (in a principled manner) to cases like maps, photographs, 3-D models, and all other examples of both analog *and* structural representations.

2.1 Mirroring & Structure

The essential feature of structural representation is the presence of some kind of structure-preserving relation between representations and what is represented. Several authors have offered characterizations along these lines, such as

- Structural representation: the pattern of relations among the constituents of the represented phenomenon is mirrored by the pattern of relations among the constituents of the representation itself. (Swoyer, 1991, p.452).

- Structural Representation: A collection of representations in which a relation between representational vehicles represents a relation between the entities they represent. (Shea, 2014, p.123).
- Structural Representation: A mechanism in which a structural correspondence between a set of vehicles and a target domain is used to accomplish [a] certain task. As a first approximation, structural correspondence refers to some kind of structural similarity between a set of entities and a target domain. (Artiga, 2023, p.1).
- In the case of structural representations, a vehicle (or a system of vehicles) must be structurally similar to the target (or target system)… Simplifying to the extreme by "structural similarity" I have in mind a relation of one-to-one correspondence between (at least) some of the elements of two distinct domains, such that a same abstract pattern of relations holds between the elements of the two domains involved in the correspondence. In the case at hand, the two domains are representational vehicles on the one hand and targets on the other. (Facchin, 2024, p.3–4).

These (and other) characterizations all gesture in the same direction, but we can be more precise and clarify a few points. But let us do this by switching our attention to analog representation first; as we develop a detailed account of analog representation, we will see that, along the way, the account is also just an account of structural representation. To begin, it will be helpful to consider a few familiar, concrete examples of analog representation just to get a feel for what is unique and interesting about it in the first place.

Consider analog thermometers, which indicate temperature via the height of a column of liquid, such as mercury or colored alcohol. The higher the height of the liquid, the higher the temperature; the lower the height, the lower the temperature. Or consider hourglasses (also known as sand timers). These devices indicate the time elapsed after they have been turned over. The greater the amount of sand in the bottom, the greater the elapsed time; the lower the amount, the lower the time.

An example of a simple analog representation in neural systems is rate coding, where a neuron's firing rate represents, for example, the intensity of a stimulus, a “preferred” stimulus orientation, or the angle of an organism’s joint Maley (2018). In some cases, the representation is straightforward: the higher the intensity of the stimulus, the faster the neuron fires. In other cases, it is slightly more complicated. Consider orientation-specific cells, which fire fastest when a visual stimulus is presented at a particular angle. When the stimulus deviates from that angle — when the stimulus is rotated — the neuron’s firing rate decreases. Plotting these responses results in a well-known U-shaped curve, with the preferred orientation at the apex. Thus, changing the angle by, say, 10°, *in either direction*, results in a decrease in the neuron’s firing rate.

It is common to think of analog representations as involving continuous quantities, and they often do. However, that is not necessary (Lewis, 1971; Maley, 2011, 2023a). The simplest example of a non-continuous (i.e., discrete) analog representation is the second hand of an analog watch that ticks in discrete steps, perhaps once every second. More interesting discrete analog representations were used in twentieth-century analog computers, where discontinuous mathematical functions were represented by discontinuous circuit elements (Maley, 2023a).

What makes these examples analog (and what makes structural representations analog, as we will see below) is that they represent what they do via some physical magnitude that varies monotonically[2] with what they represent. Analog thermometers represent temperature via the *height* of a column of liquid, such that the height of the liquid and the temperature vary monotonically—an increase in temperature results in a literal increase in the height of the liquid, where that height is what is doing the representing. Similarly for hourglasses and the hands of an analog clock: as elapsed time increases, the amount of sand or the angle of the hand literally increases. In the literature on analog representation, this is called the “mirroring” conception of analog representation: the representation “mirrors” what it represents (Beck, 2018).[3]

The mirroring relation — and also the structural similarity mentioned in discussions of structural representation — can be made precise by characterizing this relation in terms of a homomorphism. Briefly, a homomorphism is a structure-preserv-

[2]Often this variation is linear, as in all of these examples. But there are exceptions. What is important is that, more generally, increases/decreases in the representation correspond to increases/decreases in what is represented — although the converse need not be true. More details will be discussed below.

[3]Note that I am not here making any claims about how this representation is accomplished. All that matters is that the relevant relationship is in place.

ing map between two sets of elements, where "structure-preserving" can be given a precise definition. This is also a characterization of analog representation: we will use an adapted form of the one given in (Maley, 2023b). Given a representation R of Q, R is an analog representation of Q (with resolution r) if the following hold:

- There is some property P of R such that the amount/quantity/magnitude of P specifies the amount/quantity/magnitude of Q.
- The amount/quantity/magnitude of P is a monotonic function f of the amount/quantity/magnitude of Q, and f is a homomorphism from Q to P.
- An increase/decrease in P reflects an increase/decrease in Q, but an increase/decrease of Q is only reflected by an increase/decrease in P if Q has increased/decreased by more than r.

Let's unpack this, but in English. Consider the analog thermometer. At any given time, the height of the column of liquid represents a particular temperature. A thermometer only capable of representing a single temperature — say, 23 °C — would be rather useless: instead, the thermometer represents a range of temperatures. For specificity, let us say that the thermometer can represent temperatures between 0 °C and 100 °C using a column of liquid that can vary between 0 cm and 100 cm. We can naturally think of these as two sets: a set T of temperatures (which would be Q in the above formulation) and a set H of heights (which would be P in the above formulation). We will suppose that temperature varies continuously, but (for whatever reason) the thermometer can only increase its height in unit steps (i.e., r is 1 °C).

The second two clauses ensure that the property doing the representing and whatever it represents co-vary in the right way. The requirement of a homomorphism means that there is an ordering among the elements of P and the elements of Q, and the ordering of Q is preserved by the mapping from Q to P. In the thermometer, some heights are taller/shorter than others (more generally, they differ in magnitude); similarly, some temperatures are higher/lower than others. Higher temperatures are mapped to taller heights, and lower temperatures are mapped to lower heights. Specifying a monotonic function ensures that, as what is represented increases, the representing property does not decrease; it either stays the same, or increases. In many cases of analog representation, this function is linear. But for some cases, particularly when we have a discrete analog representation, an increase in what is represented may not reflect an increase in the representing property. Time may pass continuously, but a clock that ticks only every second does not change continuously.

Two points are worth highlighting here. The thermometer example we have been using, where every temperature maps to a unique height and vice versa, would constitute an isomorphism: a special case of homomorphism where the mapping is bijective (i.e., one-to-one and onto). Some accounts of analog representation rely on isomorphisms (e.g., (Lee et al., 2022)), which seems natural in some cases. However, as others have noted (such as O'Brien and Opie (2004) and Shagrir (2012)), isomorphism is too strong. The reason is that isomorphisms require that the set of representations and the set of things represented have the same number of elements. In our thermometer example, because the height only varies in unit steps from 0 cm and 100 cm, there are 101 elements in H, but (uncountably) infinitely many in T.[4] Thus there cannot be an isomorphism between the two. Moreover, it would be a mistake to take the *temperatures* that are being represented as themselves varying in unit steps. Whether they vary continuously or in minute increments is a question for physicists to answer; in either case, there are many more temperatures than there are heights to represent them, and so an isomorphism cannot exist. Now, we might map subsets of T to single elements of H, or a particular subset of T to single elements of H; or, as O'Brien and Opie (2004) note, some cases might be remedied by resorting to only partial isomorphisms. However, the simplest solution is to allow a homomorphism: a many-to-one mapping from T to H, rather than the stricter one-to-one mapping (this is explored in more detail in (Maley, 2023b)).

The second highlight-worthy point is this. Characterizing analog representation in terms of homomorphisms (rather than isomorphisms) allows for considerations of resolution. Sometimes there are increases (or decreases) in what is represented that are *not* reflected in the representation. Temperature may increase (or decrease) by a fraction of a degree without any change in the height of the liquid. However, the converse is not true: an increase (or decrease) in the height of the liquid necessarily reflects an increase (or decrease) in the temperature being represented. Or consider an analog clock: a small amount of time (the representatum) might pass *without* the second hand of the clock moving at all, but when the hand *does* move, then some amount of time has necessarily passed. The usual term for the smallest representable difference is *resolution*. Thus, homomorphisms provide a simple way to account for the resolution of different analog representations.

[4] Or, at the very least, many more than 101.

Let us examine this point more closely. Suppose that time is continuous or very nearly so; we can use stopwatches of different resolutions to represent intervals of time at different levels of precision, i.e., different resolutions. But if we limit ourselves to *isomorphisms*, we cannot make sense of this, because every point in time would have to map to one, and only one, point in the stopwatch's divisions. This is just strange: the fact that time is continuous (or nearly so) yet we measure it with discrete intervals is what allows us to make sense of the fact that we *can* represent time with different levels of resolution. When we use stopwatches of different resolutions, we are measuring the *same* thing with different resolutions, not different things (different time structures?) with different resolutions. Perhaps in a truly continuous representation, there might be an isomorphism between the representatum and the representation; that's fine: after all, isomorphisms are special cases of homomorphisms. But for many other cases, we need to avail ourselves of homomorphisms, and not just isomorphisms.

2.2 Relations and Magnitudes

Perhaps the most useful feature of analog — and structural — representations is the fact that they take advantage of particular features shared between representations and what they represent. But we have to be precise about what this means.

To motivate the issue, notice that there are a huge number of mappings between the sets of temperatures and heights: any particular temperature could be mapped to any particular height (Gładziejewski and Miłkowski (2017) make this point using a slightly different example). Thus, one might map 0 °C to 12 cm, 1 °C to 44 cm, 2 °C to 5 cm, etc. Perhaps this would be a representation, but not an interesting one, and certainly not an analog one. Further, as it stands, this is not a morphism either: we haven't specified a structure among the temperatures that is preserved when they are mapped to heights. But this is an easy fix. All that is required for a structure-preserving[5] homomorphism is that there is a relation between elements of one set that is preserved in the other. And relations, as it turns out, are easy to come by: we can just create one! So let us define a relation, $\preccurlyeq$, between temperatures such that 0 °C $\preccurlyeq$ 1 °C, 1 °C $\preccurlyeq$ 2 °C, etc. Next, define a *different* relation, $\curlyeqprec$, between heights such that 12 cm $\curlyeqprec$ 44 cm, 44 cm $\curlyeqprec$ 5 cm, etc. Now we have the ingredients for a homomorphism. The structure induced by $\preccurlyeq$ among temperatures is preserved in the

[5]More precisely, this should be called a relation-preserving homomorphism, but I will stick with the commonly-used phrase "structure preserving."

structure induced by $\preccurlyeq$ among heights using the mapping (call it f) illustrated in Figure 1. Given $t_1, t_2 \in T$ and $h_1, h_2 \in H$, where $f(t_1) = h_1$ and $f(t_2) = h_2$, if $t_1 \curlyeqprec t_2$ then $f(t_1) \preccurlyeq (t_2)$. As a concrete example, $0\,°C \curlyeqprec 1\,°C$, which is preserved when f is applied: $f(0\,°C) \curlyeqprec f(1\,°C)$. Figure 1 illustrates the homomorphism.

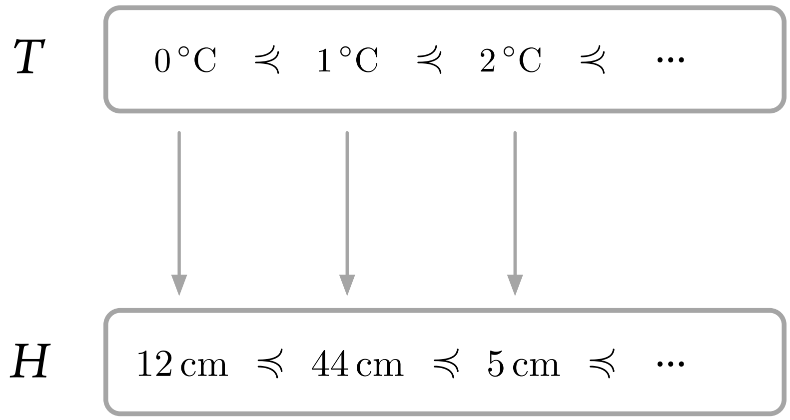

Figure 1: A structure preserving morphism, but not a structural (or analog) representation.

We thus have a structure-preserving map between temperatures and heights: each height corresponds to a particular temperature, and the relation among temperatures is preserved by the relation among heights. However, adding these structures and this structure-preserving map does not make this an analog (nor a structural) representation. What is required to make this an analog representation is rather obvious: instead of using some arbitrary relation among the elements of each set (the $\preccurlyeq$ and $\curlyeqprec$), we need to use a natural relation.

Now, "natural" is a rather vexed term, but what I mean here is straightforward. I simply mean a relation that preserves the ordering structure among *magnitudes*. Both temperatures and heights are magnitudes—they admit of greater-than and less-than relations that form a total order. The mapping between them that is required is one that preserves *this* order structure via a homomorphism. The less-than relation among temperatures maps to the less-than relation among heights. This is in contrast to the arbitrary mapping shown in Figure 1, where temperatures and heights are related by *stipulated* relations ($\preccurlyeq$ and $\curlyeqprec$) that do not track the inherent magnitude structure of either domain. Although such a mapping technically forms a homomorphism, it fails to constitute an analog or structural representation, precisely because it does not preserve the structure of the magnitudes.

Instead of the arbitrary relation in Figure 1, we can use the less-than relation among the magnitudes in each domain; using a mapping that preseres *that* structure gives us an analog — and a structural — representation. Notably, whether the magnitudes are real numbers, rational numbers, or integers, they form an ordered field with a strict total order: a mathematical structure where, for any two elements a and b, either $a < b$, $b < a$, or $a = b$, and that structure inherent in the magnitudes can be preserved by a mapping.

More explicitly, using the less-than relation among magnitudes, we obtain the homomorphism that makes an analog thermometer *analog* (as well as a structural representation). Call the function that maps temperatures to heights m; thus, $m : T \rightarrow H$. The elements of T have a structure: the elements are ordered via $<$; similarly for the elements of H. The mapping m preserves that structure: if $a, b \in T$ where $a < b$, then $m(a) < m(b)$. As an example, suppose 2 cm represents 12 °C, and 4 cm represents 14 °C. The magnitude of 12 °C is less than the magnitude of 14 °C (i.e., $12 < 14$), and the magnitude 2 cm that *represents* 12 °C is less than the magnitude 4 cm that *represents* 14 °C (again, $2 < 4$). This completely captures the nature of an analog representation.[6]

Notice that this characterization of analog representation is also an account of structural representation, albeit a rather simple type of structural representation. There is a structure inherent in what we want to represent (a relation among the elements to be represented), there is a structure among the elements of the representational scheme, and the mapping from the first to the second guarantees that the structure of the representations reflects the structure of what they represent. In the next section, we will extend this account to more complicated examples of analog representation, such as maps and photographs. Let us briefly look at a couple examples to foreshadow why understanding structural representations as analog representations is on the right track.

Consider topographic maps in neuroscience Kaas (1997). Here, the spatial relationships among neurons preserve the spatial relationships in the visual field or other receptor surfaces. These spatial relations are magnitudes: namely, distances and directions. Similarly, elevation maps[7] preserve the same type of spatial rela-

[6]As a reminder, this is not a *sufficient* condition for analog representation: homomorphisms of this type are still relatively cheap. I am not attempting here to give an account of analog representation, full stop. I am only characterizing what makes a proposed representation analog.

[7]Confusingly, these are also called topographic maps, but unlike the notion in neuroscience, they represent the elevations of an area above sea level.

tions: the map itself preserves the relations in the area represented by the map, and heights at different points (magnitudes) are represented by, say, color of an increasing intensity—again, the magnitude structures are preserved. Photographs are another example: light intensity values at points in the photograph preserve the values of corresponding points in the image being represented.

Even cases that may seem to lack a magnitude structure do, in fact, have such a structure upon closer examination. Consider simple maps that only indicate the presence or absence of certain landmarks, such as the "you are here" diagrams, sometimes found in shopping malls, hospitals, or other complex buildings. The presence or absence of a marker at a particular location on the map represents the presence or absence of the corresponding landmark; each point is, by itself, a binary, present/absent representation. But crucially, the spatial relations among the markers — which constitute the structure that makes this a structural representation — preserve the spatial relations in the area being represented.

Now, these examples of structural representation are more complex than the analog thermometer we were considering above (and we will look at these examples in more detail below). The analog thermometer has a rather simple one-dimensional structure, whereas the structural representation examples just mentioned have multiple dimensions. However, a complete account of analog representation should also be able to characterize analog representations with multiple dimesions. In the next section, we will see how to extend the account of analog representation just articulated to multiple dimensions, which then allows the account to apply equally well to prototypical structural representations.

2.3 Dimensions & Variation

The analog thermometer is analog because, as temperature (what is represented) increases/decreases, the liquid height (what is doing the representing) increases/decreases, where in both cases, the increase or decrease is of a magnitude: temperature in the first case, liquid height in the second. In contrast, a digital thermometer also represents temperature, but not by an increase/decrease in the height of a column of liquid (or any other increase in magnitude), but by a change in displayed digits.[8] A temperature of 14 °C would be represented by some pattern of the digits

[8]Of course, the thermometer might *measure* temperature by some mechanism that increases/decreases as temperature increases/decreases, but here we are concerned with how the temperature is *represented*.

"1" and "4" concatenated, where those digits are displayed via a pattern of pixels, a seven-segment display, or something like this. Increasing the temperature to 15 °C simply results in the display of a different pattern: "1" and "5". This is simply a different pattern, no "greater" or "less" than any other pattern (Maley, 2011). So far, so good.

But what is meant when we say that the temperature varies, that it increases/ decreases? Varies *with respect to what*? Closely attending to the answer reveals a subtle but useful way to extend the above account of analog representation to more complex examples. Here, I will draw on the work of (Maley, 2023b).

In the usual case, when we say that some temperature varies, we mean that it varies with respect to time. In the usual case, a thermometer remains fixed in space, and it represents the temperature at that particular point in space at a particular time; as time goes on, the change in the height of the liquid reflects a change in temperature at that point with respect to time. We have the representation — in this case, the height of the liquid. Let us call the thing represented — in this case, temperature — the *representatum*.[9] Additionally, we have an independent variable, the quantity with respect to which the representatum, and thus the representation, varies. In this case, the independent variable is time. Figure 2 shows a plot of the representation and representatum changing with respect to time.

Other common examples of analog representations work similarly: the voltage in an electronic analog computer (the representation) varies with the quantity being represented (the representatum) with respect to time (the independent variable). One minor technicality is that sometimes the representatum and the independent variable can be the same thing. Consider an hourglass: we are representing the elapsed time by the amount of sand in the bottom of the glass. But with respect to what is time varying? Maybe this is an incoherent question, but if anything, it varies with respect to, well, time. If we create a plot as in Figure 2, we would have a straight line indicating the amount of sand that increases and another straight line indicating the amount of time that has elapsed (which would be the line $y = x$). What would be the independent variable of the plot? Time, of course. Thus, our representatum and independent variable turn out to be the same quantities.

We can take these examples to be one-dimensional representations: there is a single dimension along which the representatum varies. In other words, the represen-

[9]Other authors use terms like "target," or "content." These strike some people as loaded in a way that I would like to avoid, hence the use of representatum.

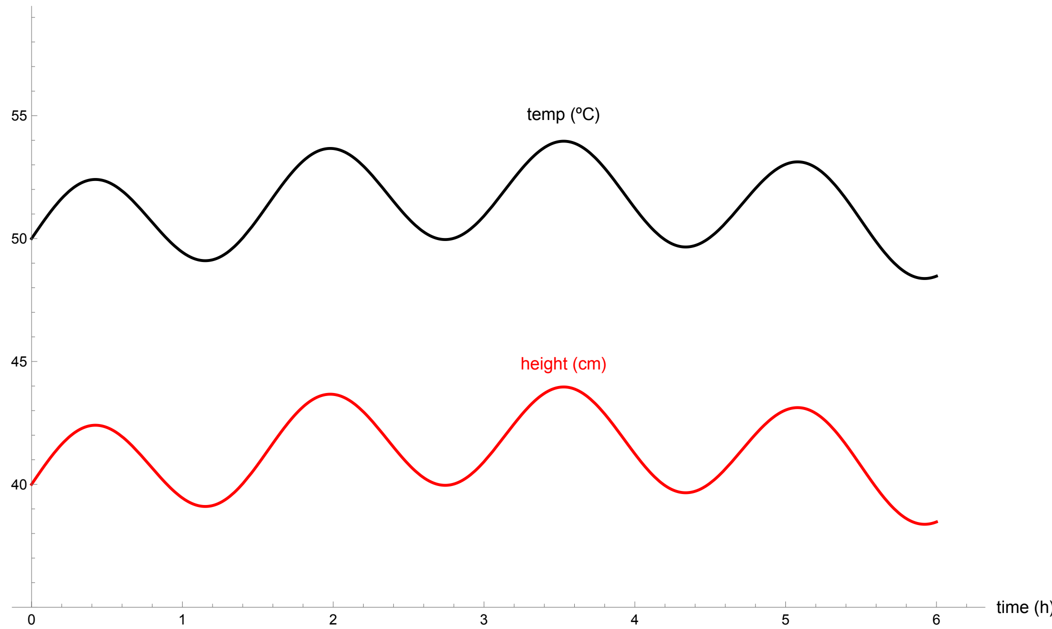

Figure 2: Temperature change and liquid height change as a function of time.

tatum (and thus the representation) is a function of a single independent variable. Things get more interesting when we look at higher dimensions. Consider a photograph (grayscale, for simplicity). In this case, we have a two-dimensional grid of points, where each point represents the gray value of the image being represented. Importantly (and obviously, once made explicit), the structure of the points of the representatum — the image represented — is preserved by the grid of points in the representation. That is: take any point of the image and where it is represented in the photograph. Move to another point left of that point on the image, and the point on the photograph representing that point is also to the left of the first. Similarly for other directions.

In a normal grayscale photograph, the gray value of a point in the photograph represents the gray value (i.e., the amount of light) at the corresponding point in the photographed image. We can zoom in on a single point in the image and, limiting ourselves to just the left-right dimension of the independent variable, see how variation as we move to the left or right corresponds to a variation in the amount of light to the left or right in the represented image. Figure 3 illustrates this point.

A different type of two-dimensional analog representation could represent the

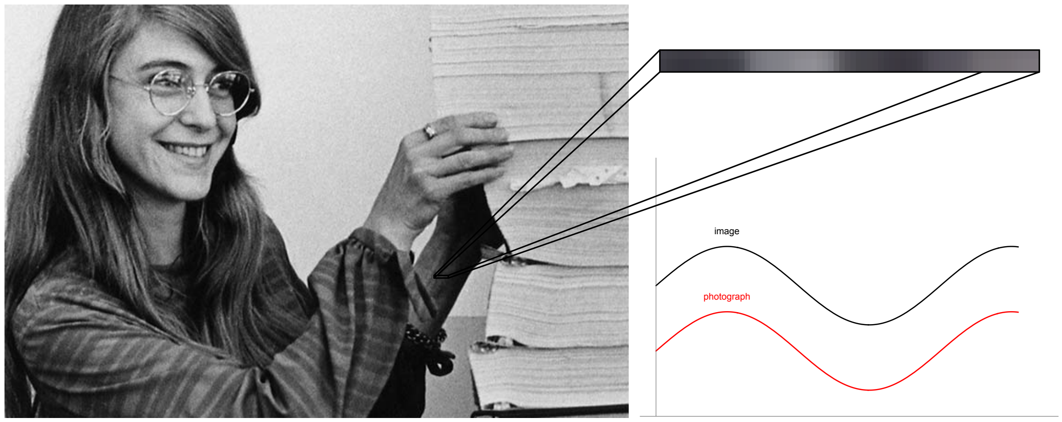

Figure 3: Grayscale photograph depicting Margaret Hamilton. Zooming in shows how variation in the gray level of the photograph corresponds to variation in the amount of light in the image.

same information by a different means. Instead of representing the amount of light in the image by a gray value, we could represent the amount of light by the height of a point. Physical artifacts such as this are not common, but we can illustrate the idea, as shown in Figure 4.

In both of these two-dimensional cases of analog representation, we have a structure in the representatum that is preserved in the representation. This is just like the thermometer case, except that this time the structure is two-dimensional: there are two dimensions of variation. But the structure is preserved all the same. Again, take two points in the image, p_1 and p_2, represented by q_1 and q_2, respectively. If p_1 is above (or to the left of, or below, or above and to the right of) p_2, then q_1 is above (or to the left of, or below, or above and to the right of) q_2. That is one aspect of the structure. There is another: as the value to be represented at each point increases (or decreases) in light level, the corresponding value in the representation increases (or decreases) in gray value (or height, in the example of Figure 4). The structure of the ordering of light values in the image is preserved by the structure of the gray levels of points in the photograph (or the structure of the heights in the graph shown in Figure 4. These two morphisms are illustrated in Figure 5.

In the case of the thermometer, we noted that the height of the liquid increases/decreases as the temperature increases/decreases; the further question we considered was: with respect to *what* are these increases and decreases? The answer in that case was time. In the case of the photograph, the gray value of individual points in-

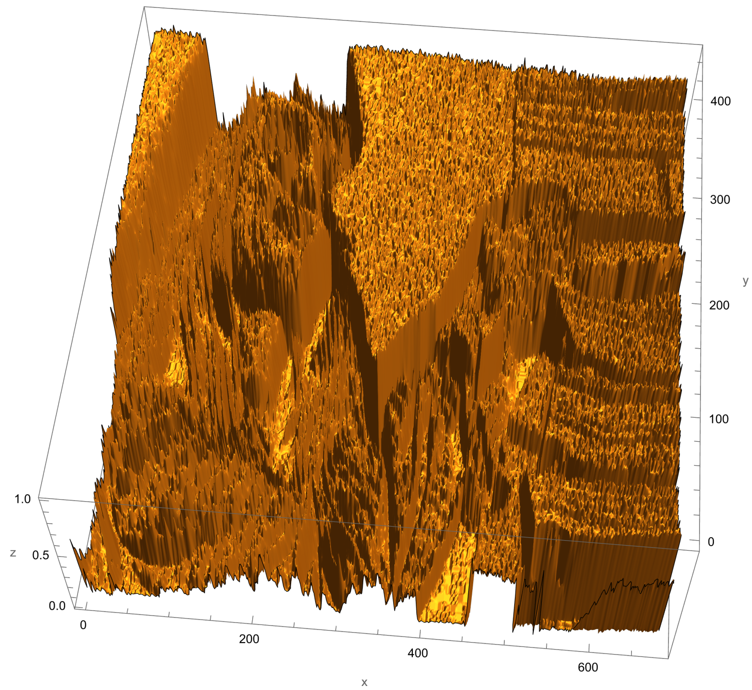

Figure 4: Light values of the image of Margaret Hamilton in Figure 3, but this time represented by height.

creases/decreases as the light at the corresponding image increases/decreases. Yet photographs are static, so it sounds strange to ask "with respect to *what* do these light values vary such that a gray value of a pixel can vary?" Point conceded: photographs are static. However, this is where the introduction of dimensions of variation—the independent variables—becomes useful. This variation occurs with respect to the change in the location of the point in the horizontal and vertical directions. Pick a point in the photograph: it has a gray value that represents the light intensity at the corresponding point in the image that the photograph represents. Looking at other points, the gray level varies with respect to changes in the independent variables.

We can characterize other analog representations with even higher dimensions. A grayscale animation—a series of static images—has three independent variables: the gray level at a given pixel can vary with respect to a change in our two spatial dimensions, but also with respect to time. A three-dimensional model varies with respect to three spatial dimensions; an *animation* of a three-dimensional model (e.g.,

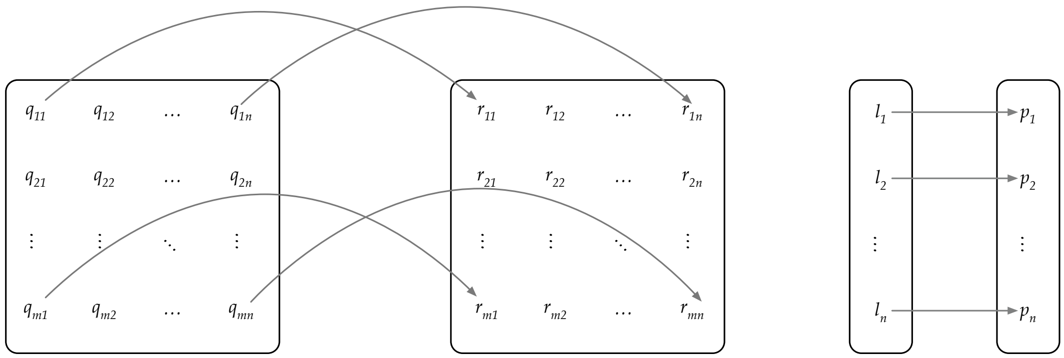

Figure 5: Two structure-preserving morphisms in the photograph. Left: the structure of the locations of light points relative to one another in the image is preserved in the locations of the pixels in the photograph. Right: at each individual point, the magnitude scale of light values is reflected in the magnitude scale of the gray values of the pixel representing that point (or the height of the graph in Figure 4.

a computational model of the heat of the surface on an airplane flying at a particular speed, where temperature is represented by the color at different points on the surface of the model) will have four dimensions of variations: three spatial plus time.

Now, this is actually a little too fast. There is an important subtlety in different types of analog representations that we should take into account. Consider again a simple map as opposed to a photograph; like the earlier example, let us take the map to be a line drawing of the interior of a building, showing only the relative locations of particular points of interest. Like the photograph, the structure between (two-dimensional) points of the representatum and the representation are preserved. But unlike the photograph, *at each point*, there is no preservation of structure: points on a line drawing do not indicate the light values at that point in the building. In fact, it is arguable that, at each point, there is no structure to be preserved in the first place. Rather, the presence of a particular point, or marker, on the map, relative to the structure of the map, corresponds to a point in the represented space relative to the structure of that space, where the two structures are preserved according to an order-preserving homomorphism. "You are here" on the map is simply a point: unlike a given point on a photograph, or heat map, it does not representation some further variable by a change in its value (e.g., gray value, or height on a three-dimensional graph).

Similarly, a simple, monochrome, three-dimensional model of an object represents only the presence or absence of points in space corresponding to the points in space of the represented object. The individual points are either there or not. In a different representation, the points *could* themselves be one-dimensional analog representations: the color of each point could represent the heat of the represented object, or perhaps a particular force acting on each object at a given time.

Using this characterization of analog representation, we could concoct rather exotic analog representations which may or may not be of any use. While it would be interesting to do so, the point of the examples mentioned here is to illustrate two ways in which analog representations can preserve structure. First, the structure of the *independent variable(s)* can be preserved; second, the structure of the variation in magnitude of the representatum can be preserved in the representation.

At this point, it is worth revisiting the ways that different authors characterize *structural* representation, mentioned at the beginning of this chapter. Here are just two: ”Structural representation: the pattern of relations among the constituents of the represented phenomenon is mirrored by the pattern of relations among the constituents of the representation itself,” (Swoyer, 1991, p. 452). ”Structural Representation: A collection of representations in which a relation between representational vehicles represents a relation between the entities they represent,” (Shea, 2014, p.123).

The characterization of multi-dimensional analog representations just presented gives us—for free—an account of structural representation. Multi-dimensional analog representations preserve structure in exactly the way that theorists claim structural representations must. Making this explicit, we have the following:

The Analog/Structural Representation Thesis: Analog representation and structural representation are precisely the same type of representation.

Once again, let us examine some cases of structural representation to see just why this thesis is true. Consider retinotopic maps in area V1 in the primary visual cortex. This area of the cortex exhibits retinotopic organization: adjacent neurons respond to adjacent locations in the visual field. This is a standard example of a structural representation (Shea, 2014). Does it fit the analog framework? Yes: we have two spatial dimensions (the independent variables) along which receptive field locations vary. At each location in V1, neural activity (the representation) varies with light intensity at the corresponding location in the visual field (the representatum). This is precisely analogous to the example of the photograph mentioned above.

Place cells in the hippocampus are another example. A place cell fires when an organism is at some particular location in an environment, with nearby place cells representing nearby locations (O'Keefe and Nadel, 1978). The spatial structure of place cells preserves the spatial structure of the environment the cells represent. Again, this fits the analog framework above: we have spatial dimensions as independent variables, and place cell activity at each location in the neural space represents a location in physical space.

There are other examples, such as the body map (the famous, exaggerated "homunculus") in the somatosensory cortex, and the representation of sound frequency and intensity in the cochlea. What these examples reveal is that the multi-dimensional analog framework developed above captures not just examples of analog representation more complex than thermometers and hourglasses, such as maps and scale models; it is also precisely the framework that captures what researchers mean when they identify neural representations as "structural." The preserved structures are magnitude structures along one or more dimensions of variation.

Put another way, the key point is that structural representations invariably involve using magnitudes and their inherent structure to preserve the magnitude structures inherent in what is being represented; and this is precisely what analog representation does as well. Or, as Peacocke (2019, p. 52) nicely puts it, "Analog representation is representation of magnitudes, by magnitudes." Moreover, a supposed structural representation that *failed* to preserve magnitude relations would either not be a genuine representation at all (perhaps like the arbitrary relation "preserved" by the mapping in Figure 1), or if it is a representation, would be characterized as some other type of representation (perhaps a symbolic representation).

The importance of this kind of magnitude structure preservation is connected to the notion of exploitability Shea (2014, 2018) discusses as central to structural representation. Relation-sensitive mechanisms can exploit structural correspondences precisely because the preserved magnitude structures allow valid inferences. Suppose representations a_1 and a_2 that represent b_1 and b_2, respectively. If some mechanism can determine (or detect, or "read off") that a_1 stands in relation R to a_2, and R preserves some magnitude relation S between b_1 and b_2, then the mechanism can *infer* that b_1 and b_2 stand in relation S. Exploitability here depends on the preservation of magnitude structures—i.e., on the representation being *analog* in precisely the sense characterized here.

Two minor points are worth noting before we move on. First, I mentioned earlier that analog representation need not be continuous; we can see that this holds for all of the analyses and all of the examples given in this section. Photographs can be continuous or pixelated; time can move continuously or in steps; three-dimensional models can be smooth or composed of discrete voxels. In any case, the characterization is the same. Second, the terminology is ultimately not the point: it does not matter much whether we refer to the type of representation I have been discussing as “analog” or “structural,” or something else entirely. What *is* important is realizing that all of the examples of analog and structural representation form a coherent, theorizable class of representational types. While examples of analog representations typically come from artifacts and examples of structural representations typically come from neuroscience, understanding that they are one and the same type clarifies how all of them work.

3 Analog Resources for Structural Representation Questions

The fact that structural representations can be understood as analog representations without remainder is interesting in its own right. But this realization also helps answer — or at least shed light on — some recent questions about structural representations. So, let us examine some of these.

3.1 Dispensing with Sufficiency Concerns

The first issue was mentioned at the beginning: we should separate an account *of* analog/structural representation (ASR) from an account of representation more generally. As mentioned earlier, everything discussed here is meant to be a characterization of a particular type of representation, but it is not meant to be an account of representation *simpliciter*. We can first understand precisely what makes a type of representation the type it is, *given* that it is a representation. In making the same point, Facchin (2024) notes that we ultimately do need an account of what makes any proposed representation a representation in the first place: that is, we need an account of representation *in general* that applies to analog/structural representations as a particular case. However, we can provide an account of a *particular* type of representation without insisting that that account must also explain why it is a representation at all. Again: the characterizations above are what I take to be *necessary* features of analog/structural representation; but there may well be mappings between two types of things that fit this characterization perfectly well, yet do not involve representation at all (e.g., the temperature of the Nikola Tesla statue in my office likely covaries with the ambient temperature of the surrounding air, yet that statue is not an analog representation of temperature). Of course, identifying such homomorphisms would indicate *candidates* for analog/structural representations, but they are not sufficient to identify whether anything is a representation in the first place.

Nirshberg and Shapiro (2020) note that some proponents of structural representation have argued that the homomorphisms present in this type of representation offer a distinct advantage over so-called "indicator" representations in determining what it is that is being represented. They then argue that this supposed advantage is not an advantage after all: structural representations have the same problem of content determination as the indicator representations. Given a proposed structural representation, it will often turn out that there is a homomorphism between many different representations; the mere fact that there is some homomorphism does not fix what the proposed representation does, in fact, represent. As far as it goes, this is completely correct and illustrates the very point at issue here. In practice, we are quite likely to have to resort to independent criteria to establish that some homomorphism is, in fact, analog/structural representation. In many cases, these criteria are obvious: photographs represent what they do because of our knowledge about how photographs work, our practices of what we use them for, etc. Liquid thermometers represent temperature, and not something else, because of how we have designed them and how we use them. Nevertheless, there is value in identifying homomorphisms that are candidates for analog/structural representation, which can then be further investigated to determine whether they are representations after all.

Similar concerns are presented in Artiga (2023), where increasingly strong characterizations of structural representation are shown to fail to identify whether some proposed homomorphism between two domains is, in fact, a representation. For example, there may be a map (i.e., a two-dimensional homomorphism) between area V1 in the visual cortex and points in the organism's visual field. But the existence of such a homomorphism may not play any role in downstream neural processing: that particular organization may be an accidental result of, say, energy minimization in the development of the visual cortex (Artiga, 2023). Increasing the stringency of the proposed characterizations of structural representation makes it more difficult to provide examples of proposed structural representations that are not representational; however, these involve invoking mechanisms, processes or operations that manipulate these representations.

Now, to be sure, many investigations of the representational capacities of a system proceed by identifying candidate representations along with the operations on those representations. In practice, these activities often go hand-in-hand. In neural systems (and in artificial systems that we do not completely understand, such as contemporary large language models and other deep neural networks) we need to see how the candidate representation changes in response to candidate representata, what mechanisms are responsible for those changes, and what downstream effects those changes have. In other words, we need independent reasons to determine that what is under consideration is a representation in the first place. In this respect, Nirshberg and Shapiro (2020) and Artiga (2023) are both correct: the proponent of structuralism should not advocate the idea that an account of structural representations *alone* is sufficient to establish that a purported representation is really a representation in the first place. At the same time, carefully identifying candidate representations and representata, including whether both have an ordered structure among their elements and whether there is a homomorphism that preserves that structure, goes a long way in determining potential representations.

3.2 Clarifying Types of Structural Representation

Recently, Facchin (2024) has provided a fourfold distinction between different types of structural representation based on different conceptions identified in the literature. Although the distinction offered there is insightful, there are some problems that can be clarified or addressed using the analog/structural account here. Most importantly, Facchin claims that the types of structural representations he articulates are different enough to qualify as genuinely distinct notions of structural representation; that claim is significantly undermined when viewed through the lens of the analog/structural thesis advocated here. Let us look at Facchin's distinctions and see what we can make of them in light of this thesis.

According to Facchin (2024), there are four different conceptions of structural representations that can be identified in the literature: STRUCTURAL MAPS, STRUCTURAL SIMULATIONS, STRUCTURAL SPACES, and STRUCTURAL DYNAMICS; I will leave off "structural" when referring to these hereafter. The MAP conception takes structural representations to be individual vehicles that represent their target by being structurally similar to their target (i.e., having the right homomorphism) and that can be decoupled from that target, exploitable by the relevant system, and capable of misrepresenting that target. Ordinary maps would be of this type. The SIMULATION conception takes any computational process that maps computational states to the states of a target process as a structural simulation. Facchin uses the example of a simple sales program: it maintains an internal state corresponding to current inventory, then takes as input a state corresponding to sales, and outputs a state that is the updated inventory after subtracting sales. Because these computational states and operations on these states (i.e., subtraction) map onto their target states and the operations on these states (i.e., sales), this is a structural simulation. Next we have the SPACE conception. These representations include liquid thermometers, where there is a graded range of "indicators" that correspond to what they represent, preserving a total ordering from targets to the indicators. Finally, there is the DYNAMICS conception, according to which structural similarity between indicators and targets is preserved in the dynamics between the two. In other words, as target states change, so do indicator states.

Facchin rightly notes that the extension of these different conceptions is different: the SIMULATION conception is limited to computational states, whereas thermometers usually are not a part of a computational system at all. The MAP conception applies to a single vehicle, whereas others apply to a number of vehicles. The SPACE conception need not involve temporal dynamics, whereas on the DYNAMICS conception accounting for temporal dynamics is necessary. However, this difference in the extensions of the different types of representation is not evidence that the types are fundamentally different; this is clear when we appeal to the richer characterization of analog/structural representation offered earlier. Thermometers, photographs, and maps differ in their extension, too, but all are analog representations.

Facchin's SIMULATION conception of structural representation adopts a notion of structure that is too weak. Virtually any computer program represents something: GPAs, tax deductions, emails, prime numbers, architectural layouts, etc. And virtually all of these things have some structure or another. It is unclear what it would take for a computer program to *fail* to be a structural simulation in this very weak sense of "structural." Moreover, any neural process that represents *something* that changes— in any way— would also be a structural simulation on this conception. Suppose that you represent the word "apple" by a pattern of neural firing in your right hemisphere, and it is made plural by activating a completely different neural pattern in your left hemisphere. There is a mapping between a word and its plural form, and whatever process gets us from the pattern in the left hemisphere to the pattern in the right hemisphere will map to the pluralization process. This will be true for any neural representations and any operations involving those representations. However, I assume that structural representations are supposed to contrast with at least *some* kind of non-structural representation, so this conception covers too much. This is simply too weak a conception of "structural," because it does not rule out *any* computational processes or neural processes that represent anything at all.

The difference between the MAP conception on one hand and the SPACE and DYNAMICS conceptions on the other is the use of a single vehicle in the former but many vehicles in the latter two. A photograph (structural map) is a single vehicle representing a single target, whereas in a thermometer (structural space), there are many vehicles (heights of liquid) representing many targets (temperatures). That seems reasonable, but it is also too fast. One way to respond would be to argue that in a photograph, there are actually many vehicles: the points that constitute the photograph (or pixels in the case of a digital photograph). Or, we could agree that there is a single vehicle, but it is composed of many parts, where each part of the representation represents the corresponding part of the target. In fact, this is precisely what the "parts principle" states regarding iconic (i.e., multidimensional-analog) representation: roughly, if R is a representation of T, then parts of R are representations of parts of T (Kulvicki (2015), Burge (2018), and Maley (2023b) offer different perspectives on this principle).

According to the characterization of analog representation given in Section 2, the difference between photographs and thermometers is simply the difference between what dimensions form the independent variables of the representation. We typically use thermometers to represent temperature variation with respect to a change in time; nevertheless, we could also use an array of thermometers to represent temperature variation over a two-dimensional space at a single point in time—this would just be a heat map, like the kind meteorologists use to display temperatures across different regions at the same instances. Additionally, we typically think of a photograph as representing an image at a single point in time; nevertheless, we can use an array of photographs to represent how an image changes over time—this would just be a video. It is not clear how we would classify these types of representation under Facchin's characterizations: if a photograph is a single vehicle, it seems that a heatmap would be, too. However, it is a strange result that a single thermometer counts as multiple vehicles, whereas a heat map representing temperatures from *many* thermometers would count as a *single* vehicle. A video would fall under both the structural map and the structural space conception. Perhaps that is acceptable, perhaps not. In any case, understanding these different types is simpler using the scheme presented here. Moreover, we can capture what makes these different analog/structural representations species of the same family, and we can do so in a non-*ad hoc* manner.

Regarding the DYNAMICS conception, Facchin mentions that theorists often combine this conception with others, even though they may be separable in principle. Note that, on the view offered here, the DYNAMICS conception is not a separate type of structural representation at all, but just an explicit recognition that time itself can serve as a dimension of variation. A thermometer tracking temperature change over time, a video tracking visual scenes, a neural population tracking a moving stimulus—all of these preserve temporal structure in addition to the other structures they preserve.

Artiga (2023) also identifies a number of ways in which structural representation may be understood, positing increasingly stringent conceptions of the notion to rule out potential counterexamples. In the end, Artiga notes that it is incumbent on theorists to be clear about what conception they have in mind, given that there are a variety of possibilities. I take the Analog/Structural Representation Thesis argued for here as a way of answering Artiga's challenge on this front.

3.3 Representation Versus Computation

Making clear that analog and structural representation are the same allows for some much-needed theoretical clarity. It is also useful to distinguish between the representations used by a system and the computations that act on those representations. Clearly distinguishing representations from computations acting on those representations is much easier in artifacts; i.e., analog computers. In neural systems—including artificial neural systems which we do not fully understand—discovering which aspects of the systems are representations and which aspects are the mechanisms that operate on those representations (i.e., computations) will often proceed in parallel, with a clear delineation available only after the relevant system is better understood. Nevertheless, beginning with lessons from systems where a clear delineation *is* already available simplifies how we can productively think about what is involved in representation-using systems. We can make things explicit as follows:

The Analog/Structural Computation Thesis: Analog/structural representation can (and, for clarity, should) be distinguished from the computations (if any) that operate on those representations.

To illustrate where this thesis can be usefully deployed, consider again the conception of structural representation offered by Artiga (2023): a "mechanism in which a structural correspondence between a set of vehicles and a target domain is used to accomplish [a] certain task." Including a mechanism in this characterization immediately discounts representations such as photographs, which are clearly not mechanisms. Now, to be fair, Artiga's concern is with the notion of structural representation as it is found in the cognitive science literature: we should not attribute to Artiga any confusion between representations and the mechanisms that operate on them; more charitably, Artiga is responding to a confusion found in much of the literature. But it is a confusion nevertheless, and although Artiga does not identify it as such, he does implicitly articulate the ways that others may have proliferated the confusion.

The larger point, illustrated by Artiga (and others), is that it is conceptually useful to separate representations from operations on those representations (i.e., computations). A genuine ASR may appear only in trivial operations: consider a simple anemometer Figure (6). This device represents wind speed by the number of rotations of its head per second (or other unit of time), and records the total number of rotations to keep a running average. This is a very simple analog representation, with a very simple operation on that representation. An organism may represent a stimulus in a similar way, where that representation does not enter into any complex downstream operations, but is a representation nonetheless. We would, of course, need grounds to establish why this particular homomorphism counts as a representation, but that is distinct from the complexity of the kinds of further computations that operate on that representation. There is also an analog representation of wind speed in the Norden bombsight, where wind speed is represented as the angle of an input dial. However, the resulting computation is very complicated: several thousand moving parts contribute to the computation of the correct heading and release point for a bomber to unload its payload.

Making clear the distinction between ASRs and the computations operating on

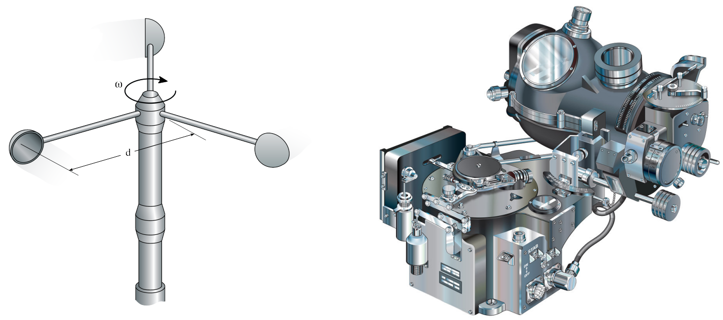

Figure 6: Left: simple anemometer. Right: Norden bombsight.

those representations helps to address a concern that Artiga articulates: that characterizations of ASRs that are too complicated may fail to identify this kind of representation as a legitimate natural kind. I do not know how concerned one should be about this, but assuming that it is a concern, the problem reduces to that of deciding whether, first, ASRs are natural kinds, and second, whether computations are natural kinds. I do not have a pony on that carousel, but these may well be simpler questions to answer than taking ASR to have mechanisms and operations built-in, alongside whatever is doing the actual representation.

3.4 Analog/Structural Representation Versus Indication

The final issue I will address is whether we can make a principled distinction between ASRs and indicators. Indicators are generally thought to be simpler than analog/structural representations: for example, an analog/structural gas gauge might represent the amount of gas by the angle of a "hand," much like a clock, whereas an indicator would simply turn on a light when the gas is below some level (thus indicating that the gas is low). Although there seems to be a clear difference between the binary nature of an indicator as opposed to the graded nature of an analog/structural representation, Nirshberg and Shapiro (2020) challenge this claim with an insightful argument that indicators are simply limit cases of analog/structural representations, and therefore the two are not different in principle. This is true for some indicators, but not all, which can be seen using the characterization of ASR developed here.

Nirshberg and Shapiro (2020) consider the example of a spring scale that varies continuously, representing the mass of an object placed on the scale in grams. Next, we can consider a modification to the scale such that it only represents single-gram increments in discrete steps, perhaps by using a different spring. We can then further modify the scale using two-gram increments. This process can continue with arbitrarily large increments, each time resulting in an increasingly low-resolution (i.e., coarse-grained) representation of an object's mass. As they say (using "spring*N*" to stand for the spring that makes the scale "maximally insensitive"): "Suppose that the crucial weight that marks the boundary point at which spring*N* 'jumps' in length is 100 kg. The scale that makes use of spring*N* is now, in a sense, maximally insensitive. It will 'detect' something hanging from its end when and only when the object's weight is 100 kg or greater," (Nirshberg and Shapiro, 2020, p. 7659). The authors argue as follows: in the first (i.e., continuous) case, we clearly have a structural representation; discretizing the representation in increments of whole numbers of grams does not alter the structural preservation needed to maintain its status as a structural representation; thus, the final case of 100 kg increments is a structural representation, too. They conclude that indicators are simply limit cases of structural representations, but structural representations nonetheless.

The authors do not mention this, but for the example to work, we have to assume that the scale in question has an upper limit, or capacity, that is below 200 kg (i.e., below the capacity that is another multiple of the 100 kg jump). Without that limit, we would still simply have a very low resolution scale: it would move in *another* discrete step for every additional 100 kg. To consider it nothing more than a detector for objects *above* 100 kg, we have to stipulate that it *cannot* go beyond a single increment and thus has an upper limit *below* 200 kg. Alternatively, if the scale does *not* have an upper limit, but simply moves one step at 100 kg but does nothing for further 100 kg increments, then, by their own lights, we would not have a structural representation to begin with. So, let us assume that for such a detector, there is an upper limit that is less than twice the increment needed for the example (e.g., for a 100 kg increment, we need a capacity less than 200 kg; change the increment to 10 g, we need a capacity less than 20 g).

Nirshberg and Shapiro (2020) are clearly on to something: it is undeniable that one can discretize a continuous structural representation without compromising its status as structural; the same point holds for analog representations, as mentioned earlier. Plus, with the right restrictions in place (just mentioned), decreasing the resolution of a structural representation effectively turns it into an indicator. However, not all indicators are limit cases of structural representations.

To see why, let us think about groceries.[10] Grocery stores often have automated conveyor belts that move continuously until an object is detected at the end of the belt. The motor connected to the belt is thus in one of two states: on or off. An object detector that is in one of two states is all that is needed to make the system work. Now, consider two ways in which this object detector might work. First, we could have a light beam projected orthogonally across the conveyor belt that hits a detector; when an object breaks the beam, it is detected (this is depicted with the green line between the boxes in 7). Second, an object could be detected by measuring the distance from a projected light source at the end of the conveyor belt, then turning off the belt when the distance is below some threshold (this is depicted with the red waves coming from the box at the end of the conveyor belt in 7).

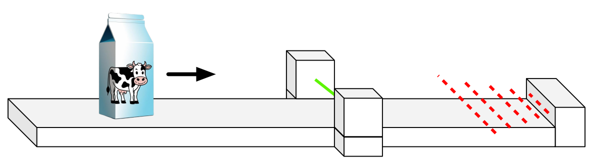

Figure 7: Conveyor belt with milk and two detectors.

To make things simpler, let us imagine that instead of being connected to a motor, these different detectors simply turn illuminate a light when an object is detected. As described, both count as indicators in the technical sense. However, they differ in the following respect: the second, but not the first, is a limiting case of a structural representation. In the second case, we could, as it were, increase the resolution of the representation (i.e., the light): instead of being only on or off, the light could be on at full brightness, medium brightness, or off completely, where the two levels of light "on" correspond to two different distances: full brightness for closest to the detector, medium brightness for further away, and off when furthest away. Or, the light could be at 10 different brightness levels, corresponding to 10 different distances. Or continuously many. The structure being preserved by the magnitude of the light is the magnitude structure of distance: closer objects produce

[10]As Donald Trump explains: "Likewise, an old fashioned term that we use — groceries. I used it on the campaign. It's such an old fashioned term, but a beautiful term. Groceries. It sort of says 'a bag with different things in it,' " 2 April, 2025.

a brighter light, farther objects produce a dimmer light, where "brighter" and "farther" are obviously ordering relations among the respective magnitudes. As long as the monotonic relationship between light intensity and distances is maintained (i.e., the brighter the light, the farther the object), we have a structural representation, as well as an analog representation.

However, for the first case (i.e., the object either breaks the light beam or does not), there is no way we can increase the resolution to recover a structural representation, because there is no underlying magnitude structure to preserve. The light beam is either broken by the presence of an object or it is not: there is no gradation that we can introduce into this mechanism, and no sense to be made of an object being more or less "present."

Now, one might argue that there is, in fact, a structure-preserving homomorphism, albeit a very simple one. Namely, one binary state (object present versus not present) is represented by another binary state (light on or off). It is true that, in a weak sense, there is a structure being preserved. However, this is not the notion of structure on offer in the literature on structural representations. As argued above, structural representations rely on *magnitude* structures being preserved: this is what allows structural representation to be exploitable. Representations that are close together in magnitude, and in a particular order, will have representata that are close together in their magnitude, in that same order. However, with only truly binary variation in cases like this one — where the binary variation is not just a very low resolution variation in a magnitude structure — there is no ordering, and no way for the system to "exploit" anything.

As before, note that everything said in these examples applies equally well to analog representations. Both analog and structural representations; i.e., ASRs, preserve magnitude structures along dimensions of variation. Simple binary detectors of binary conditions do not involve magnitude structures and thus do not count as ASRs, although they may be limit cases where magnitude structure has been reduced to the point of disappearance. Nirshberg and Shapiro (2020) are correct that many cases initially described as "binary detectors" are actually low-resolution analog representations. Their spring scale example demonstrates this perfectly. When we take an analog representation and decrease its resolution to the point where it makes only binary discriminations, it remains an analog representation in principle, even if its informational value is minimal. The difference between these cases can be stated precisely: a binary detector is a limit case of ASR if and only if there exists an underlying magnitude structure that could be represented at higher resolution. The distance detector qualifies; the light-beam detector does not.

This distinction matters for cognitive science. When investigating whether a neural system employs structural representations, we can ask whether the system preserves magnitude structure, or merely detects binary conditions. A neural population that responds gradually to increasing stimulus intensity is plausibly an ASR (even if, in a particular experiment, we only measure whether it exceeds a threshold). A neuron that simply fires or not in response to the presence/absence of a specific stimulus might not be, unless there is evidence that firing rate, temporal pattern, or population activity preserves underlying magnitude structure.

Therefore, while some indicators may be limit cases of ASRs (supporting Nirshberg and Shapiro's insight), not all indicators are ASRs (contradicting their general conclusion). The deciding factor is whether magnitude structure is being preserved, even at very low resolution.

4 Conclusion

By understanding "analog" and "structural" as two labels for the same type of representation, we can clarify our general understanding of this unique type of representation. This is particularly important in neuroscience, cognitive science, and artificial intelligence, where ideas about representation come from diverse sources and play different roles in different areas of investigation. Conceptual resources from computation, such as the notion of "representation," play a crucial role in understanding brains, minds, and AI systems. The fact that we can avail ourselves of a richer notion of analog representation by more carefully understanding analog computation has gone largely unnoticed, as has the strong connection between analog representation and structural representation.

The theoretical unification achieved by recognizing structural representations as analog representations is not merely terminological. It allows us to apply the rich framework developed for understanding analog representation — including analyses of resolution, noise, continuous vs. discrete implementation, representational capacity, and computational operations — directly to structural representations in cognitive science.

Important questions remain about how to count something as a representation in the first place; that is a difficult problem that this paper does not answer. However, when we have a clearer idea of what kinds of representation there are, we can better assess whether particular systems have elements that are candidates for given types of representation. When investigating systems that we suspect are capable of representing—and then performing computations on these representations—it is critical that we have a clear idea of what a representation could be. My hope is that the characterization of analog/structural representation given here can contribute to this effort.

References

Artiga, M. (2023). Understanding Structural Representations. *The British Journal for the Philosophy of Science*. Publisher: The University of Chicago Press.

Beck, J. (2018). Analog mental representation. *Wiley Interdisciplinary Reviews: Cognitive Science*, 10(3):e1479–10.

Burge, T. (2018). Iconic representation: Maps, pictures, and perception. In Wuppuluri, S. and Doria, F. A., editors, *The map and the territory: Exploring the foundations of science, thought and reality*, pages 79–100. Springer International Publishing, Cham.

Facchin, M. (2021). Structural representations do not meet the job description challenge. *Synthese*, 199(3):5479–5508.

Facchin, M. (2024). Maps, Simulations, Spaces and Dynamics: On Distinguishing Types of Structural Representations. *Erkenntnis*.

Gładziejewski, P. and Miłkowski, M. (2017). Structural representations: causally relevant and different from detectors. *Biology and Philosophy*, 32(3):337–355. Publisher: Springer Netherlands.

Kaas, J. H. (1997). Topographic Maps are Fundamental to Sensory Processing. *Brain Research Bulletin*, 44(2):107–112.

Kulvicki, J. (2015). Analog representation and the parts principle. *Review of Philosophy and Psychology*, 6(1):165–180. Publisher: Springer Netherlands.

Lee, A. Y., Myers, J., and Rabin, G. O. (2022). The Structure of Analog Representation. *Noûs*, pages 1–29. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/nous.12404.

Lewis, D. K. (1971). Analog and digital. *Noûs*, 5(3):321–327.

Maley, C. J. (2011). Analog and digital, continuous and discrete. *Philosophical Studies*, 155(1):117–131. Publisher: Springer.

Maley, C. J. (2018). Toward analog neural computation. *Minds and Machines*, 28(1):77–91. Publisher: Springer.

Maley, C. J. (2023a). Analogue computation and representation. *The British Journal for the Philosophy of Science*, 74(3):739–769. Publisher: Oxford University Press.

Maley, C. J. (2023b). Icons, Magnitudes, and Their Parts. *Crítica. Revista Hispanoamericana de Filosofía*, 55(163):129–154. Number: 163.

Nirshberg, G. and Shapiro, L. A. (2020). Structural and indicator representations: a difference in degree, not kind. *Synthese*, 22(3):1–18. Publisher: University of Chicago PressChicago, IL.

O'Brien, G. and Opie, J. (2004). Notes Toward a Structuralist Theory of Mental Representation. In *Representation in Mind*, pages 1–20. Elsevier.

O'Keefe, J. and Nadel, L. (1978). *The Hippocampus as a Cognitive Map*. Oxford: Clarendon Press. Accepted: 2016-10-08T00:43:54Z.

Peacocke, C. (2019). *The Primacy of Metaphysics*. Oxford University Press.

Shagrir, O. (2012). Structural Representations and the Brain. *The British Journal for the Philosophy of Science*, 63(3):519–545. Publisher: Oxford University Press.

Shea, N. (2014). VI—Exploitable Isomorphism and Structural Representation. *Proceedings of the Aristotelian Society (Hardback)*, 114(2pt2):123–144.

Shea, N. (2018). *Representation in Cognitive Science*. Oxford University Press.

Swoyer, C. (1991). Structural representation and surrogative reasoning. *Synthese*, 87(3):449–508.