Analog Representation and the Parts Principle

John Kulvicki

Abstract Analog representation is often cast in terms of an engineering distinction between smooth and discrete systems. The engineering notion cuts across interesting representational categories, however, so it is poorly suited to thinking about kinds of representation. This paper suggests that analog representations support a pattern of interaction, specifically open-ended searches for content across levels of abstraction. They support the pattern by sharing a structure with what they represent. Continuous systems that satisfy the engineering notion are exemplars of this kind because they are uninterpretable unless they are structure-preserving. Analog representations, so understood, include pictures, images, diagrams, and most graphs. This conception of analogicity also fits well with a line of thought about what makes perceptual states distinctive: they satisfy a "parts principle".

1 Introduction

Old-school photographs, mercury thermometers, and other vintage technologies exemplify analog representation. Along certain dimensions—heights of mercury, or patterns of light and dark—smooth variations yield distinct, well-formed representations. John Haugeland called this "everybody's aboriginal, intuitive idea of analog systems" (1998, 83) and he, following Nelson Goodman (1968), built his account of analog devices around it.[1]

This conception of analog representation in terms of smooth variation risks rendering the category quaint. The ubiquity of digital photos, thermometers, and graphical displays has rendered analog technology otiose. Fancy collections of increasingly cheap transistors can imitate any analog device to any degree of precision. Devotees of the analog become niche artisans: flashes of photographers cultivating the darkroom craft, or covens of audiophiles craving vinyl's warmth.

Mapping analog and digital onto the distinction between continuous and discrete is a well-motivated engineering decision. Different formal tools describe each kind of

---

[1] Haugeland's paper "Analog and analog" appeared in 1981. Page references are to the anthologized version (Haugeland 1998).

J. Kulvicki (✉)
Dartmouth College, Hanover, NH, USA
e-mail: kulvicki@dartmouth.edu

system, so it pays, from the engineer's perspective, to keep them apart. But a useful distinction between kinds of *representation* should be articulated along interpretive lines, even if doing so crosscuts engineering categories.

This paper suggests that analog representations are those that facilitate a certain kind of interaction—open-ended searches for content, across levels of determinateness—in virtue of their syntactic and semantic structures. This move re-enfranchises analog representation as central to the arts and sciences. Photos, for example, tell much about the specific color and arrangement of objects, but just as readily reveal general patterns, abstracted from their specifics. This is particularly valuable when we use representations without being certain what we want to know, and it is why analog representation is so important.

Vintage photos and mercury thermometers exemplify both the engineer's conception of analog devices and the interpretive conception suggested here. But many representations that fail the engineer's test fit well within this interpretive category. Being continuous is, in that sense, inessential to being an analog representation. The point isn't to fight over who owns the term 'analog'—the engineering conception is perfectly coherent and well-motivated—so much as to show that we miss an important interpretive category that deserves the name by focusing too much on continuity.

David Lewis (1971) suggested that analog representation involves relating the right sort of magnitude, in the right way, to the right sort of magnitude. Since the right sorts of magnitudes needn't be continuous, neither must analog representation. In that sense, this paper is spiritually in Lewis's camp. He said little about what the right relation is between magnitudes, however, and nothing at all about why it is valuable to relate magnitudes in such a fashion. What follows focuses precisely on those two points.

Section 1 unpacks the suggestion that analog representations must be smooth, or continuous, in terms of three claims. Two focus on syntactic smoothness—the way that relevant features of the representations can vary continuously—and one on a corresponding smoothness in their contents. Section 2 then shows that any readily interpretable system satisfying these three conditions must also satisfy a fourth, which is irreducibly both syntactic and semantic. This structure-preserving condition captures the right way to relate syntax and semantics in an analog system. Significantly, many representations satisfy the fourth condition even though they are not smooth and thus don't satisfy the first three.

Section 3 jettisons the three continuity constraints. Structure-preserving syntactic-semantic links are sufficient to carve out an interesting interpretive category. These links require abstractions over syntactic qualities to map readily onto abstractions over content, enabling the kind of open-ended investigation characteristic of analog representation. Continuous representations exemplify analogicity because they cannot be interpreted absent such structure-preserving features.

Section 4 takes up an idea that has appeared, on-and-off, over the last 40 years in discussions of mental imagery and non-conceptual content. Jerry Fodor is its most recent advocate: "If P is a picture of X, then parts of P are pictures of parts of X." (2008, 173) Though it doesn't seem this way at first blush, this *parts principle* nearly encapsulates what is distinctive about analog representation. It's not unique to pictures, and it's not even a true principle unless we carefully unpack the kind of part to which it refers. Suitably refurbished, however, it captures an important truth about analog representation, and ties the study of artifacts to the study of mental representation.

This paper says nothing of substance about digital representation. The analog's straight-laced cousin deserves more attention, especially given the popularity of the term 'digital' these days. Reshaping the concept of analog representation leaves it an intriguing possibility that some kinds of representation are analog and digital at once, but that intrigue must wait for another occasion.

2 Continuity

The aboriginal, intuitive thought about analog representations is that they allow for continuous variation along certain dimensions. Every height of the mercury column counts as a representation of the temperature. Different heights, even those so close it's hard to tell them apart, amount to different representations. Similarly, the pattern on a picture surface can be tweaked in indefinitely fine ways—lighten it, darken it, change the pattern, just a bit—and the result is still a picture. Set the picture on fire, and things have gone too far: the representation has been changed outside of the licensed dimensions of variation. Three points sharpen this intuitive idea.

First, this aboriginal smoothness is syntactic: defined over changes in the height of a mercury column, or patterns on a picture surface. Those features are important for what the thermometer and picture represent. It's possible to vary a picture's mass, or a mercury sample's purity, in a smooth fashion, but such variation is independent of their functions as representations, and thus not relevant to analogicity. Relatedly, it might be possible to characterize syntactic features which are continuous, at another level, as being discrete. There are only so many molecules in the mercury column, for example. But the syntactic quality is height, so we can ignore possible discontinuities introduced by the irremediable discreteness of molecules (see Moor 1978, 217–218).

Second, smoothness is meant to suggest a continuum of syntactic *types*. Each of the column heights and patterns of light are suited to representing different things. That's not, yet, to say that they in fact represent different things, but that they are suited to doing so. Contrast this with an alphabetic inscription, say of the letter 'G'. Stretch, squeeze, and wiggle the form, within a certain range, and the result is still a 'G'. These smooth changes preserve syntactic identity, rather than changing it. Representations in writing systems, exemplars of digital representation, are quite insensitive to changes in their syntactic features (Kulvicki 2006, Ch2). Analog representations, by contrast, are quite sensitive: small changes result in new representations. In both, it's easy to find smooth variations that preserve well-formed representations.

In line with these first two clarifications, Nelson Goodman cast analog representation as a special failure of *syntactic finite differentiation*:

> For every two characters K and K* and every mark m that does not actually belong to both, determination that m does not belong to K or that m does not belong to K* is theoretically possible. (Goodman 1968, 135–136 emphasis removed)

Within a finitely differentiated system of representation, any mark—an inscription on paper, for example, 'G'—can conclusively be determined not to belong to at least all but one character, or syntactic type. Some scribbles are nonsense, so they do not inscribe any letter. Some are 'G's, and some are close to being 'G' even though it's

hard to tell. They are borderline cases. Borderline 'G's, however, are not also borderline 'C's or 'F's.[2]

Any systems in which it is impossible to determine whether some mark does not belong to more than one character fail to be finitely differentiated. That is, they are structured so that some characters are so similar to one another, at least along a border, that it's impossible to tell some instances of each apart. Analog systems are a special failure of differentiation because every mark that is an inscription of some character is also such that it might be a member of indefinitely many other characters, as far as anyone can tell. For Goodman, this extreme failure of differentiation is *syntactic density*:

> A scheme is syntactically dense if it provides for infinitely many characters so ordered that between each two there is a third. In such a scheme,…no mark can be determined to belong to one rather than many other characters. (Goodman 1968, 136)

Goodman uses 'scheme' to pick out the syntactic types alone, irrespective of what kind of semantic types are attached to them in a full articulation of a representational *system*.[3]

A specific height of mercury is a syntactic type within the system of thermometer representations. But no matter how careful the measurements, there are always indefinitely many syntactic types under which any specific height might fall. Users are cut off from determinate knowledge of syntactic types when consulting mercury thermometers, old-school photographs, or any number of other analog representations.

The first two clarifications, in sum, claim that analog representations admit smooth variations in their syntactically relevant features, which amount to smooth changes in syntactic types. This blocks knowledge of the specific syntactic type into which any given representation falls.

Third, these myriad syntactic types are saturated with meaning, in the following sense. Each one of the distinct types represents something distinct from the others. Each height of mercury stands for a distinct temperature, and each pattern of light and dark represents a distinct scene. Unlike the first two conditions, this one is semantic. The vast variety of syntactic types is matched by an equally vast range of possible meanings. Analog systems are semantically rich (Kulvicki 2006, Ch2). For Goodman, this condition amounts to the claim that, like the syntactic types, the semantic types are dense: they can only be ordered such that between each two there is a third (1968, 160). Given a temperature, for example, there will always be indefinitely many other temperatures from which it cannot be distinguished, no matter how precise the measurements. This minimal semantic condition says little about how, exactly, to correlate

---

[2] Perhaps some inscriptions are such that we cannot tell whether they are 'C's, 'G's, or neither. In that case, the Roman alphabet would fail the test of syntactic finite differentiation. See Goodman (1968, 137–139) for discussion. This is not directly relevant to syntactic density, however, so we can leave the point here.

[3] Matthew Katz (2008, 405) also stresses the fact, following Goodman, that in analog representation we cannot know the syntactic identity of any given representation, but he, like Goodman, didn't notice any semantic consequences of this fact. As we will see in the next section, this is precisely where we find the action in understanding analog representation.

meanings with syntactic types, keeping open the possibility that "the compliants under an analog system may be as remote and distant as we please from the characters." (1968, 160)[4]

The first two conditions are purely syntactic while the third is purely semantic. Each helps understand the aboriginal intuitive idea of analog representation, but even together they fail to capture everything important about it. The condition we have not yet seen is irreducibly syntactic and semantic. A puzzle about how we interpret representations satisfying the three continuity conditions motivates this structure-preserving constraint. It casts serious doubt on the most extreme reading of Goodman's claim that the compliants of an analog system can be as distant as one likes from the representations.

3 Structure preservation

Analog systems of representation are such that there are indefinitely many distinct syntactic types and each type has its own associated content, distinct from the rest. Confronted with a specific representation in such a system, a user is unable to tell, just by looking, which representation it is. Indefinitely many representations cannot be distinguished from the one that's present. And yet, despite this profound lack of access to the representation's identity, it is rather easily interpreted. In practice, people readily and accurately interpret analog representations. How is this possible, if users are, perhaps in principle, shut off from analog representations' syntactic identities?

The obvious answer is that it is easy to know an analog representation's syntactic features to some limited degree of precision. The mercury clearly falls somewhere between the marks for 73 and 74°, even if it is impossible to see that the height is precisely 73.510 …. But this fact does not solve the interpretation problem, absent a very strong assumption about the link between thermometer syntax and semantics. Namely, it's only if having a height within a certain range amounts to representing a temperature within a certain range that an approximation to the syntax readily yields an accurate, if approximate, interpretation of the representation. The fourth condition, which will need some unpacking, is that

> Approximations to a representation's syntactic identity are readily mapped onto approximations to its semantic identity.

In the mercury thermometer, the column heights are isomorphic to the temperatures they represent. There is a one-one mapping between heights and temperatures, contiguous ranges of heights map onto contiguous ranges of temperatures, a height between two others maps onto a temperature between those mapped by the other heights, and so on. Failure to know the column's height does not undermine one's ability to interpret

---

[4] Haugeland conflated the second and third points just made. Analog systems are *sensitive*, in Haugeland's sense, when "not only are all variations allowed, but they all matter." (1998, 83). In line with the second point above, they all matter syntactically, and they also, in line with the third point, matter semantically. At the end of this section it will be clear that this conflation is reasonable, but to see why it's reasonable these two conditions need to be kept apart.

the representation because any approximation to that height maps onto an approximation to the specific temperature it represents. In other examples, the relationship might be less robust, and thus less readily understood. At a minimum, however, approximations to the syntax must map in a comprehensible way onto approximations to the semantics. Failing that, the obscurity of determinate syntactic identity will undermine attempts at accurate interpretation.

Consider an example that fails to meet the fourth condition. Take a thousand small ranges of temperatures and reassign them, willy-nilly, to distinct parts of a mercury column.[5] This scattershot map from heights to temperatures satisfies the first three conditions. Indefinitely fine differences in column height amount to distinct syntactic types, each of which has its own associated content. Because approximations to syntactic identity are not approximations to semantic identity, however, this system is useless. For example, though one can tell that the column is within a certain range of heights, that range of heights represents a vast variety of temperatures. One cannot grasp the mapping so as to know the temperatures represented by heights in that range. Continuous movement through the heights of mercury is a discontinuous romp around ranges of temperature.

The fourth condition suggests that an analog mapping between syntactic features and semantic features is structure-preserving. Ranges of syntactic qualities map onto ranges of semantic qualities, so that approximations to the former map onto approximations to the latter. This might seem trivial: any one-one mapping between temperatures and heights is structure-preserving in some sense. Isomorphism is cheap. Two points about the needs this structure preservation fills assuage this worry.

First, the motivation for the fourth condition is that users are unable to determine the determinate syntactic qualities of smooth representations. So, the syntactic structure in question is one understood along the lines of levels of determinateness. Determinate heights elude us, but approximations are easy to make. Which approximations are easy to make? After all, there are many abstractions over or approximations to specific heights. For example, one abstraction over a height of 1 inch is a height between ¾ and 1½ inches. Another is a height that is either 1, 5, .2, or 56 in.. The latter disjunction is less determinate than a height of 1 inch, and in that sense it is an abstraction over that height. But the latter disjunction is never perceptually important. At a distance, it might be hard to discern whether something is exactly 1 inch long, Perhaps all one can tell is that it's between ¾ and 1½ inches long. But not even a set of funhouse mirrors would leave one wondering whether it's 1, 5, .2 or 56 inches long.

Remember that the motivation for the fourth condition is lack of access, under ordinary conditions of use, to the syntactic details of a representation. The access we do have—the kinds of approximations allowed by our ordinary means of perceiving representations—determine the kinds of approximations that are relevant to the fourth condition. Demands of use determine the abstractions that matter.

The second point about structure preservation concerns the semantics. There are many ways to abstract over any specific semantic content, which mirror the options for abstracting over syntactic qualities. Some ways of correlating abstractions over the semantics are readily grasped by users, while others are not. We are interested in

---

[5] This suggestion comports poorly with the mechanism behind a mercury thermometer, but we can leave that point to one side for present purposes.

mappings that are "semantically salient" (Kulvicki 2010, 301–302) in the sense that they are readily grasped.

Though no one has framed the issue this way, many have suggested such syntactic-semantic conditions on analog representation. Kent Bach claimed that analog representations are *continuously correlative*: "…given the pictorial properties by which a symbol represents qualities of its object, a slight change in the symbol in some respect means a slight change in some respect of what is represented." (1970, 128)[6] Haugeland made much the same claim: "not only are all variations allowed, but they all matter….In general, however, the smaller the [syntactic] difference the less it matters [semantically]." (1998, 83)[7] Roger Shepard suggested "intermediate internal states [of analog representations] have a natural one-to-one correspondence to appropriate intermediate states in the external world." (1978, 135) For Shepard, intermediate states of the representation are those between others along a licensed dimension of syntactic variation. David Lewis, though taken to defend a rather heterodox conception of analog representation, says simply that "analog representation of numbers is representation of numbers by physical magnitudes that are either primitive or almost primitive…." (1971, 325) Primitive magnitudes are those expressible in some reasonable language of physics, and implicit in Lewis is that these magnitudes (numbers and physical magnitudes) relate in a structure-preserving manner. and James Blachowicz (1997, 74) and Corey Maley (2011, 123), both inspired by Lewis and Shepard, claim that analog representation involves relational mappings that are structure-preserving in ways compatible with the fourth condition.

Even Goodman, so keen to avoid resemblance talk (1972), sensed that some notion of shared structure is unavoidable with analog representation: "Can it be that—ironically, iconically—a ghost of likeness, as nondifferentiation, sneaks back to haunt our distinction between pictures and predicates?" (Goodman and Elgin 1988, 131) Shared structure is merely a ghost of likeness: thermometers need not get warm and cool in tune with the weather, and graphs of porosity need not be particularly porous. But it is near enough to resemblance to have haunted him. In fact, Goodman never endorsed the fourth condition on analog representation—"Plainly," he suggested, analog systems have "nothing special to do…with analogy" (1968, 160)—though he is one of the few who seems to have seen the key motivation for it. Ready and accurate interpretation of syntactically and semantically dense representations is implausible absent a clear structural relationship between syntactic qualities and the qualities they represent.

In sum, these last two sections considered four conditions that flesh out the "aboriginal, intuitive idea of analog systems." They allow for smooth changes in their syntactic qualities. Those changes are allowed in that the results are well-formed,

---

[6] Bach offers this specifically as an amendment to Goodman's account of pictorial representation. Goodman thought that pictorial representation was simply a species of syntactically and semantically dense representation, in which relatively many qualities—color, shape, shading—matter to a representation's identity. Pictures are *relatively replete* (Goodman 1968, 230).

[7] Haugeland avoids semantics talk in his article, because "in my view digital [and analog] devices are not necessarily representational or symbolic" (1998, 88n2). The position of a potentiometer does not represent the brightness of the lights, but dimmers are analog devices. Notice that even if some analog devices are not representational, analogicity is essentially relational. Changes in the states of one thing relate to changes in the states of another in a systematic, structure-preserving fashion. This paper focuses on analog representation, but the phenomenon is not limited to representation.

distinct syntactic types. The distinct syntactic types are saturated with semantic significance, in that each one is paired with a different content. And the pairing is structure-preserving: approximations to the determinate syntactic features of a representation map straightforwardly onto approximations to its semantic features.

4 Analogicity without continuity

Section 1 connects the engineering conception of analog systems to representations in terms of syntactic and semantic smoothness. Section 2 shows that demands of use for such systems require another, structure-preserving, condition to be met. Absent the structural condition, systems satisfying the first three are uninterpretable. This section shows that the interpretability enabled by the fourth condition has a distinctive character: such representations *support* a valuable *pattern of interaction*. Support for this pattern makes analog representations an interesting interpretive kind and provides a way to conceive of them that departs from the engineering notion. Specifically, many representations fail to satisfy the three smoothness conditions, even though they are structure-preserving in a way that supports the valuable pattern. They are all, I suggest, analog representations. Smooth representations are exemplars of analogicity because they are uninterpretable unless they support the pattern.

The *pattern* in question is an open-ended search of content across levels of abstraction. Users rarely approach representations with no sense of what they are looking for. But the representations that support such open-ended searches are well-suited to a number of ends, like figuring out whether it's too cold for swimming, or whether the roads will be slippery in the freezing rain. Photos can be important for what they reveal about specific colors, general patterns of color, combinations of color and texture, or hueless patterns of light and dark. It's possible to approach descriptions with such open-ended searches in mind, but descriptions are not structured syntactically and semantically so as to support such interaction. The remainder of this section shows how representations that satisfy the structure-preserving condition support this pattern of engagement.

According to the fourth condition, approximations to the syntactic identity of a smooth representation are approximations to its semantic identity. Equivalently, abstractions over the syntactic details correspond to abstractions over the semantic details of the representation. An abstraction over the specific height of a mercury column is simply a less determinate, but accurate, specification of the column's height. It's somewhere between the marks for 73 and 74, rather than 73.510…. Moreover, there are many abstractions over the column's height that also amount to abstractions over its content: being above the 50 mark, being between the 55 and 75 marks, and so on. Similar considerations hold for other paradigmatically analog representations, along their smooth dimensions.

These abstractions isolate syntactic properties of the representations: features responsible for doing semantic work. A thermometer whose mercury sits between 73 and 74 represents a temperature between 73 and 74°, for example, and one cannot change that content without changing where the mercury sits. Wiggling the height within those bounds leaves the content intact: it still represents a temperature between 73 and 74. The same can be said, of course, for the most specific heights and their associated

contents. Wiggling is relevant to determinate cases because small changes to the height mean changes in the specific represented temperature. The syntactic features of these representations have a structure articulated along these levels of abstraction. In fact, it would be difficult to give fully articulate voice to all of these features. Indefinitely many abstractions over the determinate syntactic details fit the bill for being syntactic qualities. The contents of such representations inherit this structure.

What is the content of the thermometer representation? One might just tease as much information as possible from it. Measure the determinate height, as carefully as possible, and state its associated temperature. This suggests that the thermometer's content is fully captured by a careful description of temperature, or even a name which indicates just that value. But stating the most specific temperature does not give voice to all of what the thermometer represents. The indeterminate temperatures—those ranges corresponding to abstractions over the specific heights—are also represented. The motivation for this claim is that there are syntactic features of the thermometer responsible for representing just such features. The thermometer's content spans levels of abstraction, and in that sense is highly *vertically articulate*. To be specific, a representation has a vertically articulate content when it represents its object as being P, but also represents it as being some Q, which is an abstraction from P (Kulvicki 2007, 359; 2010). The thermometer has a highly vertically articulate content. It represents a determinate temperature, and indefinitely many abstractions from it.

Contrast this with a linguistic expression that represents a very specific temperature: 'the freezing point of water'. The temperature described is, in fact, colder than 40° and warmer than −40, but the description does not represent the freezing point of water in that way. Facts about abstractions from the most informative thing a linguistic expression conveys are rarely part of the representational contents of such expressions. Why? There are precious few abstractions over syntactic features of the description that are responsible for doing interesting semantic work. Isolate 'the freezing point of' from the rest and you find a part of the expression responsible for representing the freezing point, just as isolating 'water' gets you the water-representing bit. That's pretty much all you get by playing the abstraction game with the description. It represents no temperatures less determinate than the freezing point of water. Those in-the-know can interpret this representation, and they might know that this temperature is between −40 and 40, but that indeterminate temperature is not represented by the description.

Concerned to prove one understands what 'chartreuse' means, one might offer a description that covers a few levels: it's a green that is yellowish, and sometimes a yellow that is greenish, which could be more or less saturated and more or less bright. Even these cases don't approach the articulateness of the thermometer, however. Indefinitely many abstractions over the determinate height of mercury are syntactic features, with their associated contents. There is no way to give fully articulate voice to the representational content of the thermometer. This is due not just to the fine details it represents, but also to the indefinite abstractions over it.[8]

---

[8] As one might expect, these two are related, which we will see. Coming at things from a different angle, Philip Kitcher and Achille Varzi say pictures and many maps are worth "a vast infinity of sentences." (2000, 377)

Analog representations are systematic in a way that makes their contents highly vertically articulate. This is absolutely essential for representations that are smooth, because their determinate details are inaccessible. Absent such a structure, such representations are uninterpretable. Many systems of representation have highly vertically articulate contents, however, even though their syntactic and semantic features bottom out at some specific level of precision. They thus support the same kinds of interaction as their smooth relatives.

Exchange the mercury thermometer for a digital version whose heights vary step-wise in one-degree increments. It's still easy to abstract over syntactic features and find abstractions over the content. Much the same happens with digital photos, whether their pixels are too small to see or just big enough to make things seem grainy. In all of these cases, the representations are easily interpreted even disregarding their finest details. And in many cases it's important to ignore such details because users' interests lay at some remove from them. Is it too warm for a jacket? Too cold for swimming? No single-degree differences matter to this question, even if the thermometer makes them available. Given any set of interests, some abstractions matter while many others do not.[9]

Both smooth representations and those which merely satisfy the fourth structure-preserving condition support these open-ended searches for content across levels of abstraction. They have syntactic qualities—being a specific height, being within a range of heights, and so on—which span levels of abstraction, and each of these qualities maps neatly onto an abstraction over the content. Different ways of focusing on syntactic qualities amount to different ways of interpreting the representations. Want to know about the reds in a field of flowers? Focus on the colors of the patches representing red flowers. By contrast, imagine a long description of the field. It's possible to focus on the names for reddish colors, but this is not accomplished via an abstraction over the syntactic qualities of the names. Insisting that all reddish color names be written in red, or prefaced by 'red-' would enable such an abstraction, but it would enable nothing like the open-ended set of abstractions possible with a color photo. This supported pattern shows up even in representations with finite grain, like digital photos and thermometers.

Increase the grain of a representation sufficiently, and it ceases to be analog in that it no longer supports this pattern of interaction. Instead of breaking the column into one-degree bits, break it in two: this thermometer is short when the temperature is below freezing and tall when it's freezing or warmer. Few abstractions over these syntactic qualities result in abstractions over the content, and the thermometer has ceased to be analog. Somewhere between it and the one-degree thermometer is a fuzzy border between the analog representations and the rest, but it's not terribly important to draw that line. Its precise location shifts with users' purposes. Someone with various interests in temperature patterns might find a thermometer too chunky to be of use, while it is analog as far as most users are concerned. In general, if the variety of abstractions over

---

[9] Notice that there are not indefinitely many abstractions available with the thermometer that has one-degree resolution. This is because there are not indefinitely many ways of abstracting over the finitely many readings the thermometer can muster. There are more than we could ever care about, but there are not indefinitely many. This is how resolution—the maximum information delivered by any given reading—interacts with the number of available abstractions.

the most specific syntactic qualities of a representation exceeds the number required of the representation in ordinary contexts of use, then the representation is analog. If not, then it's not.[10]

Matthew Katz (2008) decouples smoothness from analogicity by suggesting that the analog representations are those whose finest syntactic details are inaccessible to ordinary users in ordinary contexts of use. They seem smooth, even if they are not. Digital photos are analog because their pixles are so small. Goodman made a related claim to the effect that pixelated photos, when treated as pictorial representations, are interpreted as members of syntactically and semantically dense schemes (Goodman and Elgin 1988, 125). The current proposal suggests that in analog representation either the determinate syntactic qualities or the abstractions they make available vastly outrun the purposes to which they are put in ordinary contexts. In that sense, many aspects of analog representational content are always, as it were, out of view, but this need not be because of imperceptible grain.

Blachowicz (1997) and Maley (2011), who develop ideas in Lewis (1971) and Shepard (1978), are much closer to the current view. They suggest that analog representations are just those which satisfy a structure-preserving relation, akin to condition four in the previous section. Since this does not require smoothness, analog representations need not be smooth. Neither Blachowicz nor Maley ties his claim to the uses supported by such representations, or to the ways in which such representations have vertically articulate contents. They focus on the syntactic-semantic links to the exclusion of the engagement they enable. Both would classify the very simple thermometer—tall/warm, short/freezing—as analog. Hardly a damning admission, to be sure, but it misses the special interpretive kind isolated here.

In sum, analog representation requires structure-preserving syntactic-semantic links that result in representations with vertically articulate contents. Abstractions over the syntactic qualities map readily onto abstractions over the content, and so the content is only articulable as spanning levels of abstraction. These features enable a searching kind of engagement, across levels of abstraction, which is not found in other kinds of representation. The line between the analog and the rest is found where the abstractions made available outstrip the needs or ordinary users in ordinary contexts of use. Smooth systems exemplify analogicity because they make indefinitely many abstractions available, and are uninterpretable without them.

The class of analog representations neatly coincides with another broad category of representation. John Haugeland called the class of images, graphs, pictures, and their close relatives "icons" (1991), and I suggest we call them the "images, generally speaking" (Kulvicki 2014). The engineering notion of analog systems cuts across these categories, since some pictures, images, graphs, and diagrams are smooth while others are not. But all of these representations support the pattern of interaction at issue and all of them thus turn out to be analog. One might suggest that because we already have the category of icons or images, generally speaking, that is uninteresting to find another name for them. But this would miss the motivation for the account. We have converging reasons for thinking a certain category is interesting. Neither Haugeland (1991) nor

---

[10] I am not claiming that a representation is digital if it fails to be analog. No account of digital representation is on offer here. Analog and digital make a nice contrast, but no account of the distinction suggests that they jointly cover the range of representation. So, one needn't offer an account of both at once.

Kulvicki (2006, 2014), for example, wind up with these broad categories by considering the nature of analog representation. Haugeland (1981) even insisted that analog representations are smooth, and never noticed that they support the pattern of interaction so central to the present discussion.

5 Analog minds?

The pattern of support characteristic of analog representation is interestingly perceptual. Abstractions over syntactic qualities are essential in some cases because syntactic detail is out of perceptual reach. Discrete representations are analog when perceptually salient abstractions over syntax readily map onto abstractions over the semantics, even if their syntactic details are readily determined. These thoughts about artifacts dovetail well with some ideas about what makes sensory and/or perceptual states distinct from thoughts. Perceptual states are analog, or so the following like of thought suggests, and this might help explain the significance of analog representation in artifacts. The artifacts exploit our perceptual capacities by mimicking the relationship between perceptual states and the world they represent. This is an early modern thought that becomes much more interesting and specific in the present context.

Consider the following observation by Elliot Sober: "Every representational picture has representational subpictures; successive snippings never destroy representationality" (1976, 124). He was looking for a feature of pictures that makes them distinct from linguistic representations, and which might help explain the nature of perceptual representation. Stephen Kosslyn (1980, 33; 1994, 5) and Michael Tye (1991, 44) both made similar claims when discussing what might make a perceptual state imagistic. More recently, Jerry Fodor offered a "Picture Principle" to the effect that "If P is a picture of X, then parts of P are pictures of parts of X." Explaining the principle, he says that pictures "don't have *canonical* decompositions into parts…Take a picture of a person, cut it into parts whichever way you like; still, each picture part pictures a person part." (Fodor 2008, 173)[11] It's unclear who first made this point explicitly, but it is clearly implicit in work on pictures as far back as Alberti (1991/1435).

Strictly speaking, some ways of cutting pictures run afoul of these claims. Slicing off part of a picture's backside does not yield a picture of a part of anything. The principle applies to cuts affecting pictures' syntactic qualities. Snip a syntactically significant spatial part from a photo, and the result is a photo of a spatial part of the scene it depicts. While Sober and Fodor seem to have contiguous spatial parts in mind, even noncontiguous spatial parts of the picture depict (noncontiguous) spatial parts of the scene.

The focus on spatial parts makes sense for three reasons. First, this principle comes from philosophers of mind and psychologists concerned about mental representation.

---

[11] These claims are all formulated in terms of "pictures", but they can actually be understood a lot more broadly, as we will see. Casati and Varzi (1999) appeal to a related claim in working out a semantics of maps and Kulvicki (2006, 59n2) discusses it in relation to pictorial transparency. See Balog (2009) for a critical discussion of Fodor's proposal and Kulvicki (2014, Ch 8) where I discuss the following way to think of the parts principle, though I did not notice the connection there to analog representation.

Parts of the brain are organized in space in a manner that mirrors the organization of the retina and visual field (Van Essen, Newsome, and Maunsell 1984). The brain's colors and textures, by contrast, do not suggest that it represents color and texture imagistically.

Second, in a smooth system of pictorial representation, any spatial part of a picture, P, can represent a scene as having any quality—color, texture, shape—that P can represent, to the same degree of determinateness. In that sense, it's easy to be confident that spatial parts of smooth pictures are also *pictures*, even absent an account of the nature of depiction.

Third, the previous point can be generalized by noting that spatial locations act like independent variables, in that they are the spots filled by different values of hue, saturation, and brightness in a photo. Forming a subset of those locations leaves the range of possible values for each location intact. In that sense, the process does not "destroy representationality," as Sober puts it, even though it might destroy pictoriality to move from four million pixels to a subset of four. A four-pixel photo's limited spatial palette might undermine the claim that it is a picture, but it's not off-base to understand it as representing something.

The focus on spatial parts is unfortunate, however, because it confounds two senses of 'part', which happen to coincide in these cases: parts in the sense of spatial regions, and parts in the sense of aspects of the syntactic qualities of a representation. The latter sense is the more significant one. I've suggested elsewhere that abstracting over a representation's syntactic qualities is a way of focusing on a part of it (Kulvicki 2014, Ch8). We don't usually think of abstractions as parts, but this suggestion helps refine the Sober idea about images that has been floating around for the past 40 years. In the present context, it suggests that what makes perceptual representations distinctive is that they are analog representations.

All analog representation allows for interpretation absent detailed knowledge of its syntactic features. The previous section focused on levels of determinateness. Photos make sense even when seen blurrily, or when poor lighting makes their colors indistinct. Under these difficult circumstances, interpretation is diminished—lacking detail, commensurate with the lack of syntactic detail—but it is not inaccurate for all of that. Cutting a picture is an abstraction over its syntactic spatial features. Focus shifts from the whole to one or more contiguous portions of it, excluding the rest. In a continuous system of representation, cuts have no consequences for the level of determinateness, but in pixelated ones they do. Less spatial detail is possible in four pixels than in four million.

Once cutting a picture is understood as an abstraction over its syntactic features, we see a clear path to a refurbished parts principle, which has much greater scope and explanatory force than Sober's:

> Abstractions over the syntactic features of a representation, R, are representations of abstractions over R's content.

What does it mean for an abstraction to be a representation? Imagine two photos, alike except for that one is color and the other black-and-white. The black-and-white photo says nothing about color, but otherwise says much the same as the color photo. Now imagine ignoring the color of the color photo and focusing instead on its patterns

of light and dark. This might be difficult to do, but doing so is abstracting over the syntactic features of the color photo. The result is a representation whose content is an abstraction over that of the color photo. The claim is *not* that the black-and-white photo is an abstraction from the color photo. There is nothing abstract about the black-and-white photo. Instead, the black-and-white photo's syntactic features are an abstraction from the syntactic features of the color photo, and its content is thus also an abstraction from the color picture's content. Remember that the hues of a black-and-white photo—black, white, grays and sepias—are not syntactically relevant: lightness and darkness are all that matter. And the point of the black-and-white photo in the present context is just to illustrate how one might abstract from the syntactic details of the color photo and arrive at a representation of an abstraction from the color photo's content.

Seen in this light, patterns of hue, abstracted from the brightness and saturation that usually accompany them are abstractions from the details of a color photo, just like patterns of light and dark abstracted from hue. High-frequency spatial patterns, abstracted from low-frequency patterns, or the reverse, are also in this sense parts of pictures. And what's interesting about this in the pictorial case is that all of these representations are interpretable. These subtler, non-spatial snippings don't destroy representationality, either.

It should be clear that all analog representations satisfy the refurbished parts principle. They are structure-preserving in a way that makes abstractions from their most specific syntactic qualities map neatly onto abstractions over their contents. Indeed, the refurbished parts principle is very close to the condition that distinguishes analog representation from the rest.

Philosophers of mind and psychologists appeal to the parts principle in an effort to articulate a kind of representation that is both distinct from conceptual representation and ideally suited to playing the role it does in perceivers. It is beyond the scope of this paper to discuss mental representation in detail, but for now it pays to notice two points and one complication. First, we have already seen that such representations as these have representational contents that are quite rich, not just in the details they represent, but in the number of abstractions from them that they support. This is exactly what one would want perceptual states to do, as the needs of perceivers often span levels of abstraction from the most determinate things their eyes and ears tell them about the environment. [12] Second, this support for abstraction is notably absent from linguistic representation. Many accounts of concepts likewise suggest that their representational contents are nowhere near as rich as the contents of such analog representations. Here again, the refurbished parts principle could play the role it was designed to play in the theory of mind, and not just in the theory of artifacts.

The one complication with saying that perceptual states are analog is, of course, that we do not perceive them. The relevant abstractions over their syntactic qualities cannot be understood in terms of perceptual salience. In fact, the abstractions over syntactic

---

[12] Fred Dretske (1981, 135–141) appropriates the analog-digital distinction in the service of understanding the distinction between sensation, on the one hand, and thought-like states, on the other. He uses the terms 'analog' and 'digital' in a distinctive manner that does not fit with the present proposal. Elsewhere, however, I suggest a way of emending his view that makes it easy to show how the present conception of analog representation fits well with his account of sensory states (Kulvicki 2004, 383–389; 2007, 361–364).

qualities that matter for perceptual representations, if this is the right way of thinking about the issue, play a role in determining which abstractions over artifacts' features are perceptually salient. For now, it suffices to note that the idea behind analog representation is that syntactic structure supports certain kinds of interpretation. Perceptual states are analog if we can identify syntactic features and show that they preserve the structure of their contents in such a way as to enable abstractions over their contents. If perceptual abstraction is accomplished some other way, and this seems like an empirical question, then we cannot port the account of analog artifact representation to the mental realm. The work mentioned in this section suggests that the proposal has good, if not conclusive, reasons behind it.

6 Conclusion

Anyone committed to the view that analog representation is essentially smooth is also committed to a rather strong structure-preserving condition. Smooth representations are uninterpretable absent the proper connection between their syntactic qualities and their contents. The proper pairing requires abstractions from determinate syntactic qualities to map neatly onto abstractions from determinate contents. Such a condition in place, representations are readily and accurately interpretable even if one cannot, in principle, know the details of their syntactic features.

Representations can satisfy the structure-preserving condition without being smooth, and this suggests another account of analogicity. Analog representations are those that support a specific kind of engagement in virtue of satisfying the structure-preserving condition. In an important sense, the content of an analog representation exceeds the needs of those who use it. In virtue of satisfying the structural condition, analog representations have contents that span levels of abstraction: they are vertically articulate. Users can access these aspects of content simply by abstracting over syntactic features of the representations. They thus support fairly open-ended searches for content across levels of abstraction. As the number of abstractions available in this fashion decreases, to the point that they do not outrun the needs of ordinary users, in ordinary contexts, representations cease to be analog.

This way of understanding analog representation places it centrally among the kinds of representation we use. Pictures, images, graphs, thermometers, and many other representations are analog, even if they are not smooth. Though this paper offers no account of digital representation, it is an open possibility, given this way of conceiving things, that some analog representations are also digital. Smooth representations remain exemplars of analogicity, because they must be analog while discrete representations are only optionally so.

It turns out that there is considerable theoretical value in thinking of perceptual states as analog, as the refurbished parts principle suggests. Parts of representations, in the sense relevant here, are abstractions from the determinate details of their syntactic qualities. Parts of contents are, similarly, abstractions from the determinate details of some content. So understood, all analog representations are such that parts of representations of X are representations of parts of X. Those who introduce the parts principle typically want to explain why a certain kind of mental representation is (1) distinct from conceptual representation and (2) useful. Though it is beyond the scope of

this paper to discuss mental content in detail, the hope is that this refurbishment of the parts principle helps philosophers of mind along both fronts.[13]

References

Alberti, L. 1991 (1435). *On painting*. C. Grayson, trans. London: Penguin.

Bach, K. 1970. Part of what a picture is. *British Journal of Aesthetics* 10: 119–137.

Balog, K. 2009. Jerry Fodor on non-conceptual content. *Synthese* 170: 311–320.

Blachowicz, J. 1997. Analog representation beyond mental imagery. *Journal of Philosophy* 94(2): 55–84.

Casati, R., and A. Varzi. 1999. *Parts and places*. Cambridge: MIT Press.

Dretske, F. 1981. *Knowledge and the flow of information*. Cambridge: MIT Press.

Fodor, J. 2008. *LOT2: The language of thought revisited*. Oxford: Clarendon.

Goodman, N. 1968. *Languages of art*. Indianapolis: Hackett.

Goodman, N. 1972. Seven strictures on similarity. In *Problems and projects*, 437–446. New York: Bobbs-Merrel.

Goodman, N., and C.Z. Elgin. 1988. *Reconceptions in philosophy and other arts and sciences*. London: Routledge.

Haugeland, J. 1981. Analog and analog. *Philosophical Topics* 12: 213–225. Anthologized in Haugeland 1998.

Haugeland, J. 1991. Representational genera. In *Philosophy and connectionist theory*, eds. W. Ramsey, S. Stich and D. Rumelhart. Hillsdale, NJ: Lawrence Erlbaum. Reprinted in Haugeland 1998.

Haugeland, J. 1998. *Having thought*. Cambridge: Harvard University Press.

Katz, M. 2008. Analog and digital representation. *Minds & Machines* 18: 403–408.

Kitcher, P., and A. Varzi. 2000. Some pictures are worth 2-to-the-aleph-null sentences. *Philosophy* 75: 377–381.

Kosslyn, S. 1980. *Image and mind*. Cambridge: Harvard University Press.

Kosslyn, S. 1994. *Image and brain*. Cambridge: MIT Press.

Kulvicki, J. 2004. Isomorphism in information-carrying systems. *Pacific Philosophical Quarterly* 85: 380–395.

Kulvicki, J. 2006. *On images: Their structure and content*. Oxford: Clarendon.

Kulvicki, J. 2007. Perceptual content is vertically articulate. *American Philosophical Quarterly* 44(4): 357–369.

Kulvicki, J. 2010. Knowing with images: medium and message. *Philosophy of Science* 77(2): 295–313.

Kulvicki, J. 2014. *Images*. London: Routledge.

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

Maley, C. 2011. Analog and digital, continuous and discrete. *Philosophical Studies* 155: 117–131.

Moor, J. 1978. Three myths of computer science. *British Journal of the Philosophy of Science* 29: 213–222.

Shepard, R. 1978. The mental image. *American Psychologist* 33: 125–137.

Sober, E. 1976. Mental representations. *Synthese* 33(1): 101–148.

Tye, M. 1991. *The imagery debate*. Cambridge: MIT Press.

Van Essen, D., W. Newsome, and J. Maunsell. 1984. The visual field representation in striate cortex of the macaque monkey. *Vision Research* 24: 429–448.