

Knowing with images: Medium and message

John Kulvicki, Dartmouth College

Forthcoming in *Philosophy of Science*

Abstract

Problems concerning scientists' uses of representations have received quite a bit of attention recently. The focus has been on how such representations get their contents and on just what those contents are. Less attention has been paid to what makes certain kinds of scientific representations different from one another and thus well-suited to this or that epistemic end. This paper considers the latter question with particular focus on the distinction between images and graphs on the one hand and descriptions and related representations on the other. The paper claims that images are special in that they present information *immediately* across many levels of abstraction. Immediacy is explicated in terms of extractability of information, syntactic salience, and semantic salience. Thinking of representations in this way sheds light on the way in which images are and should be used in science.

0. Introduction

There are two problems with understanding how scientists use representations. On the one hand, what makes graphs, diagrams, pictures, images, descriptions, lists and so on *about* other things? What is it about our intentions toward them, their intrinsic features, and the relations between them that conspire to make these artifacts represent other things, and what exactly do they represent when they succeed in doing so? This question has received a good amount of attention recently. (See, e.g., Bailor-Jones 2003, Bartels 2006, Callender and Cohen 2006, French 2003, and Suárez 2003, 2004.) On the other hand, we have the problem of what makes certain kinds of scientific representations different from one another and thus well-suited to this or that end? Scientists use representations as tools for diagnosis and discovery, storage and comparison, argumentation, communication with each other, the public, and so on. Certain representational kinds are well-suited to certain ends while others are not, and comparatively little has been said about why this is so. This paper considers the latter question with particular focus on the distinction between images and graphs on the one hand and descriptions and related representations on the other. In some sense, images present vast amounts of information in a rather immediate fashion, but the obvious things to say about this immediacy are inaccurate and misleading. In what does such immediacy consist and in what sense do descriptions fail to manifest it? What are the epistemic advantages and disadvantages of using such representations?

Section one introduces the kinds of issues on which this paper focuses. Section two explains what it is for a representation to make its content *immediately* available. Immediacy consists of three things: the extractability of information, syntactic salience, and semantic salience. Immediacy is not distinctive of images and the like: any representation makes some aspects of its content immediately available. What distinguishes images, as section three explains, is that they make information across many levels of abstraction immediately available. Section four explains how representations can differ in the way they make abstractions from their determinate content immediately available, and how this helps and hinders our epistemic goals. Section five discusses why this feature of many representations, images in particular, confers such epistemic advantages.

1. *The general contrast*

It helps to begin with a plausible story, which has some helpful flaws, about the aptness of certain kinds of representation for presenting different kinds of data. Descriptions are good at providing coarse-grained, qualitative and quantitative information. They are good at this because they can present information at arbitrary levels of abstraction. A description need not, for example, to reveal the determinate values of temperature over time: it can reveal a general trend without going into any details. This is not to deny that descriptions are also good at presenting fine-grained, quantitative information, but this is not *distinctive* of descriptions. Hector Levesque points out that "the representational expressiveness of a language ... is not so much in what it allows you to say, but in what it allows you to *leave unsaid*." (1988, 370) Levesque had first-order logic in mind when he made that comment, but the point generalizes quite nicely. Descriptions are good at saying arbitrarily little.

Graphs and images, by contrast, are good at providing fine-grained, quantitative information. They cannot present information at arbitrary levels of abstraction, leaving arbitrarily much unsaid. They can, however, present vast amounts of information about a great many features of their objects. The weather report's radar image tells you just how precipitation intensity changes across New England and a graph can tell you just how temperature changes throughout the day. An fMRI image will let you know where and how brain activity differs between two tasks for an individual, while an x-ray tells properly trained individuals much about the organization of tissue and bone in a body. The plausible story is that images and the like deliver a lot of rather specific information while descriptions and their kin are able to deliver arbitrarily little. If we only need a little bit of information, descriptions are the superior means of conveying it, but if we have a lot of very specific information and we want to deliver it, images are best. The outlines of this plausible story have been traced in discussions of mental imagery—Dennett 1969, Block 1983, Tye 1991—as well as in theories of depiction—Lopes 1996, Hopkins 1998, and see Kitcher and Varzi 2000.

As promised, this plausible story has some helpful flaws that can lead us to a better understanding of why images are valuable tools for presenting information. First, there is a sense in

which descriptions are even better than images and graphs at delivering fine-grained, quantitative information. It's difficult to tell just what the numerical value of temperature is at a given time just by inspecting an accurate graph thereof. Exactly which temperature does *that* point on the graph stand for? And exactly how hard is it snowing over the Champlain valley? That's not to say that images cannot provide this information, but they provide it in a way that makes the specific quantitative bits rather difficult to extract from them. By contrast, descriptions present few such difficulties. They can list the precipitation rates and temperatures to arbitrary levels of precision and they are very easy to decode. In fact, a list of numerals, much more like a description than a graph or image, is often preferable if the specifics are what you want.

Second, images are often used in order to extract rather coarse-grained *qualitative* claims about their objects, not for determining fine-grained, quantitative information. We can see from the fMRI that there is more activity in the infero-temporal cortex than in the pre-frontal cortex. The Doppler radar shows that it's snowing hard in the Adirondacks now, but not in the Green Mountains. The specifics are often beside the point when using images. Images are valuable not because they carry very specific bits of information but because they allow us to get at precisely the abstract bits of information we want. Even when fine-grained, often quantitative information is used to *make* an image, that image is often not used *for* conveying such information. In this sense, images seem to serve a function similar to descriptions—saying very little—while descriptions are often employed to make very specific claims. There is a rift between what the plausible story says images and the like are good for and what we seem to use images for. So what, then, is the key difference between images and descriptions?

The way in which images encode such vast amounts of information allows viewers *readily* to abstract *a great number* of claims from that information. The information sought when viewing an image is often coarse-grained and qualitative, which is similar to what is sought when reading and making descriptions. The way in which images present fine-grained, quantitative information, however, makes much coarse-grained, qualitative info readily available. Such immediate availability of a great many pieces of abstract information accounts for some of the epistemological weight given to images and

graphs, not to mention photographs. Descriptions, by contrast, are very selective in the pieces of abstract information that they provide.¹ That is to say, they typically do not provide a lot of information across levels of abstraction. Sometimes they are rather abstract—the red thing—sometimes they are more specific—the crimson thing—but they rarely provide information across these levels of abstraction—the reddish, red, crimson thing. Though this point will have to be unpacked in more detail below, even in the rare cases in which descriptions provide information across levels of abstraction they do not do so in the rich way that images typically do. This means that there are limits on what one can do with descriptions, as opposed to images. In the right circumstances, of course, those limits might be just what one needs. So, we use descriptions, lists, and so on when we are already in a position to know what pieces of information are the valuable ones and which can be safely discarded. Images are at their best when we have a lot of information, but are still *looking* for what is important. They present a vast amount of information in the service of allowing us to figure out what pieces of information matter most for our purposes. In this way, images are tools for figuring out what is important while descriptions and lists are used when we already know what matters. It is not surprising, then, that images are used to diagnose problems, while they rarely count as diagnoses of problems. The diagnosis itself is a description. It presents the information that matters, and only that information.

It is surprisingly difficult to unpack just what it means to say that images provide viewers with ready access to many abstractions from their determinate content. The next section claims that representations present information to us immediately when they satisfy three conditions. These conditions distinguish representations that are alike in the information they carry with respect to how they make that information available to consumers.

¹ This epistemic distinction has little to do with, for example, the reliability of certain kinds of representation. Images, graphs, and descriptions alike can be reliably or unreliably produced, though photographs, for example, are considered more reliable than drawings in part because of their means of production (Cohen & Meskin 2004). The idea here is that given the reliability of representations we are able to come to know different things about their contents easily or with difficulty, depending on the kinds of representation that they are.

2. Immediacy

A piece of information is immediately available in a representation if and only if satisfies three conditions vis-à-vis that piece of information. First, the information must be *extractable* from the representation. When information is extractable, there is a non-semantic feature of the representation in virtue of possessing which it carries the piece of information in question and *no other, more specific piece of information*. For example, red regions of the Doppler radar image indicate stormy weather of a certain intensity, nothing more. Being red says nothing about the location of such a disturbance: the relative *location* of the red region is responsible for that. A certain kind of curve in a bubble chamber indicates a pion's trajectory, and a proper part of that curve indicates a proper part of its trajectory, nothing more. We know it indicates a *pion's* trajectory because only pions would trace a curve of that shape through the chamber. Extractability concerns how *non-semantic* features of representations are responsible for the information that they convey. Some representations are such that we can identify some of their features, like color of their surfaces, shapes of curves, or what have you, as being responsible for indicating certain aspects of the world that they represent.

This explication of extractability does not presuppose an answer to the first problem of scientific representation. The claim is not that it is in virtue of being red that some object counts as a representation in the first place. It may be that being a representation, and representing, e.g. stormy weather, are matters of a thing's causal history and its relation to our intentions. But representations of that sort, which are red, are taken to be representations of stormy weather. So, the foregoing is consistent with Callendar and Cohen's claim that non-semantic features of representations "may serve as pragmatic aids to the recognition of a representation relation that is constituted by other means" (2006, 78), as well as with claims contrary to theirs (e.g. French 2003 and Bartels 2006).² In addition, there are doubtless many ways to interpret any given artifact as a representation. Nelson Goodman (1976, xx), for example, used the example of treating a Hokusai wood block print as an

² This issue in philosophy of science has been discussed by Karen Neander (1987) and Craig Files (1996) in the context of pictorial representation. Kulvicki (2003) brings it up citing Neander, while Laura Perini (2004) mentions it citing Files. Neander, Files, and Kulvicki were writing in aesthetics/philosophy of art. Callendar and Cohen are right to point out that the philosophy of mind can fruitfully inform the first problem of scientific representation, but it is best not to forget that work on pictures from aesthetics can inform the second.

electrocardiogram. Doing so means one takes on different commitments concerning how non-semantic features of the graph relate to what they represent. This paper offers no account of how these practices conspire to generate such norms of use. Given that they are established, this paper asks questions about those practices.

Extractability is not unique to images. A list of numerals representing locations and temperatures renders many specific pieces of information extractable, since a given triple of numerals indicates a temperature at a specific location and nothing else. But consider the representation 'square', which carries the information that its object is square. Since all square things are four-sided, it carries that information too, but there is no feature of 'square' responsible for carrying information about four-sidedness and nothing more specific than that. Information about four-sidedness is not extractable, but only available via an inference. Contrast this with the representation 'four-sided, with right internal angles, and sides of equal length'. Information about four-sidedness is extractable in the latter case since 'four-sided' carries that information and nothing else. To borrow some terminology from Jill Larkin and Herbert Simon (1987, 67), these representations are "informationally equivalent" but not "computationally equivalent." Necessary for a representation making a given piece of information *immediately* available is that it carries that piece of information in extractable form.

Extractability says nothing about the *consumers* of representations. Whether a piece of information is extractable depends on (1) the content of the representation and (2) how the non-semantic features of a representation relate to its content. Immediacy cannot amount merely to extractability, because immediacy concerns how we use representations. The other two conditions, syntactic and semantic salience, relate consumers of representations to the information that they carry.

In order to present a piece of information immediately, the properties in virtue of which a representation carries that information must be perceptually salient: they need to stand out. Let's call this condition *syntactic salience*, because it says something about consumers of representations relate to their non-semantic features, not about how they relate to their contents. It's difficult to know exactly what representations must be like for their syntactic features to be sufficiently salient, but a philosophy paper is no the place to figure that out in any case. Cognitive psychologists are in a good position to

study the specifics, and it is easy enough to come up with examples of cases in which syntactic features of representations are salient and those in which they are not.

For example, imagine we want to know where a surface has temperatures within a given range: say between 98 and 102 degrees Celsius. We can make an image or graph of that surface that represents temperatures in that range with shades of red, and all others with shades of green. In this case, the features of the representation responsible for carrying the information that interests us stand out. If however, all temperatures are represented with shades of red, albeit different ones, then it will be more difficult, though still possible, to figure out which regions have the temperatures of interest because we are bad at recognizing and re-identifying specific shades of red. These two graphs differ in the syntactic salience of the properties that carry the information of interest. Similarly, if we were to make the *saturation* of the color stand for temperature, so greater saturation stands for greater temperature, but allow the use of arbitrary *hues* within such a graph, almost all syntactic salience will be lost. It is very difficult to sort colors based on saturation alone, especially when their hues differ arbitrarily. Saturation is often confused with brightness, in any case. Such a representation would carry the same information as its more useful cousin, and it would even carry that information in extractable form. But this representation would lack the syntactic salience requisite for immediacy.

Extractability and syntactic salience do not suffice for immediacy. In addition, it must be easy to learn with which properties the perceptually salient properties of the representation correlate. That is, for pieces of information to be readily available, there must be a plan of correlation between features of the representation and features of the data that is easy to grasp. This is *semantic salience*. Without semantic salience, interpreting a representation will be difficult, defeating its purpose. As with syntactic salience, semantic salience will depend on the training, backgrounds, and innate perceptual and cognitive capacities of the consumers of representations.

Imagine we have five shades of red of varying brightness that we wish to stand for five temperature ranges in a graph. The obvious way to do this is to make the temperature ranges from lowest to highest correspond to the colors from darkest to brightest, or conversely. We could make the middle color correspond to the highest temperature range, and make the brightest color

correspond to a temperature range in the middle, but that would make the graph impressively difficult to *interpret* regardless of how syntactically salient those colors are. In this case the isomorphism between temperatures and the relation of being *greater than* on the one hand, and colors and the relation of being *brighter than* on the other, is salient. We easily interpret such representations, just as we easily interpret a thermometer, for which the heights of a column of mercury and the relation *being taller than* are isomorphic to the temperatures and the relation of being *greater than*. As we will see in the next section, isomorphism is, as many have suspected, important for understanding images, but it is far from the whole story. Isomorphism is important because it contributes to immediacy, and, as the next section argues, it is important for understanding the immediacy of information across levels of abstraction. Without leaning heavily on research at this point, there is little more to say about what makes some plans of correlation salient while others are not. (See, e.g., Gattis 2001, 2002 for some empirical work on this topic.) It is easy to come up with examples, and for now that is what matters.

Immediacy is extractability, syntactic salience, and semantic salience. A bit of information is extractable or it is not, and this depends on how the non-semantic features of a representation are responsible for determining the semantic features of a representation. Syntactic salience and semantic salience are matters of degree, and they depend on the consumers of representations, their training, backgrounds, and innate perceptual capacities. It makes sense to talk about degrees of immediacy, even though not all of its components can be characterized that way. Immediacy clearly relates to, and may even capture what it is to represent something explicitly as opposed to implicitly. It is far from clear, however, that there is *one* distinction covered by the uses of 'implicit' and 'explicit', and that topic is not this paper's focus. (Cf., for example, Cummins 1983, Dretske 1988, Kirsh 1991, and Clark 1992.)

Immediacy does not suffice to show why images and graphs seem to be such special kinds of representation, however, since descriptions and lists of numerals make information immediately available as well. It turns out that what is distinctive of images and graphs is that they present information immediately *across many levels of abstraction*. Edward Tufte's work on representation is explicitly concerned with syntactic and semantic salience, and at least implicitly with extractability. In discussing data maps, he claims "Only a picture can carry such a volume of data in such a small space.

Furthermore, all that data, thanks to the graphic, can be thought about in many different ways *at many different levels of analysis....*" (1983, 16 italics mine) The next section unpacks this valuable idea.

3. *Immediacy across levels of abstraction*

Let's say we have a data set that tells us the temperature along a 2D surface. One way of presenting the data is to select an origin point and then list triples of numerals: two for coordinates on the surface and one for the temperature at that location (x, y, T). The numerals in the list present data at a *specific* level of abstraction determined by the precision of the temperature and location measurements. This information is *extractable* from the representation: we can find features of the representation, say the triple (4, 3, 45), that specify a location and temperature and nothing else.

1, 1, 06	3, 1, 33
1, 2, 22	3, 2, 39
1, 3, 27	3, 3, 45
1, 4, 33	3, 4, 54
2, 1, 27	4, 1, 37
2, 2, 24	4, 2, 46
2, 3, 35	4, 3, 45
2, 4, 49	4, 4, 65

Often the information of interest is at some remove from the determinate values of temperature at specific locations, however. We need to abstract from the data presented by the list to the desired level of detail, but would this amount to *extracting* the information? That is, if we do not care so much that the temperature is 45, but are only interested in whether it is between 40 and 59 degrees, can we *extract* that abstract bit of information from the list of numerals? You might think not. The relevant features of the list are the shapes of the numerals that constitute it and their relations to one another. It's easy to find such features responsible for carrying information about a specific temperature at a specific location—e.g., (4, 3, 45)—but it is tricky at best to find some feature responsible for carrying the information that the temperature is between 40 and 59 degrees *and nothing more specific than that*. It seems as though the parts of the list carry *more* information than we want. So, one can certainly get at the more abstract information by first decoding the list and making an inference from the determinate content—that the temperature at (4, 3) is 45 degrees—to the more

abstract information that one needs—that the temperature there is between 40 and 59 degrees—but this does not make the latter information immediately available. In a sense, the list stands *between* one and the data of interest, as any premise stands between one and the conclusion of an inference.

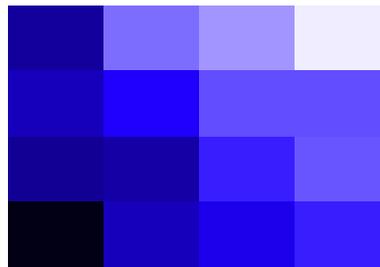
There is, however, an odd sense in which such abstract information is genuinely *extractable* from the list of numerals, appearances to the contrary notwithstanding. It is always possible to abstract over *numeral types*. This amounts to being sensitive to some abstract shape property that includes all of the shapes of the numerals that stand for values in the range of interest. So, if the list says that the value is 45 degrees at a certain point, but you are interested in values between 40 and 59 and nothing more specific, you can extract that info if you are sensitive to the abstract shape property '*40'-or-'41'-or-'42'-or-...-or-'59'*'. The problem with the list above vis-à-vis the info that we want is not extractability, it is syntactic salience. Abstractions over numeral types are generally not at all syntactically *salient*, so for this reason when we want to get at the abstract data of interest, we decode the list and then make an inference from the determinate data. While the list presents some information immediately, it does not present information across levels of abstraction immediately.³

We can improve matters somewhat by making the spatial relations between values more immediate in that it is easier to extract information about relative locations. This makes features of the representation that are responsible for coding relative locations isomorphic to the locations that they represent in a semantically and syntactically salient manner:

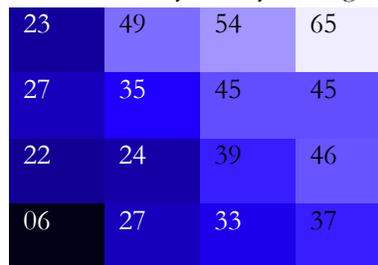
³ An anonymous reviewer helpfully pointed out that sometimes the information of interest might very well be available via a syntactically salient feature of the list. If the temperature values of interest just happen to correspond, for example, to temperatures between 50 and 59, then it is easy to look for all of the temperatures represented by two-digit numerals that have '5' in their first position. Some abstractions from determinate data are syntactically salient, even in a list, though it is unclear that these typically correspond to features of interest to us. If we were to use base-20 numerals, different abstractions over the determinate information would be rendered syntactically salient. Similar remarks apply to spatial locations represented by pairs of numerals, and anything else one might be able to put into a list. Consider, for example, how lists of numerals in base 10 make some information about primeness easy to extract. If a numeral ends in '0', '2', '4', '5', '6', or '8', then it is not prime.

23	49	54	65
27	35	45	45
22	24	39	46
06	27	33	37

Things get even easier if we consider a 2D image of the temperature along that surface. This is just another way of presenting the same data. The darker a region of the graph, the colder the corresponding region of the represented surface is:



This image carries all of the info that the list and the grid carry, and no more information, but it is much easier to abstract from the image's detail. With the image, abstractions over the data are extractable and syntactically salient. The plan of correlation between colors and temperature is semantically salient as well. It could be reversed, with darker regions standing for warmer temperatures, but it works either way. A haphazard correlation between colors and temperatures would not work, however. This image makes it comparatively difficult to extract information about specific temperatures, however. We could remedy this by leaving the numerals on the image:

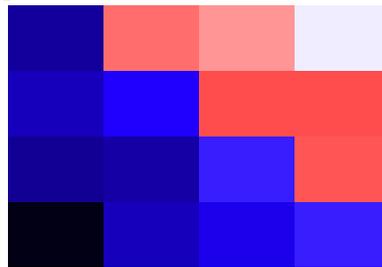


Without doing that, the image is still rather useful, especially if one is less concerned with the specific values of temperature and more concerned with relations between them.

The image is *much* better for drawing abstractions over the data than the list. For example, one could scour the list to figure out that the extremes of temperature occur in the bottom left and top right corners, and that the bottom left is cold while the top right is warm. It is much easier to figure

this out using the image because the regions of the image corresponding to the corners of the surface represented are clearly the darkest and lightest on the image. It is obvious that the bottom left is darker than the top right. There is no need to decode the image or a part thereof in all of its specific detail before abstracting to these more general claims. Abstracting over features of the graph itself and *then* decoding it gets you to abstract features of the data. This explains Tufte's point that we can get at the data in an image "at many different levels of analysis..." (1983, 16).

It's easy enough to say that this results from there being an isomorphism between regions of the image and their features and regions of the represented surface and their features, namely temperatures. Many agree—e.g. Barwise and Etchemendy 1995, Gurr et al. 1998, Stenning 2002—that isomorphism, or the more general notion of homomorphism, distinguishes graphs and images from other kinds of representation. Without being inaccurate, this covering term often used to describe images and graphs misses what makes *this* kind of isomorphism so interesting, vis-à-vis the goals of those who need to use the representation. Isomorphisms are multifarious and ubiquitous, which means that they are in and of themselves unhelpful in addressing either the first or the second problem of scientific representation. (Goodman 1976, Suárez 2003) The graph exhibits isomorphisms *across* many levels of abstraction from the determinate data points. Moreover, these particular isomorphisms are syntactically and semantically salient. So, regardless of how abstract or specific one's information needs are, those bits of information are extractable from the graph of temperature along a surface in a syntactically and semantically salient manner. If we want the particular temperature range between 40 and 59 degrees to be even more salient, we can code those temperatures with a different hue than the rest. Then they stand out with respect to the rest and relative to one another:



The foregoing puts us in a position to draw a general lesson about the contrast between graphs and images on the one hand and descriptions and lists of numerals, on the other. With lists, numerals, and descriptions, the rule is: *decode first, ask questions later*. Only once we have figured out the specific content of the list can we abstract from its details to something we are interested in. As a result, the list itself is of little if any *help* in getting from the specific data to our more abstract goals. It dumps its determinate content onto us and we are left to sort out the mess. Graphs and images endure quite a bit of interrogation *before* they need to be decoded. For this reason, such representations can help us get from the most determinate details represented to where we want to be. Reasoning *with* the graph—drawing abstractions over its features—allows us to draw conclusions about the graph's content. That's not to say lists are never more useful for forming some abstractions than images are. Note 2 above considers some cases in which the lists might be preferred. By and large, however, images have the advantage of presenting much information across many levels of abstraction in an immediate manner. Images and graphs are tools for discovery and diagnosis, interestingly enough, because they present a wealth of information in such a way as to allow us to ignore what simply does not matter. Descriptions are not as helpful in this way, and they are thus best suited to stating the conclusions we draw rather than presenting the data on the basis of which we draw them. The next section looks at how we can use immediacy better to understand how we use different ways of representing data.⁴

4. Floors and Ceilings

Any representation of temperature picks out some temperature at some level of detail. Excellent thermometers and detailed graphs represent temperatures accurately to hundredths or thousandths of degrees. Others are content to put us in the ballpark of the nearest degree, and some are even more coarse-grained than that. So, any representation has a determinate *floor*, characterized by the most specific piece of information about some determinate, such as temperature, that it carries. One can

⁴ I work out some closely related ideas in the context of *mental, perceptual* representations in my 2007. The relation between artifacts like images and perceptual states in terms of representing qualities across levels of abstraction constitutes one of the most interesting senses in which perceptual states are image-like.

always abstract from the determinate floor, of course, as when one reads the thermometer that says 79 degrees Fahrenheit just to get at the information that it's warm outside. Just as representations have determinate floors, they also have abstract ceilings, beyond which information is not *immediately* available. As immediacy is a matter of degree, these ceilings are somewhat less well-defined than the floors. The idea is that one can access abstractions above the ceiling, but doing so amounts more to a "decode first, abstract later" process rather than something more immediate. For example, a digital thermometer that reads '79' makes information that the temperature is in the 70s Fahrenheit immediately available, because it is relatively easy to focus merely on the numeral in the tens place, viz. '7', to the exclusion of the rightmost numeral '9'. The information that temperature is in the 50s-or-60s-or-70s is less immediately available since an abstraction over those three numeral types in the tens places is less syntactically salient than an abstraction over the '7's alone. One can compare representations with respect to the distances between their floors and ceilings and the number of salient steps between them.

While lists of numerals make information about their determinate floors and a number of steps above them immediately available, information of interest is usually obtained via an inference after the list has been decoded. By contrast, the 2D image of temperature has the same floor as the list but makes information at many levels of abstraction immediately available without the need for an inference from determinate content. The image has a richer set of steps between its floor and ceiling than the list does. One can tailor a representation for certain purposes and certain audiences by manipulating the salient steps between its floor and ceiling. For some representations, abstractions up to a certain point are perceptually salient or easily learned while beyond that point they are not. Sometimes abstractions are salient for just about any consumer of the representation, but sometimes they are only salient for a select few with the requisite experience using such representations. The upper limit on what is immediately represented can be set by syntactic or semantic salience, as well as by extractability.

For example, one can color the numerals in a list, chart, or graph depending on the abstract ranges of values they represent. All of the numerals representing temperatures below 35 degrees can

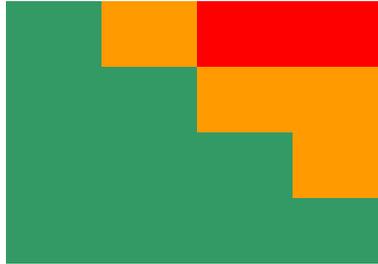
be colored green, those between 40 and 49 orange, and those above 50 red. For example, one could modify the graph from above as follows:

23	49	54	65
27	35	45	45
22	24	39	46
06	27	33	37

This raises the list's ceiling by making abstract pieces of information extractable in a salient manner. The numerals themselves do not make these steps salient because there is no syntactically salient abstraction over the numerals' features corresponding to the temperature ranges of interest. But notice that this modification raises the ceiling without making many intermediate steps up from the floor immediately available. The modified representation makes its determinate floor immediately available as well as the abstract comparisons between large ranges of temperatures, such as between temps in the 30s and those in the 40s. It does not make comparisons within two-degree ranges immediately available, however, even though those are closer to the determinate floor than the more abstract comparisons just mentioned. This is a good thing, since often we care about, for example, the determinate values and only certain coarse grained comparisons between them. Imagine a speedometer that reads out speeds in green numerals when below or at the speed limit and in red numerals when above it. Comparisons that are in between the floor and rather high ceiling of such a representational system are not immediately available. By contrast, the image of temperature discussed above in which a different shade of blue stands for each of the determinate temperatures renders comparisons at many of the abstraction steps from floor to ceiling immediately available. That image does not make some coarse-grained comparisons--such as between the 30s and the 40s--more salient than others, however, in the way that the image that uses shades of red and blue does.

Sometimes we don't care about the determinate floor at all. Perhaps one does not care how fast one is going, beyond knowing whether it is above or below the speed limit. Similarly, one may not care about precise oil pressure aside from whether it is dangerously low, just as sometimes one only cares whether a temperature is in the "red zone" or not. In these cases, one raises the *floor* of a

representation by eliminating the determinate information that it carries in favor of the abstract stuff one cares about. Doing so for the image we have been working with would yield something like this:



When one raises the floor, one changes the information that a representation carries, eliminating much of it. The representation no longer tells us the specific temperatures, opting instead to indicate general ranges into which those temperatures fit. These ranges are not of uniform size: they correspond to the ranges that interest us. When one raises the ceiling or highlights some important steps between it and the floor, by contrast, one leaves the information carried alone and merely makes abstract bits of it immediately available that were not so available earlier. When we think of whether a representation is useful, misleading, both, or something in between, it is often a consideration of its floor, ceiling, and the steps in between that helps to figure it out. Similarly, these considerations are precisely what we need if we want to *make* representations more or less useful for certain purposes.

5. *Images' epistemic weight*

We perceive all representations—linguistic, imagistic, or otherwise—but the foregoing suggests that images present their contents to us in a way that mimics the way in which we perceptually acquire information more generally. We do not simply *perceive* the contents of images, since they are often the kinds of things that are imperceptible, at least by visual means. But the images are visible, and the way in which their perceptible features relate to features of their contents renders our contact with those contents much like our perceptual contact with the image itself.

More specifically, the features of such representations that are responsible for them having the contents that they do are perceptible, and in immediate representations those features are salient. The way in which such features correlate with features of the representation's content is also easily grasped:

the representation is semantically salient. And finally, abstractions over the syntactic features of the representation stand for abstractions over the determinate content of the representation. That is, information is immediately available across levels of abstraction. By noticing that one region of the image is lighter than another--and thus abstracting away from the determinate brightness of each region and from the specific degree to which the regions differ from one another--one isolates a feature of the image that carries the information that one represented region is cooler than the other. When it comes to images and related representations, as opposed to linguistic representations, many more perceptual abstractions over the syntactic features of the representation saliently relate to abstractions over features of the representation's content.

Perception gives us access to many properties in our environment, at many levels of abstraction. We are not only able to focus on the determinate, dark green color of a tree, but we can also see it as green, or merely as dark. In fact, it seems as though our access to such abstract properties, like being green as opposed to being any particular shade of green, is just as direct as our access to more determinate properties. Similar facts hold for our perception of spatial properties, audible properties, and so on. There is no need to offer an account of direct perception here and it can even remain an open question whether perception really reveals both determinate and abstract properties of things in an equally direct manner. Perhaps, for example, it is only by a fast and often unnoticed inference that we come to know about greenness perceptually, as opposed to some determinate shade of green. For now what matters is that we *seem* to have ready perceptual access to properties across many levels of abstraction. Depending on our interests at the time, different bits of information about our environment are important or worth ignoring. Our perceptual systems give us access to a lot of information in a manner that allows its selective use.

The foregoing begins to explain the special epistemic place that pictures, images, and graphs can hold in relation to other kinds of representation. It is not just that such representations are particularly useful for certain purposes, but that this is so because interpreting them mimics the way in which we glean information about our environment perceptually. A searching, perceptual investigation of an image can straightforwardly be translated into a searching, perceptual investigation of its content.

Kendall Walton (1990) explains the differences between kinds of representation in terms of the make-believe interaction that they support. He thinks that we are able to make-believe that in perceiving pictures we perceive their contents, and his evidence is in part that the way in which we acquire information about those contents mimics the way in which we would acquire information about them perceptually. (Walton 1990, 305-9) It seems as though immediacy across levels of abstraction can figure in an explanation of why such representations support the kinds of make-believe that Walton thinks they do. This provides a clearer sense of the difference between representations of the decode-first-ask-questions-later variety--such as descriptions and lists--and images. We can reason *with* images rather than just decoding them and reasoning about their contents, in the sense that perceptual abstractions from their determinate details can lead us to conclusions about their contents. Interrogating the images themselves yields insights regarding their contents.

Images can make us feel rather reliably and intimately connected with their contents because they rely on the perceptual resources that we have on hand for investigating the world at large. These resources are tried and true, as far as most of us are concerned. This does not mean, in general, that we will regard the *contents* of images as particularly reliable, but rather that we will regard ourselves as being reliably and intimately in touch with those contents, regardless of their accuracy. We have reason to trust our *grasp* of the content, obtained as it was through perception-like means. Our grasp of abstractions from the content of a list, by contrast, even if that content is the same as the determinate information carried by an image, is not perceptually acquired in the way that our grasp of such abstractions is when we look at an image.

The foregoing discussion is far from complete. The myriad uses to which representations are put resists a compact, general treatment. The point of this essay is to suggest that thinking of representations in terms of their determinate floors, abstract ceilings, and the salient steps between them can aid the discussion of why certain choices for presenting data are and indeed should be favored, or shunned.⁵

⁵ Versions of this paper were presented in 2006 at Michael Dietrich's Leslie Humanities Seminar *Pedagogy in the Life Sciences* at Dartmouth, the Pacific APA in San Francisco, the *Beyond Mimesis and Nominalism* conference at LSE, and the CNRS Summer School on scientific images in Roscoff, Brittany. I thank the participants for their

References

- Bailer-Jones, D. 2003. When scientific models represent. *International Studies in Philosophy of Science* 17: 59-74.
- Bartels, A. 2006. Defending the structural concept of representation. *Theoria* 21: 7-19.
- Barwise, J. and Etchemendy, J. 1995. Heterogeneous logic. In J. Glasgow, Narayan, N., and Chandrasekaran, B. (eds.). *Diagrammatic Reasoning*. Menlo Park, CA: AAAI/MIT Press.
- Block, N. 1983. The photographic fallacy in the debate about mental imagery. *Noûs* 17(4): 651-61.
- Callendar, C. and Cohen, J. 2006. There is no problem of scientific representation. *Theoria* 21: 67-85.
- Clark, A. 1992. The presence of a symbol. *Connection Science* 4: 193-205.
- Cohen, J. and Meskin, A. 2004. On the epistemic value of photographs. *Journal of Aesthetics and Art Criticism* 62(2): 197-210.
- Cummins, R. 1983. *The Nature of Psychological Explanation*. Cambridge: MIT Press.
- Dennett, D. 1969. *Content and Consciousness*. London: Routledge.
- Dretske, F. 1988. *Explaining Behavior*. Cambridge, MA: MIT Press.
- Files, C. 1996. Goodman's rejection of resemblance. *British Journal of Aesthetics* 36: 398-402.
- French, S. 2003. A model-theoretic account of representation (or, I don't know much about art...but I know it involves isomorphism). *Philosophy of Science* 70 (Proceedings): 1472-83.
- Gattis, M. 2001. Mapping conceptual and spatial schemas. In Gattis, M. (ed.) *Spatial Schemas and Abstract Thought*. Cambridge, MA: MIT Press.
- , 2002. Structure mapping in spatial reasoning. *Cognitive Development* 17: 1157-83
- Goodman, N. 1976. *Languages of Art*, second edition. Indianapolis: Hackett.
- Gurr, C., Lee, J. and Stenning, K. 1998. Theories of diagrammatic reasoning: distinguishing component problems. *Minds and Machines* 8(4): 533-557.
- Kirsh, D. 1991. When is information explicitly represented? In Hanson, P. (ed.) *Information, Language, and Cognition*. Vancouver: University of British Columbia Press.
- Kitcher, P. and Varzi, A. 2000. Some pictures are worth 2-to-the-aleph-null sentences. *Philosophy* 75: 377-81.

helpful comments. I also thank Adina Roskies for comments on an early draft of this paper. Work for this paper was supported by a grant from the Leslie Center for the Humanities at Dartmouth College.

- Kulvicki, J. 2003. Image structure. *Journal of Aesthetics and Art Criticism* 61(4): 323-40.
- , 2007. Perceptual content is vertically articulate. *American Philosophical Quarterly* 44(4): 357-369.
- Larkin, J. and Simon, H. 1987. Why a diagram is (sometimes) worth 10,000 words. *Cognitive Science* 11: 65-99.
- Levesque, H. 1988. Logic and the complexity of reasoning. *Journal of Symbolic Logic* 17: 355-389.
- Neander, K. 1987. Pictorial representation: a matter of resemblance. *British Journal of Aesthetics* 27: 213-226.
- Perini, L. 2004. Convention, resemblance, and isomorphism: understanding scientific visual representations. In G. Malcolm (ed). *Multidisciplinary Approaches to Visual Representations and Interpretations*, Elsevier.
- Stenning, K. 2002. *Seeing Reason: Image and Language in Learning to Think*. Oxford: Oxford University Press.
- Suárez, M. 2003. Scientific representation: against similarity and isomorphism. *International Studies in the Philosophy of Science* 17: 225-44.
- , 2004. An inferential conception of scientific representation. *Philosophy of Science* 71 (Proc.): 767-79.
- Tufte, E. 1983. *The Visual Display of Quantitative Information*. Cheshire, CT: Graphics Press.
- Tye, M. 1991. *The Imagery Debate*. Cambridge, MA: MIT Press.
- Walton, K. 1990. *Mimesis as Make-Believe*. Cambridge, MA: Harvard University Press.