

Dartmouth College

Dartmouth Digital Commons

Dartmouth Scholarship

Faculty Work

9-2009

Knowing with Images: Medium and Message*

John Kulvicki

Dartmouth College

Follow this and additional works at: <https://digitalcommons.dartmouth.edu/facoa>



Part of the [Philosophy of Science Commons](#)

Dartmouth Digital Commons Citation

Kulvicki, John, "Knowing with Images: Medium and Message*" (2009). *Dartmouth Scholarship*. 1904.
<https://digitalcommons.dartmouth.edu/facoa/1904>

This Article is brought to you for free and open access by the Faculty Work at Dartmouth Digital Commons. It has been accepted for inclusion in Dartmouth Scholarship by an authorized administrator of Dartmouth Digital Commons. For more information, please contact dartmouthdigitalcommons@groups.dartmouth.edu.

Knowing with images: Medium and message¹

John Kulvicki, Dartmouth College

Draft May 19, 2006

Please do not quote or cite without permission.

There is something special about the way in which images, graphs, and the like present their contents to us. It is easy enough to say that in some sense our access to the information carried by such representations is rather direct or immediate, but rather difficult to unpack just what this special kind of directness or immediacy consists in. The fact that images seem so special is the source of the privileged place they hold as a means of presenting data as well as of the suspicion with which many regard them. What are the epistemic advantages and disadvantages of using such representations? And along what dimensions relevant to these issues do different kinds of representation--images, graphs, descriptions, lists, etc.--relate to one another?

The overall project of which this paper is a part is both descriptive and normative. How do representations actually work, and what makes some better than others, given a set of goals? Scientists have many goals in presenting data, from gaining understanding of a phenomenon to convincing others that they are correct. This paper does not sort out those goals and discuss the uses of images with respect to them in detail. Rather, it presents some dimensions along which representations differ from one another that seem particularly useful in addressing how scientists should and should not use 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 the same thing as explicitness, though the two are related, and immediacy is not in and of itself *distinctive* of images and the like. Almost any representation makes some aspects of its content immediately available. What distinguishes images and the like, as section three explains, is that they make information across many levels of abstraction immediately

¹ Thanks to Adina Roskies for extensive, helpful comments on this paper and in particular for helping to avoid a serious misstep in section four. Remaining problems are my problems.

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 makes a speculative claim about just 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 how we use representations for presenting data. Descriptions are good at providing coarse-grained, qualitative and quantitative information about a data set. They are good at this because they can present information at arbitrary levels of abstraction. A description is not bound 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 some 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.

Graphs and images, by contrast, are good at providing fine-grained, quantitative information about a data set. They cannot present information at arbitrary levels of abstraction: they cannot leave arbitrarily much unsaid. They can, however, present vast amounts of information about a great many features at, minimally, some levels of abstraction. The weather report's image tells you just how precipitation intensity changes across the Midwest 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 those who are properly trained very much about the different kinds of tissue in a body.² For now, the plausible story is that images and the like must 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.

As promised, this plausible story has some helpful flaws that can lead us to a better understanding of why images are valuable tools for

² Perhaps diagrams occupy some middle ground between images on the one hand and descriptions on the other, but the terminology in this area is nowhere near regimented.

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* shade of red 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. In fact, a list of numerals, much more like a description than a graph or image, is preferable to the latter if these 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 Upper Valley. The specifics are often beside the point when we use images. What makes images valuable is not that they carry many very specific bits of information but the fact that they allow us to get at precisely the abstract bits of information we want. We need to determine fine-grained, often quantitative information about something in order to *make* an image of it, but the image is not used *for* presenting 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. 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 in the process of diagnosing 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 a greater number of pieces of information. The next section claims that representations present information to us immediately when they satisfy three conditions. These conditions are a way of explicating the distinction between what Jill Larkin and Herbert Simon (1987, 67) call informational equivalence of representations and their computational equivalence. Representations can be alike in the information that they carry but differ in how they make that information available to their consumers. The topic of sections two and three is thus isolating the relevant dimensions along which representations can differ computationally while remaining informationally alike.³ Section four looks at how one can use these features of representations to explain why some are better for some tasks than others, and section five addresses why these features of images lend them a privileged epistemic place among representations, which can be used or abused.

2. Immediacy

First, a piece of information is immediately available in a representation only if it is *extractable* from that representation. Extractability just means just that 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

³ Larkin and Simon do not draw this distinction in the best way that they could, it seems. Computational differences, for them, include only the inferences that one is able to make easily based on the representation (1987, 67). As we will see below, one thing that distinguishes representations from one another is whether an *inference* is required at all in order to get at certain abstractions from a representation's determinate content.

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. 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 carry. 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, and nothing more specific than that.

Extractability is a feature of *any* kind representation and not just images. A list of numerals representing coordinates and determinate temperatures renders many specific pieces of information extractable. One can identify a given triple of numerals that indicates a temperature at a specific location and nothing else. If, however, we were to represent the entire data set with a name, 'Ralph', we would be unable to find features of the latter representation responsible for carrying information about a specific temperature at a specific location. This is true even though the name and the detailed list are about the same thing: they are informationally equivalent but not computationally equivalent, in Larkin and Simon's terminology. An inference is required from the representation 'Ralph' to an abstraction from its determinate content, while no inference is needed from the list because the relevant information is extractable. Necessary for a representation making a given piece of information *immediately* available is that it carries that piece of information in extractable form.⁴

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.

⁴⁴ For more on extractability, but in the context of perceptual, mental representations, see Kulvicki (2004, 2005). Extractability relates, albeit at some remove, to what Levesque (1988) calls "vivid" knowledge representation, but it would take us too far off course to unpack that relation here. Similarly, extractability relates to what is *explicitly* represented, as opposed to implicitly represented--see Cummins 1983, Dretske 1988, Kirsh 1991, and Clark 1992--but making that connection clear is beyond the scope of this paper.

The other two conditions, syntactic and semantic salience, relate consumers of representations to the information that those representations 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 now it suffices to point out that the properties in question are the kinds of things that your average, perhaps appropriately trained perceiver could be in a position to notice without much effort, and to give some examples of this.

So, 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 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. Many are unaware of what saturation is, in any case, and it is often confused with brightness. 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. (See, e.g., Gattis 2001, 2002)

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 dullest to brightest, or conversely. We could make the middle color correspond to the highest temperature range, and the brightest color correspond to a temperature range in the middle, but that would make the graph impressively difficult to *interpret*. It just so happens that 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 semantically 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, to the immediacy of information across levels of abstraction. Without leaning heavily on research, there is little more to say about what makes some plans of correlation salient while others are not, but it is easy to come up with examples on both sides just as we could with syntactic salience.

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.⁵

⁵ Edward Tufte's work on representations, especially his *Envisioning Information* (1990), and *The Visual Display of Quantitative Information* (1983), is rather explicitly concerned with syntactic and semantic salience, and implicitly concerned with extractability. In discussing some data maps, he claims: "Only a picture can carry such a volume of data in such a small space. Furthermore, all that data,

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*. The next section unpacks this 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 as a list of triples of numerals: two for coordinates on the surface and one for the temperature at that location. 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 (1, 1, 42) that specify a location and temperature.

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 42, but are only interested in whether it is between 39 and 42 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., (1, 1, 42)--but it is tricky at best to find some feature responsible for carrying the information that the temperature is between 39 and 42 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 in the top left corner is 42 degrees--to the more abstract information that one needs--that the temperature there is between 39 and 42 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.

thanks to the graphic, can be thought about in many different ways at many different levels of analysis...." (1983, 16) That last remark of Tufte's will be quite relevant in Section 3.

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 42 degrees at a certain point, but you are interested in values between 39 and 42 and nothing more specific, you can extract that info if you are sensitive to the abstract shape property '*39'-or-'40'-or-'41'-or-'42'*'. The problem with the list 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.

By contrast, 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 can carry all of the info that the list carries, 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 (not to mention semantically salient). For example, one could scour the list to figure out that region A is warmer than region B and that the difference between A and B is greater than the difference between B and C. It's easier to figure this out using the image because the region of the image corresponding to A is lighter than the region corresponding to B and the difference in lightness between A and B is greater than the difference between B and C. 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 is what Tufte was pointing out when he claimed that we can get at the data in an image "at many different levels of analysis..." (Tufte 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, marks off 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.⁶ The graph of temperature exhibits isomorphisms *across* many levels of abstraction from the determinate data points. Moreover, this particular isomorphism is 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 39 and 42 degrees to be even more salient, we can just code those temperatures with a different hue than we use to code the rest. Then they stand out with respect to the rest, and relative to one another as well.

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. 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 helpful in this manner, 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

⁶ Nelson Goodman (1976) famously pointed out that representations will in general *resemble* what they are about in indefinitely many ways, and he used this to argue against the claim that resemblances between representations and what they represent could in and of themselves do any interesting work in explaining representational kinds. Here I am making a similar point about isomorphisms, which are, if anything, more ubiquitous than bona fide resemblances. Absent an account of why certain kinds of isomorphism are interesting and relevant, it is unhelpful to point out that certain representations are isomorphic to what they are about. I make a similar point about perceptual, mental representations in (Kulvicki 2004).

use immediacy to better understand how we use different ways of representing data.

4. Floors, Ceilings, and Raising the Roof

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 that it carries. One can always abstract from the determinate floor, of course, as when one reads the thermometer that says 79 degrees 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. One can access abstractions above the ceiling, but doing so amounts to a "decode first, abstract later" process rather than something more immediate. Two dimensions along which we can compare representations, then, concern the distances between their floors and ceilings and the number of salient steps between them.

Lists of numerals make information about the determinate floor immediately available, but information more abstract than the floor is usually obtained via an inference from the decoded determinate information. These lists have very low ceilings, so most of what we do with them involves decoding first and then asking questions. By contrast, the 2D image of temperature makes information at many levels of abstraction immediately available without the need for an inference from determinate content. The image has a much higher ceiling than the list, and there are many ways to manipulate the number of salient steps between the image's floor and its ceiling. Exactly what one has to do in order to manipulate these features of representations will depend on the particular kind of representation one is using, and one's audience. 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 depending on the abstract ranges of values they represent. All of the numerals representing

temperatures between 40 and 49 degrees can be colored orange, those between 35 and 39 green and those between 50 and 55 red. This raises the list's roof (or ceiling) by making abstract pieces of information extractable in a salient manner. But notice that this modification raises the roof 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, imagine an image of temperature over a surface where a different shade of blue stands for each of the determinate temperatures. This renders comparisons at many of the abstraction steps from floor to ceiling immediately available, though it does not make some coarse-grained comparisons--such as between the 30s and the 40s--more salient than others.

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. When one raises the floor, one changes the information that a representation carries, eliminating much of it. When one raises the ceiling, 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* a representation more useful for certain purposes, or, of course if we want to make it more misleading. The next section sketches how these features of

representations affect the epistemic weight we take them to have and our choices concerning how to represent data.

5. Images' epistemic weight

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, 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 needs 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.

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. It's wrong, of course, to claim that we simply *perceive* the content of an image, since the contents of images are often the kinds of things that are imperceptible, at least by visual means. But the images themselves, like all representations, are perceptible, and the way in which their perceptible features relate to features of their contents renders our contact with that content 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 stand out. This is just what syntactic salience amounts to. 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, the result of abstracting over the syntactic features of the representation is a syntactic feature that stands for an abstraction 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. Perceptual abstractions over the syntactic features of the representation saliently relate to abstractions over features of the representation's content.

The foregoing helps 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 interpreting many of 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.⁷ 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. Asking questions of 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 that content, regardless of whether the content is accurate. 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

⁷ Cf. Kendall Walton (1990) who 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 our perceiving of pictures is perceiving their contents, and that this seems right because 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) See also my (2006, 239-44) for a discussion of Walton on this topic. It seems as though the structural feature of images discussed in this paper can figure in an explanation of why such representations support the kinds of make believe that Walton thinks they do.

carried by an image, is not perceptually acquired in the way that our grasp of such abstractions is when we look at an image.

While we feel in touch with an image's content, it might be quite difficult to isolate a scientist's important claim amidst the wealth of data an image can present. The point of scientific papers--the take-home message--is often much more abstract than the wealth of data that an image presents. In fMRI studies, for example, the point is often that there is greater activity in some areas of the brain while performing a given task than there is while performing a certain baseline task. So the particular shape of the region of activity, and the particulars of exactly how much greater than the baseline the activity happens to be often do not figure in the conclusion that the researchers want to draw. Such particulars matter even less in studies that average results for a few test subjects and present the results in one image. The particulars of the data are *data*, of course. They are worthy of presentation and oftentimes the best way to present such a wealth of data will be in an image. But as with many images the point is often not the particulars of that wealth of data but a conclusion considerably more abstract.

The abstract conclusions of scientific experiments, especially fMRI studies are rarely presented in imagistic form. They are written out, and the reader is expected to be able to see the conclusion in the more detailed image that represents the whole wealth of data. It is certainly possible to make an image that more closely matches the conclusion at its level of detail. All we need to do is raise the floor of the representation, so that we throw away information that is irrelevant to the conclusion at hand. In such a case, the most determinate claim that the image makes is in line with the conclusion that the scientists want to draw. Such an image could be presented next to the image that carries all of the information about the data, so that readers could see how the abstraction relates to the wealth of information that the scientists acquired.⁸

The problem with doing this, which might be a reason such abstract images are shunned, is that such images are apt to seem cartoonish. Rather than being characterized by subtle changes in color and brightness, such as characterize photographs as well as ordinary fMRI images, these

⁸ Michael Lynch (1988, esp. 157-60) discusses this in the context of biological illustrations of cells, in which a photograph of a cell is juxtaposed with a diagram thereof. The diagram raises the floor of the representation--although Lynch does not use this terminology--and thus makes it easier for viewers of the photograph to see the features of the cell in the photo. That is, the diagram makes the subset of information in the photo that is relevant stand out.

images would be composed of rather large regions of uniform hue. Such a representation would accurately capture the result that the scientists are interested in, and presenting it in such a way would aid in the audience's understanding of those results. As an aid to understanding, however, it might fail to seem particularly scientific, since cartoonishness is generally regarded as a mark against a representation. Cartoons are not sources of knowledge so much as sources of entertainment.

The fact that photos and other representations that carry vast amounts of information immediately, and thus have very low floors and rather high ceilings makes them quite appealing as media for representing data. When we raise the floors of such representations, which often brings them in line with the information that we wish to convey, we lose the appeal of the richly informative representation and wind up with something that can seem much less scientific because it seems diagrammatic and even cartoonish. This is not to say we should not make use of such representations, however. They can actually be just what we need, especially if the consumers of such representation, who in the case of fMRI are often the public at large, are unable readily to extract the relevant bits of information from the all-too-informative image.

The foregoing discussion is far from complete. The point was merely to suggest how knowledge of representations' floors, roofs, and the salient steps between them can aid the discussion of why certain choices for presenting data are favored or shunned and how certain favored or shunned practices might be put to better use.

Appendix: Raising the roof, and then some

The following extended example illustrates the ideas explicated above.
First consider a list:

(x-coordinate (1-4), y-coordinate (1-4), temperature (0-6))

1, 1, 0	3, 1, 3
1, 2, 2	3, 2, 3
1, 3, 2	3, 3, 4
1, 4, 3	3, 4, 5
2, 1, 2	4, 1, 3
2, 2, 2	4, 2, 4
2, 3, 3	4, 3, 4
2, 4, 4	4, 4, 6

The next step raises the roof a bit 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.


0	2	2	3
2	2	3	4
3	3	4	5
3	4	4	6

Raise the roof a bit more by color-coding groups of numerals based on the ranges to which they belong. This is a goal-directed modification of the graph in that each color stands for some range of temperatures, but the sizes of the ranges are far from uniform. Green covers 0-2 but yellow only covers 3 while red covers 4-5 and purple 6. There is a sense in which this coding is homomorphic to the temperatures, as long as one posits an appropriate similarity relation among the relevant colors, but we cannot call this an isomorphism because of the different temperature ranges covered by each color. A graph like this could tell us which regions of a surface are at a safe temperature (green), which are borderline (orange), which are dangerous (red) and which are, say, critically dangerous (purple). This

coloration does not change the information carried by the original graph, so much as make certain abstractions from the determinate data more syntactically and semantically salient. Thus, this contributes to immediacy of the data of interest.

0	2	2	3
2	2	3	4
3	3	4	5
3	4	4	6

The next move does not raise the roof so much as it raises the floor. Determinate information about temperature is discarded in favor of a colored matrix that picks out the temperature ranges of interest. Unlike raising the roof, this move changes the information carried by the representation by disposing of a lot of it. If the determinate temperatures are only relevant insofar as they fit into one of the four color categories below, however, there is no harm in raising the floor to make the relevant information more easily accessible: the irrelevant details cannot be a bother once you raise the floor. The downside of raising the floor is that the result can seem a bit cartoonish. To the extent that one wants to represent rather abstract bits of information using an imagistic medium, one flirts with making a representation that looks like a cartoon. Since cartoonish representations are not typically the kind of representation on which we rely for knowledge, especially when it is knowledge of a world that we know to be vastly complex, the risk with raising the floor is that one will not be taken seriously.



References

- 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.
- 1998. Computers, visualization, and the nature of reasoning. In Moor, J. and Bynum, T. (eds.) *The Digital Phoenix: How Computers are Changing Philosophy*. London: Blackwell: 93-116.
- Clark, A. The presence of a symbol. *Connection Science* 4: 193-205.
- Cummins, R. 1983. *The Nature of Psychological Explanation*. Cambridge: MIT Press.
- Dretske, F. 1988. *Explaining Behavior*. Cambridge, MA: MIT Press.
- 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.
- Kulvicki, J. 2004. Isomorphism in information-carrying systems. *Pacific Philosophical Quarterly* 85(4): 380-95.
- 2005. Perceptual Content, Information, and the Primary/Secondary Quality Distinction. *Philosophical Studies* 122(2): 103-132.
- 2006. *On Images: Their Structure and Content*. Oxford: Oxford University Press.
- 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.
- Lynch, M. 1988. The externalized retina. In Lynch, M. and Woolgar, S. (eds.) *Representation in Scientific Practice* Cambridge, MA: MIT Press.
- Stenning, K. 2002. *Seeing Reason: Image and Language in Learning to Think*. Oxford: Oxford University Press.
- Tufte, E. 1990. *Envisioning Information*. Cheshire, CT: Graphics Press.
- 1983. *The Visual Display of Quantitative Information*. Cheshire, CT: Graphics Press.
- Walton, K. 1973. Pictures and Make-believe. *Philosophical Review* 82(3): 283-319.
- 1990. *Mimesis as Make-Believe*. Cambridge, MA: Harvard University Press.