

# 6

## Bare Bones Content

Part I accounts for what makes a representation pictorial in terms of structural features of the systems to which representations belong. The main opponents there are perceptual theories that distinguish pictures from other representations by appeal to how they come to have the contents they have. For the perceptual theorist, what makes a representation a picture of a horse is that there is some special relation between the perception of horses and the perception of the representation. There is a third way to go, however. John Haugeland's (1991) account is unique in that it distinguishes pictures and images from other kinds of representation exclusively in terms of their contents without making reference to their structure or how we perceive them. For Haugeland, the bare bones contents of images have features that distinguish them from the bare bones contents of descriptions, and these distinguishing marks suffice for understanding what makes different kinds, or 'genera', of representation distinct from one another.

Haugeland's account raises a worry for the structural account offered in Part I. If he is right, and we can effectively distinguish images or pictures from other kinds of representations solely in terms of their contents, then structural or perceptual considerations may be interesting but beside the main point. Since bare bones content is an essential feature of the account on offer here, what real work is there left for the structural aspects of these representational systems to do? The structural account may just be a footnote to Haugeland's more fundamental considerations.

This chapter shows that Haugeland fails to identify features unique to the contents of images, and thus fails to account for what makes a

representation pictorial, imagistic, or what have you. Since bare bones content does not form the core of a successful account of pictorial or imagistic representation in Haugeland's hands, more needs to be said. The shortcomings of Haugeland's account can be traced to a failure to appeal to the structure of such representations as worked out in Part I. Furthermore, the specific way in which Haugeland's account fails points directly to transparency, and indirectly to sensitivity and repleteness, as features distinctive of pictures. The upshot is that Haugeland has introduced an important and unfortunately neglected tool for understanding the distinctions between images and other kinds of representations, but the true value of his contribution can only be realized by an appeal to the structure of images as well as their content.

It is also worth keeping in mind that Haugeland's aim is a somewhat coarser sorting of representational systems than that achieved in Part I. He is looking for the 'representational genera', not the species. The 'iconic' genus includes the images plus some graphs and diagrams, while the 'logical' genus includes descriptions, and other logico-linguistic kinds of representations.

## Bare Bones Content

Chapter 3 appeals to the bare bones content of pictures in order to explicate transparency, but it does not mention the bare bones contents of descriptions, diagrams, and graphs. For Haugeland to claim that features of bare bones content distinguish pictures from linguistic and other kinds of representations, he needs to give some account of their bare bones contents.

For example, the claim 'Johnny went to bed with a frown and without his supper' has a rather rich fleshed out content. Arguably one does not fully get that content without being able to draw some connection between the frown and the lack of a meal, without being able to guess that Johnny is likely a child or being portrayed as one, and so on. The sentence relies on the reader to draw out what is implied, implicated, and otherwise suggested by and alluded to in the remark without being built into the sentence itself. What one has to build upon, however, is the bare bones content of the sentence,

which amounts to nothing more than that someone, Johnny, went to bed, that he was frowning at the time, and that he did not have any supper.

One could ultimately challenge whether there really is anything like the bare bones content of logical representations as Haugeland construes it. If this is a problem, however, it is a problem for many more philosophers than Haugeland. The bare bones/fleshed out distinction builds on, for example, the distinction between what is implicit and what is explicit in a logical representation. (Haugeland 1991, 186).<sup>1</sup> While bare bones content closely tracks what is explicit in a representation, fleshed out content seems to include not just what is, strictly speaking, implicit, or implied by what is explicit. It also includes what is implicated, and ‘Only if “implication” is stretched to the full gamut of what can be expected from conversational skills, topical associations, critical judgments, affective responses, and so on can it hope to encompass [fleshed out content]’ (Haugeland 1991, 187).

By way of contrast, the bare bones/fleshed out distinction is distinct from the literal/figurative distinction. There is nothing figurative about what is implied by a representation, but fleshed out content includes what is implied. Haugeland (1991, 186) also notes that if something is figurative, then it is not literal, and conversely, while fleshed out contents include bare bones contents. Bare bones content is certainly literal, but so is fleshed out content. Whether there are bare bones to a representation’s figurative content is interesting, but will not be pursued here and Haugeland does not pursue it either.

The bare bones contents of pictures are by now more familiar than those of descriptions. Pictures are often taken to be ambiguous or at a minimum rather indeterminate in content. Those pictures in one’s family album certainly represent family members on different occasions and at different locations, but there is a sense in which they are indeterminate as to whether they represent one’s family or cleverly designed mannequins on cleverly designed sets. Thought of in this way, pictures are quite indeterminate in what they represent, and if one takes this to its extreme, ‘all the photos “strictly” represent is certain

<sup>1</sup> In what follows all page references are to the version of Haugeland’s (1991) that is reprinted in his (1998).

variations of incident light with respect to direction' (Haugeland 1991, 189). In terms of linear perspective, for example, these pictures pick out some information about color and illumination and spatial features that are invariant under projective transformations. This is no more and no less than what all of the different interpretations of a given picture in LP have in common, so it is a good candidate for a picture's bare bones content: that upon which any precisification of the content must build. This is a far cry from the content that we usually ascribe to pictures—dogs, trees, people, places, etc.—but ascribing these fleshed out contents to pictures is consistent with them having a rather abstract bare bones content. For Haugeland, 'The point is not to deny the obvious, but rather to distinguish, within the undeniable contents of everyday representations, a substructure that is skeletal—in the special sense that it does not draw upon the user's antecedent familiarity with the situation and the world' (Haugeland 1991, 189).

## Distinguishing Genera

Now that we have a more expansive sense of what bare bones content is, and how it might apply to descriptions as well as pictures, what are the relevant differences between the bare bones contents of images and those of logical representations? Haugeland's answer is deceptively simple: '[T]he primitive elements of logical contents... are always identifiable separately and individually. That is, they can enter into... contents one by one, without depending on their concrete relations to one another, if any' (Haugeland 1991, 191). Earlier we characterized Johnny as frowning, which was quite independent of any other characterization of him, including those concerning whether he had had his dinner. What is important here is that it is possible to identify the elements of the bare bones content of a logical representation wholly on their own terms, independently of how these elements may relate to the other elements of that content. The claim is not that there are no relations between these elements, or that there can be no relations between these elements. The point is just that what it is to be an element of the content of a logical representation does not essentially include any relations to the other elements of such content.

With images, the situation is quite different. '[I]conic contents might be conceived as variations of values along certain dimensions with respect to locations in certain other dimensions.' (Haugeland 1991, 192). The bare bones content of a color photograph involves the differing values of hue, saturation, and brightness along two spatial dimensions. Unlike logical representations, the elements of an image's bare bones content are not specifiable without reference to the other elements of that content. Each location is relative to the other locations depicted, and does not count as an element of the picture's content without being placed relative to those other locations. When we say of a picture's bare bones content that it includes something like 'red outline there', the immediate question to ask is 'where is there?'. Understanding where the red outline is depicted as being requires understanding where everything else depicted is depicted as being. Not only are all depicted locations relative to the other depicted locations, but each location is depicted in terms of color. What unites all of the regions in the picture is that each is specified in terms of a hue, saturation, and brightness. The 'elements [of an image's content] are always organized relative to one another in some regular structure that every representational token of the pertinent scheme essentially presupposes' (Haugeland 1991, 192). So, given that one element of the picture's content is some color at some relative location, we know that some other element of the picture's content is a color at a relative location.

It is quite difficult but absolutely essential to understand the foregoing as claims about the *contents* of representations, and not as claims about the representations themselves. Note, for example, that every sentence in English is such that its elements are 'organized relative to one another in some regular structure that every representational token of [English] essentially presupposes', since all languages have a grammar and syntax. This is beside the point. The elements of the contents of sentences of English are not like this, since one can specify the frown of Johnny quite independently of any other elements of the content of a sentence containing it. The rest of the sentence can discuss any kinds of features of just about any object one likes. There is no spatial, topical, or other kind of unity that the contents of descriptions all have in common, and this is what distinguishes descriptions

from images. One cannot specify an element of the content of a picture independently of its concrete relations to other elements of that content because every representation in the system presupposes a structure whereby all elements are related to one another. Every depicted location is assigned some hue, saturation, and brightness, and it is part of the nature of certain pictures that they do so. That is not to say that the color at a certain location places constraints on what colors appear at other locations, but just that if color-at-a-location is an element of a picture's content, then other elements of that content are also colors at locations relative to that one. Likewise, if temperature-at-a-time is an element of a graph's content then other elements of that content are also temperatures at times.<sup>2</sup>

## Problems with the Account

This is an interesting and subtle observation about the nature of the contents of images, graphs, and related representations, but does it really distinguish this class from others like descriptions? At first glance, and no less in light of Part I, it seems as if part of what makes a representation pictorial is the structure of the representations themselves and not just what they are about. These intuitions may be false, of course, but if so, one is obliged to explain them away. It is not as though the account in Part I fails to grate against intuitions. In line with Part I, Haugeland thinks that there is nothing essentially perceptual or even visual about pictorial representations. For all that Haugeland says, images could be auditory, tactile, or what have you. What makes Haugeland's claim quite radical is that, in addition, he thinks it is only in terms of content that we need to distinguish the iconic genus from other kinds of representations, while Part I appeals to structure and content. There are problems with the account as Haugeland works it out, however, which suggest that Part I is on the right track.

<sup>2</sup> It is worth pointing out, even if it is a bit of an aside, that Haugeland's (1991) account makes room for digital pictures, graphs, and diagrams much as the account of Part I does. There is nothing about this feature of pictures' contents that requires them to be in an analog medium, since dimensions of variation can be discrete as well as continuous.

According to Haugeland's account, one could create an imagistic representational system by assigning undeniably iconic contents to the syntactic types of any system at all. Chapter 1 distinguishes semantic features of representational systems from their syntactic features and points out that once a set of syntactic types is in place, one is quite free to assign any of a number of contents to them. This immediately renders Haugeland's way of identifying images a bit dubious. For example, the representational system in which the *contents* of color photographs are paired with the sentences of English is an imagistic system on Haugeland's account. The contents of color photographs are paradigmatically iconic, but we can assign what contents we like to a representation, including the sentences of English. So this kind of reinterpretation of English results in a pictorial, or at least iconic, representational system on Haugeland's account, which is an unfortunate result.<sup>3</sup>

What seems to make pictures pictorial is not just that they represent 'variations of values along certain dimensions with respect to locations in certain other dimensions', but that they do so in virtue of having non-semantic elements that are related to one another as variations of values along different dimensions. The picture does not just represent variations of hue, saturation, and brightness along two spatial dimensions, it also does so by *being* an instance of a two-dimensional surface at each point of which there is a certain value of hue, saturation, and brightness. Another example that Haugeland uses is a graph of temperature versus time. The graph has a bare bones content characteristic of an image. Each element picks out a time and assigns a temperature to it. But we loathe calling anything a graph of temperature versus time independently of how the properties of the representation manage to tie it to the content that it has. Such graphs usually use spatial dimensions to indicate the time and temperature, so the elements of the representation itself—its syntactically relevant properties—are related to one another in ways, perhaps isomorphic (see Chapter 4),

<sup>3</sup> Note that there are ways of reinterpreting English so as to make its representations pictorial. Just take each representation to be a depiction of an inscription of an indeterminate kind. So the sentence 'the cat is on the mat' picks out an inscription of the English sentence, which may be written in any of a number of styles, fonts, hands, and so on. This reinterpretation seems plausible, whereas Haugeland's does not. Both are odd.

to how the elements of the content are interrelated. But the foregoing are all structural facts concerning how parts of representations relate to one another and how syntactic types within a system are paired with semantic types. This is not to say that structural features of representations, completely independently of their semantic features, are sufficient to identify the pictorial. Part I ties pictures' structure to their bare bones contents. The point here is just that it cannot be exclusively in terms of content that images are distinguished from other representations.

A closely related problem for Haugeland's account is that it seems to allow rather ordinary logical representations, without *recherché* reinterpretations, to have the same contents as images. For example, a list of quintuples of numerals corresponding to two spatial coordinates and numerical values for hue, saturation, and brightness is not a pictorial representation, though in a sense it carries the same information as a pictorial representation. Likewise, we can construct a list of pairs of numerals—one for a moment of time, another for temperature—which would contain all of the information carried by the aforementioned graph. The list, however, is no image, graph or otherwise. This is a problem. Lists cannot have the same contents as images, at least according to Haugeland's account, since it is precisely in terms of content that such representations are supposed to be distinguishable from one another.

## Recording and Representing

Haugeland is not unaware of the second objection just raised, which he responds to by drawing a distinction between recording and representing. For Haugeland, the lists just mentioned are not representations with the same content as their corresponding icons, but instead are logical *recordings* of those icons (Haugeland 1991, 179). One can recover the image or graph from the list, because the list is an accurate representation of the token representation in question, not because the list has the same content as the original. Notice that this is plausible given what Haugeland has said so far. It is not the case, it seems, that the elements of the content of the list are essentially related to one another in some scheme that every list presupposes. But this does

seem true for the photograph and the graph of temperature. The distinction between recording and representing is supposed to capture the mistake behind thinking that a description or list can have the same kind of bare bones content as an image or graph.

Haugeland's defense of this claim is initially quite convincing. First, consider things the other way around. Instead of a list with the same content as an image, let's try to make an image with the same content as a description. One might think that a simple way of doing this is to make an image *of* that description. Photograph an inscription of 'The cat is on the mat' and you get an image with the content that a cat is on the mat. This is clearly though cleverly wrong: photographing an inscription that happens to be about a cat on a mat results in an image with an inscription as its content, and the image does not *eo ipso* inherit the content of the inscription. Such an image is what Haugeland calls a recording of a logical representation in an iconic medium. The recording does not share content with the representation recorded.

Likewise, Haugeland argues, given an image one could always generate a list of points and their corresponding values of hue, saturation, and brightness, but this would constitute a logical recording of an imagistic representation and the former would not share content with the latter. After all, 'explaining perceptual recognition—how, for instance, a system can look at a scene (or picture) and produce an articulate verbal description of what it sees—is a profound outstanding problem in psychology and artificial intelligence' (Haugeland 1991, 177). Simply listing the pixels in a picture could not have the rich content that we usually associate with images, so it would also be wrong in this case to suppose that logical and imagistic representations share content. One is simply a recording of the other. We even use such recordings to generate images all the time on our computer. We tend to think that the image is not stored in the computer's memory in imagistic form, but recorded there in logical form, ready to be unpacked on the screen in all of its pictorial glory.

There are three big and related difficulties with this analysis of the problem. First, recording is representing a representation, not representing just what the representation represents. Haugeland thinks the intuition that pictures and descriptions can have the same contents is rooted in mistaking a logical recording of an image for a

representation of what the image represents. This is unconvincing at best. The lists proposed above were not meant to contain descriptions of parts of images but rather descriptions of parts of the *contents* of images. The list for the graph of temperature versus time includes representations of time and temperature, not representations of regions of a graph, just as the list of regions of differing hue, saturation, and brightness is not about the image, so much as what the image is about. It may be that the photograph of the inscription ‘the cat on the mat’ does not have the same content as the description so inscribed, and it may even be that no picture has just the content that the cat is on the mat.<sup>4</sup> That is a far cry, however, from the claim that no description can have the same content as a picture or vice versa. It seems that Haugeland just acknowledges a difference between representing a representation and representing what the representation represents. He needs to do more to show that the best we can do to get a description to have the same content as an image is to record that image in a logical medium.

Haugeland has a reply to at least some of this. He is happy to admit that we can describe the content of an image—e.g. color and location of the depicted scene—rather than the image itself—color and location within the picture. In such a case ‘we must say that the resulting sentences describe the scene, since there is no image in play for them to describe, or to record logically.’ (Haugeland 1991, 179). Nevertheless, ‘even when the sentences are generated directly from the scene, what they are doing is logically recording an image-like representation—a virtual image, we could call it—... and it is this which the sentences describe, pixel by pixel.’ (Haugeland 1991, 179). So even if the descriptions pick out features of the scene depicted, they do so from a point of view, much like a picture would, and they thus record a virtual image of the scene. The description is nevertheless just a recording of that virtual image in a logical medium. The description of the scene captures only the varying intensity of color from a certain point of view within a certain solid angle. But, Haugeland points

<sup>4</sup> No doubt, this is because pictures tend to represent more detailed states of affairs than descriptions represent, as discussed at the end of Chapter 5. But this difference between pictures and descriptions cannot be analyzed in terms of the recording/representing distinction.

out, describing the variations of color within a solid angle is a far cry from describing the content of an image. To repeat a remark mentioned before, the descriptions in question ‘don’t describe the scene in anything like the way an articulate observer would describe it’ (Haugeland 1991, 179), and it is quite counterintuitive that such descriptions of (virtual) pixels have the same content as an image. Pictorial content is rather rich and complicated, and it most certainly is not captured by a mere list of locations and colors.

The foregoing leads to the second and third problems mentioned above. Second, even if we grant Haugeland this proposed solution, it does not apply to the list generated from the graph of temperature versus time. The descriptions in that list have temperatures and times as their contents. Neither the graph nor even a virtual graph of temperature versus time is plausibly taken to be a part of the content of those descriptions. Haugeland still has a problem, then, because it seems as if the list of descriptions and the graph itself can have the same content. This problem applies to a huge class of what Haugeland takes to be iconic representations. He could always exclude graphs on these grounds from the iconic genus, but this would cut strongly against the grain of his argument. In any case, he would have to revise his view substantially, because it seems as though the elements of the content of a graph are organized with respect to one another in precisely the way he takes to be characteristic of the bare bones contents of icons.

The third and most important problem is that his reply to the first problem misses the mark, even if it is true. He claims that the lists of locations and colors do not track what ‘an articulate observer’ of a scene or image would describe as an image’s content. This should make us doubt that the description and the image have the same content. But the insightful core of Haugeland’s account is that we are *not* supposed to be focusing on what articulate observers regularly take to be the contents of images. Imagistic representation, on Haugeland’s view, must be understood in terms of bare bones content and not fleshed out content. While the fleshed out content of an image may be a cat on a mat, its bare bones content is nothing more than ‘variations of values along certain dimensions with respect to locations in certain other dimensions’ (Haugeland 1991, 192). For Haugeland, it is only insofar as we consider images in terms of their bare bones content that

we can distinguish them from descriptions. But if we take images in this way, then the lists we have been considering seem as though they can have exactly the same contents as images. It is precisely variation of temperature with respect to time or color with respect to position that these lists describe: no more, but certainly no less. Haugeland is therefore wrong to treat, on the one hand, making a picture of an inscription and, on the other, describing a picture as analogous cases. A picture of an inscription does not inherit the content of the inscription, just as certain descriptions of pictures do not inherit their content. But a description of the bare bones content of a picture inherits that picture's bare bones content just as a picture of what an inscription describes can have the same content as the description.

Might Haugeland reply that the elements of the description's content nevertheless 'can enter into... contents one by one, without depending on their concrete relations to one another, if any' (Haugeland 1991, 191) while this is still not true of the elements of the picture's content? Perhaps, if Haugeland takes the lists in question to be mere parts of English, since in English it is not the case that the 'elements [of content] are always organized relative to one another in some regular structure that every representational token of the pertinent scheme essentially presupposes' (Haugeland 1991, 192). This reply does not rescue his position, however. On the one hand, the descriptions pick out relative locations, so in that sense the elements of their content can only be understood in relation to one another. On the other hand, one could easily construct a representational system whose purpose it is to do nothing but describe variations of color with respect to position. Every possible token in this system is a list that indicates position, hue, saturation, and brightness. The contents of representations in such a system would all essentially involve the 'variations of values along certain dimensions with respect to locations in certain other dimensions' (Haugeland 1991, 192). This system would not be iconic, despite the contents of its representations being what Haugeland takes to be characteristic of images. The representations in this system are sets of lists of numerals: paradigm examples of logical representations.

Perhaps Haugeland can claim that even in the specialized representational system just adumbrated, constituents of the content can be

understood independently of the other constituents. Just looking at the list, for example, tells one about color in a certain location without reference to other locations or the fact that the rest of the content is colors at locations. The rest of the list might say nothing about colors or locations. This too is implausible because, first, if the list is accurate it locates colors at relative locations, not locations in any interesting absolute sense. Second, one can come to know about elements of the picture's content by focusing on a portion of the picture, without taking in the whole. If this does not challenge the claim that a picture has a certain bare bones content, then the corresponding situation, in which one focuses on a portion of the list, doesn't challenge the claim for the list, either.

One more possibility is that Haugeland could claim that the lists in this special representational system are indeed images or iconic representations. It is counterintuitive, certainly, but philosophical accounts sometimes need to bite the bullet and be counterintuitive. The problem with this response is not that it is counterintuitive, however. The problem is that it undercuts the motivation for the view in the first place. The representing/recording distinction was introduced precisely to show why a list like the one we have been discussing does not have the same kind of bare bones content as an image. To the extent that this approach is successful, it allows Haugeland's account of representational genus in terms of bare bones content to track our intuitive sorting of representational systems into classes. It is fine to give up on the intuitive way in which we sort representational systems as long as one has an explanation of why our intuitions have led us astray, and why the counterintuitive sorting does some good theoretical work. Haugeland does not do this, however, and it is not clear how he would. This response seems more motivated by a desire to keep a flawed theory than a desire to account for what makes a representation iconic as opposed to logical.

At this point, we could be pulled in either of two directions: toward perceptual accounts or toward structural accounts. Haugeland has failed to explain imagistic representation solely in terms of content. If we were to add some convincing perceptual conditions, however, they may, along with Haugeland's discussion of bare bones content, yield a workable view of what makes a representation an image. For

example, Haugeland could claim that it is not just features of the bare bones content of an image that make a representation an image, but bare bones content plus the way in which a picture makes its content available for fleshing out. The reason the description of variations of color at relative locations is not an image is that we cannot flesh that content out to a dog in a field just by looking, while we can do just that with an image. Chapters 8 and 9 consider the role that bare bones content plays in fleshing out pictures' contents. But in order to understand that role, we need to understand the structure of such representations. Haugeland's claim goes wrong precisely where his account points to the structural conditions worked out in Part I, transparency in particular. That is, once you see that pictures are transparent, it is easy to see why Haugeland's deployment of the recording/representing distinction in defense of his claim that the bare bones contents of images are distinctive is bound to fail.

## The Link to Transparency

Why did Haugeland think that the recording/representing distinction could save his account of representational genera? This is an important question because, on the face of it, a recording represents a representation in another medium, not what that representation is about. If the problem is to figure out whether imagistic and descriptive representations can have the same content, the question we should be asking is whether a description of a depicted scene can have the same content as the picture. The question is not whether a description of a depiction can have the same content as the depiction. Furthermore, this approach immediately gets Haugeland into the uncomfortable position of calling descriptions of depicted scenes, as opposed to descriptions of pictures, descriptions of virtual pictures. It turns out that this approach seems plausible mainly because pictures are transparent. Ironically, precisely what makes Haugeland's proposal *prima facie* plausible is what renders it unworkable upon closer inspection.

Recall that in transparent systems representations of representations are of the same syntactic, and hence semantic, types as their objects. That is, a representation of another representation in a transparent system has the same bare bones content as the original. One

consequence of this is that pictures satisfy their own bare bones contents. Descriptions, graphs, and even many imagistic representations fall well short of meeting this condition. Transparent systems are also mimetic: pictures, among other representations, are similar to their objects in semantically significant ways. So whenever a picture's bare bones content includes a certain shape, the picture itself is that shape.

Notice, though, that if a picture is an instance of its own bare bones content, then there are many descriptions of the picture itself that wind up being accurate descriptions of its bare bones content. Describe the projective invariants and colors of a picture's surface and you not only describe its syntactically relevant properties, you describe features of its bare bones content. So this structural fact about pictorial systems makes it plausible that one could confuse a recording of a picture in a logical medium with something that represents just what the picture represents.

Only *some* descriptions of a picture's SRPs are descriptions of its bare bones content, though. Recall from Chapter 3 that the syntactic identity of a picture in, say, linear perspective depends on its projective invariants and not the determinate shapes of the regions of its surface. If one describes a picture's shapes and colors in all of their determinate glory, one is describing more than the bare bones content of the picture. The picture surface is square, but its bare bones content includes only a four-sided region of space. Since bare bones content is rather abstract, no object just satisfies bare bones content. So, pictures are instances of their bare bones contents, but they are *fleshed out* instances of it. There is more to the picture surface than the picture itself determinately represents. So describe the picture's SRPs in all of their determinate glory and you describe a fleshing out of that bare bones content.

Transparent systems are special for many reasons, and the present discussion brings out another one. We already know that within such a scheme if you represent another representation the result shares bare bones content with the original. The foregoing points out that if you *describe* a representation in such a scheme, the result can share either bare bones or even fleshed out content with the original. Contrast this with describing descriptions, in which case there is no reason to think that the result shares any content with its object.

In this sense pictures translate well into a descriptive medium even when it is the picture itself one describes. This is ultimately why the recording/representing distinction is of no avail in saving Haugeland's view from the objections made earlier.

Moreover, descriptions and other non-transparent representations do *not* translate well into pictorial media. Haugeland found it appealing to illustrate his point about recording versus representing by considering a photograph of an inscription. Even if the photograph captures all of the SRPs of the description so inscribed, the picture will not share content with the description. This asymmetry between depicting descriptions and describing depictions can be traced back to the transparency of pictures, and the fact that transparency entails mimesis. Pictures are instances of what they represent, so any accurate representation of their SRPs, in whatever medium you choose, will share fleshed out or bare bones content with the picture itself. Descriptions do not hold onto their content in such a way as to make them easily translatable into other media.

Transparency, and not just mimesis, is at the heart of this feature of pictures. Non-pictorial images—representations that are mimetic while failing to be transparent—fail to translate well into other media. Describing the weather radar does not yield a characterization of its bare bones content. The radar image is reddish over Chicago, but only yellowish over Flint. One could certainly describe the content of the image, in terms of intensity of precipitation, but access to that content does not come for free once one has access to the representation itself. This isn't really because one has to 'learn' the rules that assign colors to intensities in order to 'read' the representation correctly. We needn't assume that the describer of a picture has any sense of what it depicts or, more generally, any sense of how to read pictures in order for the description to capture the picture's content. It is because the picture is an instance of its bare bones content, not because it is easy to learn the scheme that pairs sets of SRPs of a picture with its content, that pictures translate well into other media.

Recall that the recording/representing distinction, though initially plausible as a defense of Haugeland's view of pictorial content, did not even get off of the ground when it came to graphs and diagrams. There is no reason to think that a description of the

graph of temperature versus time captures anything like a virtual graph of temperature versus time. Graphs encode information about temperature and time in a spatial medium. They may encode this information isomorphically, but they do so neither mimetically nor transparently, so one cannot describe a graph and *eo ipso* describe its content. Any accurate description of a graph—a recording of it—must describe the spatial dimensions of variation in the graph, not the temperature and time. And any description that picks out only spatial features of something, and not temperature and time, fails to share content with the graph. But this does not mean, as mentioned above, that there is no description that shares content with the graph. After all, one can always describe changes of temperature over time.

This discussion of graphs and non-pictorial images highlights the fact that the intuitive appeal of Haugeland's claims relies on the systems he is discussing being transparent. Only in transparent systems does it seem plausible in the first case that we are recording a representation rather than representing what the representation does. The irony here is that in these cases even if we are describing the representation, we are doing so in a way that does in fact capture its content. The resulting description therefore can have the same bare bones content as the representation it describes. Though Haugeland has given us an important way of describing pictures' contents, he has not given us a way of singling out pictures, or even what he calls 'icons', from other kinds of representational systems. Precisely where his account has problems, though, transparency steps in to fill the gap. In the end, Haugeland was not wrong to look to the contents of images and descriptions for an answer to how they differ from one another, but he was wrong to eschew reference to the structure of those systems as well. Transparency, as we have seen, is a feature of the structure of representations in a system vis-à-vis what they represent. Only proper attention to structure and content can account for what it is to be a picture.



Illustration 7 Rembrandt Harmensz van Rijn, *Adam and Eve*, 1638. Hood Museum of Art, Dartmouth College, Hanover, New Hampshire; gift of Jean K. Weil in memory of Adolph Weil, Jr., Class of 1935. Photography: Jeffrey Nintzel