8

IMAGES IN MIND

Many tools for explaining pictorial representation are of great interest to philosophers of mind, and progress in that area can readily inform the study of depiction, and maybe even the study of images more generally. To the extent that we better understand recognition, for example, we can more readily understand whether and how pictures evoke recognition responses. The deeper our understanding of perceptual experience, the clearer our sense of how experiences of pictures might be distinctive. These last two chapters look at things the other way around. How does the study of images inform the philosophy of mind? First, we consider the problems of mental imagery and perceptual content, then we move on to consider photographs and what they might tell us about object perception.

Are images *in* the mind? If images are essentially things that people make, then of course there are no images in the mind. But if images are a kind of structure-preserving representation, it's possible that the mind makes use of them. There are two phenomena that are usually taken to involve "mental imagery": imagining things in a richly perceptual way, and perceiving generally. When asked to imagine the house you grew up in, you call to mind a kind of visual state, which seems similar to a state of seeing the house. It makes sense to ask you from what perspective you are imagining the house—inside, outside, from the front or back—and to ask how vivid an imagining it is—can you see the color of the shingles? How many

windows on the first floor face the street? Some suggest that these episodes give excellent evidence for the presence of imagistic mental representation, in no small measure because it seems as though one can learn from them much what one could learn from perceiving the house directly. This suggestion generalizes. Perhaps the mind makes use of imagistic representations when perceiving the world generally. When you look out and see what's going on, you are doing so in part by forming image-like representations in your brain.

There are, of course, no paintings, drawings, sketches, graphs, or pie charts in the brain. In another sense, however, it is an interesting and important question whether and to what extent perceptual states are like images, as opposed to, for example, linguistic representations. Two dimensions along which one can ask this question concern the structure and content of images. First, are perceptual states structured, syntactically and semantically, like images are structured? Disputes along these lines fall under the rubric of the "Mental Imagery Debate." As with the discussion of scientific images, philosophers and psychologists have developed accounts of images independently of the work done in philosophy of art. We will consider those views and relate them to what we have already done. Second, do perceptual states have representational contents that are closer ken to images' contents than the contents of linguistic representations? This is often called the "Non-Conceptual Content Debate." The latter issue has been more popular over the last ten years, but it is a direct descendant of the imagery debates that were most active in the last few decades of the twentieth century. One focuses on structure, the other on content. Both of these issues are distinctive in that the study of artifacts can help with sorting them out, as we will see.

8.1 Artifacts and the mind

Artifacts are distinct from anything going on in the mind or brain, so we need to be cautious about extending work on artifacts to the mind. First, what makes something a representation in the first place is clearly distinct in the case of mental states and artifacts. Collaborations of intention and skill make artifacts represent things, whether they are pictures, images, graphs, or inscriptions. Mental states do not represent in virtue of the intentions of their makers. They have no makers, and artifacts must have makers. In fact, artifacts represent in large part because of the mental states of those who make them. Second, mental representations are not *consumed*

in the manner of artifacts. They are not read, like poems, or contemplated, like paintings. Mental representations are not among the things we perceive. Instead, they partly *constitute* perception of the world: perceiving is having the right kinds of mental representation.

Most accounts of artifact images do not generalize to mental representation. Mental images do not evoke experiences with distinctive phenomenologies because we do not, in any straightforward sense, experience our mental states at all. Mental states might *be* experiences, but they are not the objects of perceptual experiences. Similarly, we do not experience mental states as resembling things in outline shape, since we do not experience mental states as having outline shape at all. We might experience a surface as resembling the Empire State Building in outline shape, as Robert Hopkins suggests, but we do not experience our mental image as resembling anything in outline shape. Kendall Walton's view applies across sense modalities and across kinds of representation, but it cannot explain mental imagery. While mental states are sometimes props in games of make-believe (e.g. Walton 1978) they are not pictorial props. We do not imagine of our seeing of mental states that it is the seeing of other things, because we do not see our mental states.

Recognition theories might seem quite helpful for thinking about mental images. Recognizing something involves applying a concept to it: that's a giraffe! The recognition theory, recall, suggests that pictures evoke recognition responses for what they depict. The ability to recognize things in the wild also helps recognize them in pictures. Now notice that two people can have similar experiences, and differ in what they recognize, because they have different concepts. Perhaps both of us look out the window, and I recognize a cat-shaped animal while you recognize an ocelot. In a sense, we had similar experiences, even though our responses to the scene were different. Perhaps we can understand these *experiences* as the kinds of things that evoke certain recognition responses, and in that sense understand them on the model of images. The idea is that we distinguish experiential states based on the recognition responses they can evoke in suitably prepared perceivers. This would not be to show that these experiential states are images, however. They are not among the things we perceive, so much as things that partly constitute perception. Recognition is keyed to things in the environment, not things in the brain.[1]

None of this is meant to be an argument against these accounts of depiction, since it's hard to blame someone for trying to explain artifacts in terms of the responses they elicit from perceivers. From their perspective,

results from psychology and neuroscience might lend weight to their accounts of depiction, or undermine them, but they do not offer accounts of mental state types. It might be that some mental states have features in common with artifact images, and that would be interesting, but insufficient to justify the claim that there are images in the head.

The two kinds of view that have a chance of extending their reach to mental representation are resemblance views and structural views. It's possible that brain states systematically resemble their objects, and that this plays a central role in explaining what they represent. Similarly, it's possible that brain states have syntactic and semantic features that align them more with artifact images than with other kinds of artifact representations. In both cases, what is responsible for a brain state being a representation is going to be different from what makes charcoal on paper a representation, but that would not undermine the structural or resemblance claims.

The next three sections look at how debates over mental images interact with the structural and resemblance views of images, and then we move on to look at non-conceptual content. These issues within philosophy of mind have developed almost independently of the literature on artifacts, just as the literature on artifacts has developed largely independently of the philosophy of mind. The hope is to bring these strands of thought together in a mutually beneficial manner.

8.2 Mental mirroring I: spatial resemblance?

Every account of mental images suggests that some brain states are structure-preserving representations. They differ over whether such states resemble their objects, in limited respects, or whether they merely share a structure with them. Let's start with two related proposals about the nature of images, one from psychologist Stephen Kosslyn and the other from philosophers Elliot Sober and Jerry Fodor, and see how they zero in on a certain kind of resemblance. While they use the terms "depiction" and "picture," in the present context we should understand them as describing images, broadly speaking.

> In a depictive representation, each part of an object is represented by a pattern of points, and the spatial relations among those patterns in the functional space correspond to the relations among the parts themselves. Depictive representations convey meaning via their resemblance to an

> object, with parts of the representation corresponding to parts of the object ...
>
> (Kosslyn 1994: 5)[2]

The key ideas here are parts, correspondence, and the somewhat cryptic "spatial relations … in the functional space … ." We will consider the parts of Kosslyn's claim in turn, but before doing that it helps to look at a related claim, this one from philosophers interested in mental images:

> Every representational picture has representational subpictures; successive snippings never destroy representationality.
>
> (Sober 1976: 124)

> If P is a picture of X, then parts of P are pictures of parts of X.
>
> (Fodor 2008: 173)[3]

These claims also mention parts of a picture and the fact that they correspond, somehow, to parts of what they depict. With all of these claims, getting clear about the notion of *part* is key to what follows.

Kosslyn says, "each part of an object is represented by a pattern of points." This makes intuitive sense, but we should proceed carefully. First, his idea seems to be that an image, on the *whole*, is a pattern of points that represents objects. A photograph is a pattern of color across two spatial dimensions, for example. The points, in that case, are very small areas of the photo's surface, each of which has a certain color. Now notice that any spatial part of the photo's surface represents a spatial part of the scene it depicts. For example, you can snip, as Sober says, the photo into four parts, and each of them would be a photo of *part* of what the whole represents. This claim is not restricted to *contiguous* parts of the image. Any disjoint part of the image—i.e. two regions that are not connected by spatially adjoining regions—also represents a part of the depicted object. The upper left corner of the photo, along with the lower right corner, together represent a (disjoint) part of the scene, for example. This gives us something like the claim by Sober and Fodor: parts of a picture of X depict parts of X. The photographic examples are helpful, but it is very important not to take them too literally. What counts as a pattern of points might be very different in a brain than it is on the surface of a photo, and we will return to that issue shortly.

Next Kosslyn says "spatial relations among those patterns in the functional space correspond to the relations among the parts themselves." Ignore "in the functional space" for now, and keep the photo in mind. All of those spatial parts of the image relate spatially to one another. Those spatial relations correspond to spatial relations in the represented scene in that the former represent the latter. What's more, they correspond in a resembling or structure-preserving way. These representations "convey meaning via their resemblance to an object." Kosslyn is more specific about the respects of resemblance in earlier work:

> Images occur in a spatial medium ... [and] the pattern formed in the spatial medium is a topographic mapping from the represented object such that (a) each local portion (set of contiguous points) of the image corresponds to a portion of the represented object as seen from a particular point of view, and (b) the interval relations among the portions of the image implicitly represent ... the interval distances among the corresponding portions of the represented object.
>
> (Kosslyn 1980: 33)[4]

Images resemble their objects with respect to contiguity and interval relations. Intuitively, a set of points is *contiguous* if from any point in that set, you can get to any other, without leaving the set.[5] The claim is that all contiguous parts of the image map onto contiguous parts of the image's content. Place a circle anywhere you like on the photo, and that region maps onto a contiguous region of the photo's field of view. Contiguous regions can overlap, and Kosslyn claims that *all* contiguous regions of the representation map onto contiguous regions of the scene. It is less clear what Kosslyn means by "interval relations" if this is not captured by the first claim about contiguity. Because it's not central to present concerns, we can just focus on the contiguity claims and let interval relations fall where they may.[6]

Michael Tye winds up in a similar place describing images. Instead of contiguity, he focuses on adjacency. A representation is an *image* only if

> (i) every part of R that represents anything represents a part of O visible from V; (ii) a sufficient number of apparent adjacency relationships among parts of O visible from V are represented in R; (iii) any apparent adjacency relationship among the parts of O that is represented in R is

> represented in such a manner that the parts of R representing those O parts are (literally) adjacent to one another.
>
> (Tye 1991: 44)

The first condition is an explicit endorsement of the Sober/Fodor claim mentioned above. The second claim is semantic: some adjacency relationships are represented. And the third claim is syntactic–semantic: those adjacency relationships are represented by the corresponding parts of the image being adjacent. In the present context, preserving adjacency amounts to much the same thing as preserving contiguity, so these conceptions are not far apart. To see why, notice that Tye must have contiguous parts in mind when he talks about "every part of R." The reason for this is that it is unclear how we should define adjacency for parts that are not contiguous. So, Tye is already suggesting that mental images preserve contiguity, and doing so will preserve adjacency as well.[7]

Kosslyn and Tye, in what we have seen so far, suggest that mental images share qualities, like contiguity and adjacency, with the scenes they represent. In that sense, these mental representations fall into the class of non-pictorial images, akin to fMRI scans and radar images. They are structure-preserving, but in many significant respects they do not share qualities with their objects. So far, we have just been trying to clarify what Kosslyn, Tye, and others mean by "image." It also turns out that there is empirical motivation for their claims.

We have good evidence that many brain regions are "retinotopically mapped," in the sense that they have a structure that mimics the retina. In the retina, contiguous regions correspond to contiguous parts of the visual field, adjacent regions in the retina correspond to adjacent regions of the visual field, and so on. In a well-known paper, Van Essen, Newsome, and Maunsell suggest that:

> One important feature of the cortical representation ... is that smooth sequences of recording sites within the cortex invariably correspond to smooth progressions of receptive field centers. Thus, at this level of resolution the cortical representation preserves all neighborhood relationships within the visual hemifield. ... Moreover, this representation is highly regular, insofar as there is an orderly spacing between receptive field centers for recording sequences in which the sampling was at fairly constant intervals.
>
> (Van Essen, Newsome, and Maunsell 1984: 434)

Kosslyn and Tye are not just speaking in the abstract about what makes something an image. They are working in light of what we know about functional neuroanatomy. This might not be the most helpful way of characterizing mental images, however, even given the empirical facts, for two reasons.

First, we know that mental images cannot resemble their objects in any but a few respects. Tye (1991) goes on, for example, to characterize mental images as "symbol-filled arrays" suggesting that most of what these images represent, aside from things like spatial adjacency, is not represented imagistically. That might, at the end of the day, be true, but we don't want to suggest that it is true simply because we have an overly restrictive notion of images, and we have already seen that images, broadly speaking, need not share qualities with what they represent. Second, one point of asking if there are images in the brain is to account for how perceptual states (1) give us information about the world and (2) are different from other mental states, like thoughts. This general goal is appealing, but it applies across sense modalities, and the talk of contiguity and adjacency seems geared specifically towards vision. These conceptions of images are in danger of applying exclusively to vision.

Areas of the auditory system, for example, from the brain stem through the midbrain, thalamus, and cortex are arranged *cochleotopically*, not retinotopically, which means that cells in a neighborhood respond preferentially to a given range of frequencies (Imig and Reale 1980; Brugge 1992; Clary et al. 1992). As one moves along an area of auditory cortex, one moves through the range of frequencies one might hear, and not across the spaces at which one might hear things. In this sense, the auditory system is structured quite differently than the visual system. Barn owls are an interesting contrast to the human case, because their auditory cortices are laid out spatiotopically, much more like our retinotopic visual systems, no doubt because spatial hearing is so important to an animal that hunts in the dark (e.g. Maczko, Knudsen, and Knudsen 2006). It would be nice, however, if our account of mental images could include auditory representations too. If it does not, then it's hard to see the theoretical value of talking about mental images at all.

8.3 Mental mirroring II: functional space

The previous section kept tabling discussion of Kosslyn's remark: "the spatial relations among those patterns *in the functional space* correspond to

the relations among the parts themselves" (Kosslyn 1994: 5, emphasis added). That hedge is important. We find a similar one here:

> ... each part of the depiction must correspond to a visible part of the object such that the distances among the points on the object are preserved by the corresponding "distances" (including empty locations) among their representations.
>
> (Kosslyn 1984: 107; cf. Kosslyn, Thompson, and Ganis 2006: 12)

Kosslyn pulls back from the strong claim that mental images share spatial qualities with their objects. It's not intervals or distances, but "distances" or distances in "the functional space" that matter. There might be a system that *functions* just like the ones we find in the brain but does not as a matter of fact preserve contiguity or adjacency. We could say of such a system that it preserves a kind of functional contiguity and adjacency, even though it does not preserve the actual relationships.[8] This suggests we do not need genuine resemblance, even in limited respects like contiguity or adjacency, to have mental images. The theoretical tools for explaining how he might soften the resemblance claims are already in hand: perhaps mental images are merely isomorphic or homomorphic representations. They do not share qualities with their objects, but they share a structure with them.

In a structure-preserving system, relations between groups of neurons representing different portions of the visual field map onto the spatial relations of corresponding regions of the visual field.[9] Such a representation can function just like the resembling one because both have parts and relations between them corresponding to parts and relations in the represented domain. Actions taken with respect to the representation, such as being sensitive to one part of it, to the exclusion of other parts, can map without remainder onto the thing it represents.

Why pull back from resemblance? As mentioned in the last section, we want an account of mental images to work across modalities, but spatial relations are the only decent candidates for resemblance. This would leave us suggesting that there are no auditory mental images, even though there are auditory mental representations that do much the same work as visual mental images. It seems, that is, as though the resemblances that we do find in the visual case are *contingent* and that what does the work is shared structure. For example, Kosslyn and colleagues suggest that contingent facts about neural function speak in favor of representations of the visual field being laid out in a space that literally shares contiguity and adjacency

with the visual field itself. Let's look briefly at what those contingent features might be.

Much of what the earliest areas of visual and auditory processing do is detect difference. Where is there a high-contrast border? How is it oriented? Neurons manage this feat in part by a process of lateral inhibition enabled by spatial proximity.[10] The details of lateral inhibition are irrelevant for present purposes. What matters here is that neurons work well together when they are near to one another. So, it pays, neurologically speaking, to represent a spatial layout of high-contrast boundaries with an ensemble of cells that is spatially similar to the scene it represents. But Kosslyn realizes that if neurons worked differently we might have representations of visual space that do not mimic the contiguity and adjacency relationships of the visual field. Such states could nevertheless be mental images. Differently structured states—structure-preserving but not resembling—could do the same things for perceivers.

This fact about what visual representation might have been like, if our neurons had been different, carries over nicely to what auditory representation *is* like. Lateral inhibition mechanisms, so important for sorting out high contrast boundaries in vision, are deployed to distinguish component frequencies of the sounds we hear. In the barn owl, as we noted earlier, the auditory cortex represents spatial features in a way similar to how our visual cortexes do. In humans and barn owls alike, however, we can make sense out of the idea that there is auditory imagery, because images in the broad sense are structure-preserving representations, which need not resemble their objects, even though some of them do.

The thought that got us here in the first place was a worry that the resemblances between the mental image and its object were incidental: merely a result of contingent facts about how neurons work. Once we have made the move to isomorphism, however, we also notice that we cannot deny that a mental representation is imagistic just because it fails to resemble its objects. Tye (1991), for example, insists that mental images are symbol-filled arrays, because there is no interesting sense in which brain states share color or texture properties with what they represent. True: the brain is mostly gray and white, and it's got a limited palette of textures too. But that does not undermine the claim that representations of color in the brain share a structure with the colors. In different ways, Robert Cummins (1996), David Rosenthal (2005) and the present author (Kulvicki 2005) suggest that brain states share structure with the colors

they represent. Brain states can share a structure with things' shapes, sizes, colors, and anything else of interest.

The world must cooperate, of course. If it were shown that there are no shared structures between perceptual states and perceptible qualities, the hypothesis that there are mental images, and that they play an important role in our psychological economies, would be in peril. A review of neuroscience is beyond the scope of this chapter. The goal is not to establish the claim that there are mental images, so much as clarify the claim and assess its prospects. The results about retinotopic and tonotopic organization mentioned above provide a direction for further investigation.[11] For now, our focus will be on why the hypothesis is helpful. Why is it so important that brain states, specifically brain states that occur as part of perception, share a structure with what they represent? One answer is that these perceptual states must make aspects of the world available to perceivers, and structure-preserving representations are excellent at doing just that.

Fodor once said that "what perception must do is so represent the world as to make it available to thought" (1983: 40). We are on our way to understanding how images, in the broad sense, might allow just that. These last two sections asked what the syntactic–semantic structure of a brain state must be like in order for it to count as an image. The main claim is that, at a minimum, such representations must be structure-preserving. The next section looks at what philosophers have taken to be distinctive about perceptual *contents*. Are the contents of perceptual states distinct from the contents of other mental states, and if so, how? The final section brings these thoughts together and shows how structure-preserving states "so represent the world as to make it available to thought."

8.4 Perceptual content

Not as much is written these days in philosophy about the imagery debate as is written about the nature of perceptual content. This "non-conceptual content debate" stems partly from the imagery debate and partly from a distinction between two influential theories of perception formulated in the 1960s.[12] It focuses on whether perceptual states can have contents that do not correspond to any of the concepts a perceiver has. Alternatively, is the best way of articulating the content of a perceptual state at least sometimes, and perhaps often, going to use concepts that the perceiver does not have?

One argument in favor of the claim that perceptual states have contents that are non-conceptual is that they are *richer* and *of a finer grain* than those of typical linguistic representation or thought (Evans 1982: 229; Heck 2000: 489f.; Tye 1995: 139; 2005: 519). Richness just captures the idea that perceptual states present viewers with information about many things and many of their qualities.[13] We learn about colors, shapes, distances, textures, illuminations, and so on, simply by looking at the world. We need not attend to most of this information: it is there, ready for use should we need it. Similarly, we receive *a lot* of information about particular qualities. We learn about the colors of things, for example, at a fine-grained level that outstrips most of our conceptual repertoire: "… a perceiver can be in perceptual states a faithful report of whose content would necessarily employ concepts she did not possess" (Heck 2000: 484, citing Evans 1982). Imagine seeing a specific shade of red. It is a shade of red, and you probably have the concept "red," but the color you see is also crimson. Even if you have the concept "crimson," and know how to use it, do you have a concept for the specific shade of crimson you see? We are typically in a position to conceptualize our experiences to some degree, but there are always things with which we are presented, for which we seem to have no concepts. The same points that apply to colors apply to shapes, textures, sounds, smells, and everything else you perceive. Do you have a concept for the specific shape of the last apple you ate? Richness and fineness of grain are related, of course, in the sense that fineness of grain is just getting a lot of information about a specific feature of some object.

Some linguistic representations, and some concepts we have, manage to pick out colors and shapes at impressively fine levels of detail. Diana Raffman (1995: 303), for example, points out that it is relatively easy to acquire the concept of unique green. This is a shade of green that is neither bluish nor yellowish. This shade of color is highly determinate, so it's not as though we *cannot* have concepts for such things. Similarly, the concept of square is quite determinate: a figure that is right, equilateral, and four-sided. Such determinate concepts applied to perceptible qualities are somewhat rare—we see a lot of squarish things, but rarely see something as being square in the highly specific sense just articulated—and they come nowhere near to covering the range of determinate features we are in a position to perceive. Similarly, though we can find some highly specific concepts, we never deploy all of the concepts that could reasonably be applied to any given perceptible scene. Perceptual states might make a lot

of information *available* to thought, but it never happens that all of the available information makes its way into thought.

One more interesting feature of perceptual states is that they seem equally good at providing quite specific information about the environment as they are at providing rather abstract information about it. It's easy, for one in the know, to realize that the sweatshirt is colored Dartmouth green, but that very same person has an easy time classifying the shirt as dark green, or just green. The greenness of the sweatshirt seems just as perceptually available as its more specific shade. The idea here is that perceptual content is "vertically articulate" (Kulvicki 2007) in that the representational content of a perceptual state seems to span levels of abstraction. The *concept* of Dartmouth green picks out a color that is also a shade of green, but it does not thereby represent greenness, as the concept of green does. It's uncommon for the content of a concept to be articulate along this vertical dimension, but quite common for the contents of perceptual states to be so.

8.5 Content and structure

As we saw in the last chapter, many scientific representations are useful because of how they make content extractable in a syntactically and semantically salient manner. Structure-preserving representations like pictures, non-pictorial images, and diagrams can make lots of information available across a number of levels of abstraction. This is exactly what perceptual representations do. And even though resemblance is questionable as a tool for understanding perceptual representation, structure preservation does not require resemblance.

Notice that when we say perceptual states have rich and fine-grained contents, and that they are structure-preserving, we are suggesting that the richness of the content is mirrored by richness in syntactic structure.[14] That syntactic structure is what allows one to make use of the myriad parts of the content easily. Information about color is readily available for our conceptual faculties because aspects of the perceptual representation are responsible for representing color and nothing else. Ditto for shapes, textures, and so on. Not only is information about the specific qualities readily available, but information about less specific things is readily available as well. For example, let's say a set of qualities (r_1, r_2, r_3, … r_n) is such that each member stands for one of the specific shades of red. The following quality does the job of representing red things, and

nothing else: r_1-or-r_2-or-r_3-or- … -r_n. Any state that has one of those qualities is a state representing a red thing. So, if all that matters is whether something is red, one can be sensitive to nothing but the fact that a state has one or another of those syntactic qualities. Some more restricted disjunction of those qualities does the job of representing all the shades of vermilion, another, the shades of dark red, and so on. Structuring a representation so that it is isomorphic to the set of determinate shades of red also means it will be isomorphic with respect to many abstractions from those determinate shades (Kulvicki 2007: 363).

Focusing on structure preservation also gives sense to the claim that mental images are very replete and sensitive representations. Syntactically simple representations can have quite detailed contents, just like syntactically complex ones. Perceptual states need to represent detailed states of affairs, however, in a way that makes aspects of that detail available to thought. And the model on the table at this point is that they accomplish this by sharing a structure with the complex states of affairs they represent. They have parts responsible for carrying information about distinct aspects of the complex environment they represent. Absent a rather impressive palette of features that are responsible for carrying information, they are unable to share structure with what they represent, and thus unable to make those aspects of information readily available. Since other mental representations are not responsible for doing such things, there is no reason to think they are as syntactically replete as perceptual representations.

Claims about sensitivity follow in similar fashion. Perceptual states are geared toward representing a range of qualities along dimensions of similarity. The colors are the kinds of things we can grasp perceptually, for example, along the dimensions of hue, saturation, and brightness (e.g. Rosenthal 2005). Other mental representations, like concepts and the thoughts that they constitute, can also represent colors, but they are not charged with capturing the subtle ways in which they differ from one another. The subtlety of those dimensions suggests that perceptual states that differ only slightly from one another can represent different states of affairs, which is just the mark of syntactic sensitivity discussed in Chapter 5.

None of the foregoing establishes that there are structure-preserving representations in the head, but it does show two other important things. First, if there are such representations operative in perception, we get a tidy explanation of the features distinctive of such states. That is, we know

that if states are structured as suggested here then they can have quite rich, fine-grained contents that make information about the environment available at many levels of abstraction. Second, we see how the tools developed to understand images, in the broad sense, in Chapters 5 and 7 can fruitfully be applied to the mental imagery/non-conceptual content issue.

8.6 Revisiting pictures and parts

Before closing, let's revisit the Sober/Fodor claim about pictures that we encountered toward the beginning of our discussion:

> Every representational picture has representational subpictures; successive snippings never destroy representationality.
>
> (Sober 1976: 124)

> If P is a picture of X, then parts of P are pictures of parts of X.
>
> (Fodor 2008: 173)

Key to understanding the significance of these claims is clarifying the notion of a part, and it's essential to note that they are talking about the broader class of images, and not specifically about pictures. Sober snips his figures, which suggests he is thinking of spatial parts of images, and Fodor's discussion of his "Picture principle" reveals the same focus. There is nothing wrong with focusing there, but doing so misses a deeper truth.

Recall that Fodor's statement had to be qualified: not any part of a photo is a photo of a part of its subject, because the back portion of a photo represents nothing at all. Fodor has to focus on the syntactic features of the image, not its incidental ones. Tye does this precisely when he says "every part of R *that represents anything* represents a part of O visible from V" (1991: 44, emphasis added). The focus on syntactic features reveals that the reason Sober and Fodor can snip is that spatial parts of many images are syntactic, and snipping amounts to focusing on some of them, to the exclusion of others.

But there are many ways of focusing on just some of the syntactic features of an image. We have seen some, albeit not under this banner before. For example, we saw that we can ignore the color of a photo, focusing instead on lightness and space, and still have an interpretable image on our hands. We can relieve the specific shapes on a surface of the

burden of representing specific shapes of things, and still have something we can interpret. This feature of pictures, specifically, was important for understanding what makes kind realism interesting, as we saw in Chapter 6. But in the present context—images in mind and science—it takes on a deeper significance. Structure-preserving representations make parts of their contents readily available across levels of abstraction. It's precisely this that accounts for their usefulness in science and in the mind. The way in which they do this is by sharing a structure with what they represent. But how does sharing a structure permit them to do this? Precisely by having syntactic parts that are responsible for representing parts of their overall content. In a much more general sense than that envisaged by Sober and Fodor, parts of images are images of parts.

In this context, "part" has no specifically spatial connotations. A part of an image could just be the pattern of color it manifests at a certain level of abstraction, or a limited palette of its spatial features, even if it's not a contiguous part. Ignoring determinate color and shape, we wind up with an image that represents an indeterminately colored and shaped scene: a part of that scene, in the extended sense of the word. This is a feature of structure-preserving representations generally, not just pictures. It helps explain why they are good at doing what they do, as we saw for scientific representations and mental representations over the past two chapters. It might help explain why pictorial realism, per se, is so interesting aesthetically. It's very hard to see any of this absent the tools for structurally analyzing representational kinds developed in this chapter, and in Chapters 5 and 7.

Summary

Mental representations are neither produced nor consumed in the same way as artifacts. Most accounts of images fail to extend straightforwardly to the mind because they understand images in terms of their consumers' responses, not their syntactic and semantic structure. We do not experience, recognize, or engage in make-believe with our mental states in the way we do with artifacts. It's possible, by contrast, to describe artifacts in terms of their syntactic and semantic structure. Unsurprisingly, when psychologists and philosophers tried to understand mental imagery, they naturally gravitated toward structural accounts, articulated in terms of limited, spatial resemblances, or isomorphism.

Both of these approaches are plausible. In one case, there are non-pictorial images in the brain, while in the other, there are merely structure-preserving representations there. In both cases, we have images, in the broad sense intended in this book, in the head. Because not all perceptual modalities are as richly spatial as vision, we cannot expect spatial resemblances in perceptual representation across the sense modalities. In fact, the preservation of spatial structure in the visual system might be due to contingent features of how neurons process information, and thus be completely coincidental. More important, in the visual cases as well as other modalities, is the fact that there is a shared structure between perceptual representations and what they represent.

This shared structure is important because, as we saw in the previous chapter, structure-preserving representations make many parts of their content readily available, across levels of abstraction. This is precisely what we want from perceptual representations. They need to present the world to perceivers, so that they can think about it in ways that serve their ends. We saw, too, that this way of understanding the value of mental images sheds light on the Sober/Fodor claim about pictures. They focus on the way in which spatial parts of pictures of X are pictures of spatial parts of X. But both Sober and Fodor are committed to thinking of parts in terms of the syntactic features of images. Recognizing this, we can broaden the notion of part to encompass limited groups of the syntactic features of these representations. Indeed: pictures of parts of X are parts of pictures of X, but this is not limited to spatial parts.

Further reading

Zenon Pylyshyn (1981, 2006) has been the most prominent and persistent critic of the claim that there are images in the brain. Ned Block's collection *Imagery* (1981) is a *locus classicus* for papers on the topic. York Gunther (2003) edited a very helpful collection of papers on non-conceptual content. Ruth Millikan (2004) has an interesting perspective on the problem of perceptual content, and Katalin Balog (2009) responds to Fodor's discussion of the picture principle. And recent work by Michael Rescorla (2009a) and Dominic Gregory (2010b) approaches the problem of mental imagery from the perspective of what makes the contents of images distinct from the contents of other kinds of representation. Richard Watson (1995) provides an historical overview of the role that images play in our philosophical understanding of perception.

Notes

1 Mark Rollins suggests that the recognition theory should play a part both in an account of pictorial representation (1999) and in an account of mental imagery (2001). He does this by suggesting that "an image is just activity in areas of the brain that are dedicated to fundamental visual processes" (2001: 282–83; cf. 1989: 103–4). Rollins's suggestion is that much of the same machinery is involved in understanding mental images as is involved in understanding pictorial representations. This does not undermine the worries about applying the recognition view as a way to show that there are imagistic representations in the brain. It might, however, help address the other problem of mental imagery: the problem of what happens when we visually imagine something. Specifically, Rollins suggests that when we visually imagine seeing something we generate states in us that lead to recognition responses akin to those that states of perceiving the object do.

2 Though not focused on Kosslyn's theory of images, Rollins (2001) offers a helpful discussion of the moving parts of Kosslyn's (1994) overall view.

3 See also Fodor (2007). Not any part of a photo represents a part of the scene it depicts: for example, part of the photo's back surface. Fodor needs to restrict his notion of part to the *syntactic features* of the photo. This restriction gets interesting, as we will see later on. This idea is implicit in a lot of work related to pictorial representation.

4 And compare Kosslyn (1984: 107) and Kosslyn, Thompson, and Ganis (2006: 12).

5 This is called "path-connectedness." I thank Craig Sutton for his help here.

6 Roberto Casati and Achille Varzi (1999) suggest that we can understand maps and related representations in terms of structure preservation like this. They also suggest that this approach can be extended to pictures and images (1999: 187) in general. Rescorla (2009b) proposes a different account of cartographic representation.

7 To see this, it's easy to imagine mapping two adjacent contiguous regions to two non-adjacent contiguous regions, but it's not easy to imagine doing this for all contiguous regions. Specifically, the contiguous region that includes both of the original adjacent areas also gets mapped onto a contiguous region, which is incompatible with mapping the original areas onto non-adjacent areas of the image.

8 Kosslyn, Thompson and Ganis (2006: 17). But see Tye (1991: 35–38) for worries about how Kosslyn (1980) unpacks the idea.

9 Psychologists have identified isomorphism as important for understanding perceptual states for quite some time, though they have not been completely univocal in the way they understand the notion, and it has not always been the notion used by mathematicians and logicians, like Carnap (1958). In this vein, see Shepard and Chipman (1970), Brown and Herrnstein (1981), and also Kulvicki (2004: 393–94 n. 10) for discussion.

10 Rollins (2001: 281–82); Kosslyn et al. (2006: 17 and 23 n. 5), citing Chklovskii and Koulakov 2004.

11 The hypothesis gains more support from the extended discussion in Kosslyn (1994), and Kosslyn et al. (2006).

12 Namely, David Armstrong's (1961, 1968) and Daniel Dennett's (1969) cognitivist approaches and Fred Dretske's (1969) non-cognitivist approach to perceptual states. This makes for a philosophically rich story, in fact, but it's sadly beyond the scope of this book to cover it in any detail.

13 This is not the same as semantic richness, which we encountered in Chapter 5.

14 Remember, in this connection, a point that came up first in Chapter 3 and resurfaced in Chapter 5. Those who insist on resemblance or structure preservation more generally can identify a level of the representation's content with just those respects in which it resembles the object, is isomorphic to it, and so on.