

Mental imagery: In search of a theory

Zenon W. Pylyshyn

Rutgers Center for Cognitive Science, Rutgers University, Busch Campus,
Piscataway, NJ 08854-8020. zenon@ruccs.rutgers.edu
<http://ruccs.rutgers.edu/faculty/pylyshyn.html>

Abstract: It is generally accepted that there is something special about reasoning by using mental images. The question of how it is special, however, has never been satisfactorily spelled out, despite more than thirty years of research in the post-behaviorist tradition. This article considers some of the general motivation for the assumption that entertaining mental images involves inspecting a picture-like object. It sets out a distinction between phenomena attributable to the nature of mind to what is called the cognitive architecture, and ones that are attributable to tacit knowledge used to simulate what would happen in a visual situation. With this distinction in mind, the paper then considers in detail the widely held assumption that in some important sense images are *spatially displayed* or are *depictive*, and that examining images uses the same mechanisms that are deployed in visual perception. I argue that the assumption of the spatial or depictive nature of images is only explanatory if taken literally, as a claim about how images are physically instantiated in the brain, and that the literal view fails for a number of empirical reasons – for example, because of the cognitive penetrability of the phenomena cited in its favor. Similarly, while it is arguably the case that imagery and vision involve some of the same mechanisms, this tells us very little about the nature of mental imagery and does not support claims about the pictorial nature of mental images. Finally, I consider whether recent neuroscience evidence clarifies the debate over the nature of mental images. I claim that when such questions as whether images are depictive or spatial are formulated more clearly, the evidence does not provide support for the picture-theory over a symbol-structure theory of mental imagery. Even if all the empirical claims were true, they do not warrant the conclusion that many people have drawn from them: that mental images are depictive or are displayed in some (possibly cortical) space. Such a conclusion is incompatible with what is known about how images function in thought. We are then left with the provisional counterintuitive conclusion that the available evidence does not support rejection of what I call the “null hypothesis”; namely, that reasoning with mental images involves the same form of representation and the same processes as that of reasoning in general, except that the content or subject matter of thoughts experienced as images includes information about how things would look.

1. Why is there a problem about mental imagery?

1.1. The pull of subjective experience

Cognitive science is rife with ideas that offend our intuitions. It is arguable that nowhere is the pull of the subjective stronger than in the study of perception and mental imagery. It is not easy for us to take seriously the proposal that the visual system creates something like symbol structures in our brain since it seems intuitively obvious that what we have in our mind when we look out onto the world, as well as when we close our eyes and imagine a scene, is something that *looks like* the scene, and hence whatever it is that we have in our heads must be much more like a *picture* than a description. Though we may know that this cannot be literally the case, that it would do no good to have an inner copy of the world, this reasoning appears to be powerless to dissuade us from our intuitions. Indeed, the way we describe how it feels to imagine something shows the extent of the illusion; we say that we seem to be *looking at something* with our “mind’s eye.” This familiar way of speaking reifies an observer, an act of visual perception, and a thing being perceived. All three parts of this equation have now taken their place in one of the most developed theories of mental imagery (Kosslyn 1994), which refers to a “mind’s eye” and a “visual system” that examines a “mental image” located in a “visual buffer.” Dan Dennett has referred to this view picturesquely as the “Cartesian Theater” view of the mind (Dennett 1991) and I will refer to it as the “picture theory” of mental imagery. There has been a tradition

of analyzing this illusion in the case of visual perception, going back to Descartes and Berkeley (it also appears in the seventeenth century debate between Arnaud and Malebranche – see Slezak 2002a), and revived in modern times by (Gibson 1966), as well as computationalists like (Marr 1982). More recently, O’Regan (1992) and O’Regan and Noë (2001) have argued against the intuitive picture-theory of vision on both empirical and theoretical grounds. Despite the widespread questioning of the intuitive picture view in visual perception, this view remains very nearly universal in the study of mental imagery (with such notable exceptions as Dennett 1991; Rey 1981; Slezak 1995); (see also

ZENON PYLYSHYN is Board of Governors Professor of Cognitive Science at the Rutgers Center for Cognitive Science, which he helped to found in 1991. Prior to that he was at the University of Western Ontario, where he was associated with the departments of psychology, computer science, philosophy, and electrical engineering. He is recipient of the Donald O. Hebb Award and is a fellow of the Canadian Psychological Association, the American Association for Artificial Intelligence, and the Royal Society of Canada.

Pylyshyn’s research is on vision and visual attention, and he has also written on theoretical issues. He is author of over 120 articles, book chapters, and books. His new book *Seeing and Visualizing: It’s not what you think* will be published this summer by the MIT Press.

the critical remarks by Fodor 1975; Hinton 1979; Pylyshyn, forthcoming; Thomas 1999).

Why should this be so? Why do we find it so difficult to accept that when we “examine our mental image” we are not in fact examining an inner state, but rather are contemplating what the inner state is about – that is, some possible state of the visible world – and therefore that this experience tells us nothing about the nature and form of the representation? Philosophers have referred to this displacement of the object of thought from the (possible) world to a mental state as the “intentional fallacy” and it has much of cognitive science in its grip still.

What I try to do in this article is show that we are not only deeply deceived by our subjective experience of mental imagery, but that the evidence we have accumulated to support what I call the “picture theory” of mental imagery is equally compatible with a much more parsimonious view, namely, that most of the phenomena in question (but not all – see below) are due to the fact that the task of “imagining” invites people to simulate what they believe would happen if they were *looking at* the actual situation being visualized. I will argue that the alternative picture theory, or depiction-theory, trades so heavily on a systematic ambiguity between the assumption of a literal picture and the much weaker assumption that visual properties are *somehow* encoded. I will also argue that recent evidence from neuroscience (particularly the evidence of neural imaging) brings us no closer to a plausible picture theory than we were before this evidence was available.

1.2. The imagery debate: What was it about?

There has been a great deal of discussion in the past thirty years about “the imagery debate.” Many people even believe that the debate has, at least in general outline, been put to rest because we now have hard evidence from neuroscience showing what (and where) images are (see, e.g., Kosslyn 1994; and the brief review in Pylyshyn 1994a). But if one looks closer at the “debate” one finds that what people think the debate is about is very far from univocal. For example, some people think that the argument that has been settled is whether images, whatever their nature, are fundamentally different from the form of representation involved in other kinds of reasoning, or whether there are two different systems of mental codes. For others, it is the question of whether images have certain particular properties – for example, whether they are spatial, or depictive, or analogue. Others feel that the question that has been settled is whether imagery “involves” the visual system. I will argue that none of these claims has been sufficiently well posed to admit of a solution. In this article I will concentrate primarily on a particular class of theory of mental imagery, which I refer to as “picture theories,” and will consider other aspects of the “debate” only insofar as they bear on the alleged pictorial nature of images.

In this article I defend the provisional view, which I refer to as the “null hypothesis,” that at the relevant level of analysis – the level appropriate for explaining the results of many experiments on mental imagery – the process of imagistic reasoning involves the same mechanisms and the same forms of representation as are involved in general reasoning, though with different content or subject matter. This hypothesis claims that what is special about image-based thinking is that it is typically concerned with a certain

sort of content or subject matter, such as optical, geometrical, or what we might call the appearance-properties of the things we are thinking about. If so, nothing is gained by attributing a special format or special mechanisms to mental imagery. While the validity of this null hypothesis remains an open empirical question, what is not open, I claim, is whether certain currently popular views can be sustained.

In the interest of full disclosure I should add that I don't really, in my heart of hearts, believe that representations and processes underlying imagery are no different from other forms of reasoning. Nonetheless, I do think that nobody has yet articulated the specific way that images are different and that all candidates proposed to date are seriously flawed in a variety of ways that are interesting and revealing. Thus using the null hypothesis as a point of departure may allow us to focus more properly on the real differences between imagistic and other forms of reasoning.

1.3. Plan of the article

Section 2 reviews some observations that have led many people to hold a picture theory of mental images (although a detailed discussion of what such a theory assumes is postponed until sect. 5). Section 3 introduces a distinction that is central to our analysis. It distinguishes two reasons why imagery might manifest the properties that are observed in experiments. One reason is that these properties are *intrinsic* to the architecture of the mental imagery system – they arise because of the particular brain mechanisms deployed in imagery. The other reason is that the properties are *extrinsic* to the mechanisms employed – they arise because of what people tacitly believe about the situation being imagined, which they then use to simulate certain behaviors that would occur if they were to witness the corresponding situation in reality. This distinction is then applied to some typical experiments on mental imagery where I argue that such experiments tell us little about *special* dedicated imagery mechanisms. Since section 4 discusses some material that has been published elsewhere, readers who have followed the “imagery debate” may wish to skim this section.

Section 5 discusses two widely held views about the nature of mental images (Kosslyn 1994): that images are “depictive” and that they are laid out in a “functional space.” I claim that the preponderance of evidence argues against the inherent spatial nature of mental images. An exception is evidence from experiments in which subjects project their images onto a visual scene. In this case I claim (sect. 5.3) that the use of visual indexes and focal attention provides a satisfactory explanation for how spatial properties are inherited from the observed scene, without any need to posit spatial properties of images. In section 2, I argue that the notion of a functional space is devoid of any explanatory power, since such a “space” is unconstrained and can have whatever properties one wishes to attribute to it (unless it is taken to be a simulation of a real spatial display as in the model described in Kosslyn et al. 1979, in which case the underlying theory really is the *literal* picture theory). Section 6 discusses a claim that is assumed to be entailed by the depictive nature of images; namely, that information in an image is accessed through vision. Although there is evidence for some overlap between the mechanisms of imagery and those of vision, a close examination of this evidence shows that it does not support the assumption of a spatial display in *either* vision or imagery. Section 7 consid-

ers evidence from neuroscience, which many writers believe provides the strongest case for a picture theory. Here I argue that, notwithstanding the intrinsic interest of these findings, they do not support the existence of any sort of depictive display in mental imagery. Finally, section 8 closes with a brief discussion of where the “imagery debate” now stands and on the role of imagery in creative thinking.

2. What is special about image-based reasoning?

Imagery seems to follow principles that are different from those of intellectual reasoning and certainly beyond any principles to which we have conscious intellectual access. Imagine a baseball being hit into the air and notice the trajectory it follows. Although few of us could calculate the shape of this trajectory, none of us has any difficulty imagining the roughly-parabolic shape traced out by the ball in this thought experiment. Indeed, we can often predict with considerable accuracy where the ball will land (certainly a properly situated professional fielder can). It is very often the case that by visualizing a certain situation, we can predict the dynamics of physical processes that are beyond our ability to solve analytically. Is this because our imagery architecture inherently and automatically obeys the relevant laws of nature?

Opposing the intuition that one’s image unfolds according to some internal principle of natural harmony with the real world, is the obvious fact that it is *you* alone who controls your image. Perhaps, as Humphrey (1951) once put it, viewing the image as being responsible for what happens in your imagining puts the cart before the horse. In the baseball example above, isn’t it equally plausible that the reason the imagined ball takes a particular path is because, under the right circumstances, you can recall having seen a ball inscribe such a path? Surely your image unfolds as it does because you, the image creator, *made* it do so. You can imagine things being pretty much any size, color, or shape that you choose and you can imagine them moving any way you like. You can, if you wish, imagine a baseball sailing off into the sky or following some bizarre path, including getting from one place to another without going through intervening points, as easily as you can imagine it following a more typical trajectory. You can imagine all sorts of physically impossible things happening – and cartoon animators frequently do, to our amusement.

Some imagery theorists might be willing to concede that in imagining physical processes we must use our tacit knowledge of how things work, yet insist that the optical and geometrical properties of images are true intrinsic properties, despite that fact that the dynamic properties of images that are often cited in studies of mental images – properties such as mental rotation, mental scanning, or “representational momentum” discussed in sections 3.1 and 4. Nonetheless, the suggestion that the intrinsic properties of images are geometrical rather than dynamic makes sense both because spatial intuitions are among the most entrenched, and because there is evidence (Pylyshyn 1999) that geometrical and optical-geometrical constraints are built into the early-vision system, as so-called “natural constraints.” While we can easily imagine the laws of physics being violated, it seems nearly impossible to imagine the axioms of geometry and geometrical optics being violated. Try

imagining a four-dimensional block, or how a cube looks when seen from all sides at once, or what it would look like to travel through a nonEuclidian space. However, before concluding that these examples illustrate the intrinsic geometry of images, consider whether your inability to imagine these things might not be due to your not *knowing*, in a purely factual way, how these things might look (that is, where edges, shadows, and other contours would fall). The answer is by no means obvious. It has even been suggested (Goldenberg & Artner 1991) that certain deficits in imagery ability resulting from brain damage are a consequence of a deficiency in the patient’s knowledge about the appearance of objects. At the minimum we are not entitled to conclude from such examples that images have the sort of inherent geometrical properties that we associate with pictures.

We also need to keep in mind that, notwithstanding one’s intuitions, there is reason to be skeptical about what one’s subjective experience reveals about the form of a mental image. After all, when we look at an actual scene we have the unmistakable subjective impression that our perceptual representation is of a detailed three-dimensional panoramic view, yet it has now been convincingly demonstrated that the information available to cognition from a single glance is extremely impoverished, sketchy, and unstable and that very little is carried over across saccades (see, e.g., Blackmore et al. 1995; Carlson-Radvansky 1999; Carlson-Radvansky & Irwin 1995; Intraub 1981; Irwin 1993; O’Regan 1992; O’Regan & Noë 2001; Rensink 2000a; 2000b; Rensink et al. 1997; 2000; Simons 1996). Indeed, there is now considerable evidence that we visually encode very little in a visual scene unless we explicitly attend to the items in question, and that we do that only if our attention or our gaze is attracted to it (Henderson & Hollingworth 1999; although see O’Regan et al. 2000). There are remarkable demonstrations that when presented with alternating images, people find it extremely difficult to detect a difference between the two – even a salient difference in a central part of the image.¹ This so-called *change blindness* phenomenon (Simons & Levin 1997) suggests that, notwithstanding our phenomenology, we are nowhere near having a detailed internal display since the vast majority of information in a visual scene goes unnoticed and unrecorded. It would thus be reasonable to expect that our subjective experience of mental imagery would be an equally poor guide to the form and content of the information in our mental images.

3. Why images exhibit certain properties: Cognitive architecture or tacit knowledge?

Nobody denies that the content and behavior of our mental images can be the result of what we intend our images to show, what we know about how things in the world look and work, and the way our mind or our imagery system constrains us. The important question about mental imagery is: which properties and mechanisms are *intrinsic* to, or *constitutive* of having and using mental images, and which arise because of what we believe, intend, or attribute to the situation we are imagining.

The distinction between effects attributable to the intrinsic nature of mental mechanisms and those attributable to more transitory states, such as people’s beliefs, utilities, habits, or interpretation of the task at hand, is central not only for understanding the nature of mental imagery, but

for understanding mental processes in general. Explaining the former kind of phenomena requires that we appeal to what has been called the cognitive architecture (Fodor & Pylyshyn 1988; Newell 1990; Pylyshyn 1980; 1984; 1991a; 1996) – one of the most important ideas in cognitive science. It refers to the set of properties of mind that are fixed with respect to certain kinds of influences. In particular, the cognitive architecture is, by definition, not directly altered by changes in knowledge, goals, utilities, or any other representations (e.g., fears, hopes, fantasies, etc.). In other words, when you find out new things or when you draw inferences from what you know or when you decide something, your cognitive architecture does not change. Of course, if as a result of your state of beliefs and desires you decide to take drugs or to change your diet or even to repeat some act over and over, this can result in changes to your cognitive architecture; but such changes are not a direct result of the changes in your cognitive state. A detailed technical exposition of the distinction between effects attributable to knowledge or other cognitive states, and those attributable to the nature of cognitive architecture, is beyond the scope of this article (although this distinction is the subject of extensive discussion in Pylyshyn 1984, Ch. 7). This informal characterization and the following example will have to do for present purposes.

Suppose we have a box of unknown construction, and we discover that it exhibits particular systematic behaviors (discussed in Pylyshyn 1984). The box emits long and short pulses according to the following pattern: pairs of short pulses most often precede single short pulses, except when a pair of long-short pulses occurs first. What is special about this example is that it illustrates a case where the observed behavior, though completely regular when the box is in its “ecological niche,” is not due to the nature of the box (to how it is constructed) but to an entirely extrinsic reason. The reason this particular pattern of behavior occurs can only be understood if we know that the pulses are codes, and the pattern is due to a regularity in *what they represent*, in particular, that the pulses represent English words spelled out in International Morse Code. The observed pattern does not reflect how the box is wired or its functional architecture; it is due entirely to a regularity in the way English words are spelled (the principle being that generally *i* comes before *e* except after *c*). Similarly, I have argued that in most of the core experiments on mental imagery – such as the mental scanning case described in section 4.1 – the pattern does not reveal the nature of the mental architecture involved in imagery, but reflects a principle that observers know governs the world being imagined. The reason why under certain conditions the behavior of both the code box and the cognitive system does not reveal properties of its intrinsic nature (i.e., of its architecture), is that both are capable of quite different regularities if the world they were representing behaved differently. They would not have to change their architecture in order to change their behavior. The latter observation, concerning the plasticity of nonarchitectural properties of thought, is the key to a methodology I have called “cognitive penetrability” for deciding whether tacit knowledge or cognitive architecture is responsible for some particular observed regularity (see sect. 3.2).

In interpreting the results of imagery experiments, it is clearly important to distinguish between cognitive architecture and tacit knowledge as possible causes. Take the following example. You are asked what color you see if you

look through a yellow filter superimposed on a blue filter. The way that many of us would go about solving this problem, if we did not know the answer as a memorized fact, is to imagine a yellow filter and a blue filter being superimposed; we generally use the “imagine” strategy when we want to solve a problem about how certain things look. What color do you see in your image when the two filters are overlapped? Now ask yourself why you see *that* color in your mind’s eye rather than some other color? Some people (e.g., Kosslyn 1981) have argued that the color you see follows from a property of imagery, presumably some property of how colors are encoded and displayed in images. But since there can be no doubt that you can make the overlapping part of the filters be *any* color you wish, it can’t be that the image format or the architecture involved in representing colors is responsible. What else can it be? It seems clear in this case that the color you “see” depends on your tacit knowledge of the principles of color mixing, or a recollection of how these particular colors combine (having seen something like them in the past). In fact, people who do not know about *subtractive* color mixing generally give the wrong answer: mixing yellow light with blue light produces white light, but overlapping yellow and blue filters allows green light through.

When asked to do this exercise (as reported in Kosslyn 1981), some people claim that they “see” a color that is different from the one they report when they are simply asked to say (without using imagery) what would happen. Results such as this have made people leery of accepting the tacit knowledge explanation. There are indeed many cases where people report a different result when using mental imagery, than when asked to merely answer the question without using their image. It is not clear what moral ought to be drawn from this, however, since it is a general property of reasoning that the way a question is put and the reasoning strategy that is used to get to the answer, can affect the outcome. Knowledge can be organized and accessed in many different ways (see sect. 4.3 for more on the relevance of this to mental imagery studies). Indeed, it need not be accessed at all if it seems like more work than it is worth. For example, consider the following analog of the color-mixing task. Close your eyes and imagine someone writing the following on a blackboard: “ $759 + 356 =$.” Now, imagine that the person continues writing on the board. What number can you “see” being written next? People may give different answers depending on whether they believe that they are supposed to work it out, or whether in the interest of speed they are supposed to guess or merely say whatever comes to mind. Each of these is a *different task*. Even without a theory of what is special about visual imagery, we know that the task of saying what something would look like can be a *different task* from that of solving a certain intellectual puzzle about colors or numbers.

In most of the cases studied in imagery research, it would be odd if the results did not come out the way picture theorists predict. For if the results were inconsistent with the picture-theory, the obvious explanation would be that subjects either did not know how things would work in reality or else they misunderstood the instructions to “imagine *x*.” For example, if you were asked to imagine in vivid detail a performance of the Minute Waltz, the failure of the imagined event to take approximately one minute would simply indicate that you had not carried out the task you were supposed to. Since taking roughly one minute is constitutive of a real performance, it is natural to assume it to be indicative of a realistic imaginary re-creation of such a performance.

3.1. What knowledge is relevant to the tacit knowledge explanation?

The concept of tacit knowledge plays an important role in cognitive science (see, e.g., Fodor 1968), though it has frequently been maligned because it has to be inferred indirectly. Such knowledge is called “tacit” because it is not always explicitly available for, say, answering questions. There may nonetheless be independent evidence that such knowledge exists. This is a point that has been made forcibly in connection with tacit knowledge of grammar or of social conventions, which typically also cannot be articulated by members of a linguistic or social group, even though violations are easily detected. In our case the role of tacit knowledge can sometimes be detected using the criterion of cognitive penetrability, discussed below.

Not only is the notion of tacit knowledge often misunderstood, but in the case of explaining mental imagery results, the *kind* of tacit knowledge that is relevant has also been widely misunderstood. The only knowledge that is relevant to the tacit knowledge explanation is knowledge of *what things would look like* to subjects in situations like the ones in which they are to imagine themselves. Many writers have mistakenly assumed that the tacit knowledge explanation refers to one of several other kinds of knowledge. For example, although tacit knowledge of what results the experimenter expects (sometimes referred to as “experimenter demand effects”) is always an important consideration in psychological experiments (and may be of special concern in mental imagery experiments; see Banks 1981; Intons-Peterson 1983; Intons-Peterson & White 1981; Mitchell & Richman 1980; Reed et al. 1983; Richman et al. 1979), it is not the knowledge that is relevant to the tacit knowledge explanation, as some have assumed (Finke & Kurtzman 1981b). The tacit knowledge explanation does not assume that people know how their visual system or the visual brain works (as Farah [1988] apparently thought). It is also not the knowledge people might have of what results to expect from experiments on mental imagery (as assumed by Denis & Carfantan 1985). Denis and Carfantan studied “people’s knowledge about images” and found that people often failed to correctly predict what would happen in experiments such as mental scanning. But these sorts of questions invite respondents to consider their folk psychological theories to make predictions about psychological experiments. They do not reflect tacit knowledge of what it would look like if the observers were to see a certain event happening in real life. The tacit knowledge claim is simply the claim that when subjects are asked to “imagine *x*,” they use their knowledge of what “seeing *x*” would be like (as well as their other psychophysical skills, such as estimating time-to-collision) and they simulate as many of these effects as they can. Whether a subject has this sort of tacit knowledge cannot always be determined by asking them, and certainly not by testing them for their knowledge of psychology!

Notwithstanding the importance of tacit knowledge explanations of imagery phenomena, it remains true that *not all imagery results are subject to this criticism*. Even when tacit knowledge is involved, there is often more than one reason for the observed phenomena. An example in which tacit knowledge may not be the only explanation of an imagery finding can be found in Finke and Pinker (1982). The example concerns a particular instance of mental scanning (one in which it takes more time to judge that an arrow

points to a dot when the dot is further away). Finke and Pinker argued that these results could not have been due to tacit knowledge because, even though subjects correctly predicted that judgments would take more time when the dots were further away, they failed to predict that the time would actually be longer for the shortest distance used in the study. But this was a result that even the authors failed to anticipate, because the aberrant short-distance time was most likely due to some mechanism (perhaps attentional crowding) different from the one that caused the monotonic increase of time with distance.

Another example in which tacit knowledge does not account for some aspect of an imagery phenomenon is in what has been called “representational momentum.” It was shown that when subjects observe a moving object and are asked to recall its final position from memory, they tend to misremember it as being displaced forward. Freyd and Finke (1984) attributed this effect to a property of the imagery architecture. On the other hand, Ranney (1989) suggested that the phenomenon may actually be due to tacit knowledge. It seems that at least some aspects of the phenomenon may not be attributable to tacit knowledge (Finke & Freyd 1989). But here again there are other explanations besides tacit knowledge or image architecture. In this particular case there is good reason to think that part of the phenomenon is actually visual. There is evidence that the perception of the location of moving objects is ahead of the actual location of such objects (Nijhawan 1994). Eye movement studies have also shown that gaze precedes the current location of moving objects in an anticipatory fashion (Kowler 1989; 1990). Thus, even though the general phenomenon involving imagined motion may be attributable to tacit knowledge, the fact that the moving stimuli are presented visually may result in the phenomena also being modulated by the visual system. The general point in both these examples is that even in cases where tacit knowledge is not the sole determiner of a result in an imagery experiment, the phenomena in question need not reveal properties of the architecture of the imagery system. They may be due to properties of the visual system, the memory system, or a variety of other systems that might be involved.

3.2. Methodological note: “Cognitive penetrability” as a *litmus*

How is it possible to tell whether certain imagery effects reflect the nature of the imagery architecture or the person’s tacit knowledge? In general, methodologies for answering questions about theoretical constructs are limited only by the imagination of the experimenter. Typically, they involve convergent sources of evidence and independent theoretical motivation. One theoretically motivated diagnostic, discussed at length in Pylyshyn (1984), is to test for the *cognitive penetrability* of the observations. This criterion is based on the assumption that if a particular pattern of observations arises because people are simulating a situation based on their tacit beliefs, then if we alter their beliefs or their assumptions about the task, say by varying the instructions, the pattern of observations may change accordingly, in ways that are rationally connected with the new beliefs. So, for example, if we instruct a person on the principles of color mixing, we would expect the answer to the imaginary color-mixing question discussed above to change appropriately. We will see other examples of the use of this crite-

tion throughout this article (especially the examples in sect. 4).

Not every imagery-related phenomenon that is genuinely cognitively impenetrable provides evidence for the nature of mental images or their mechanisms. Clearly, many beliefs are resistant to change by merely being told that they are false. Nonetheless, this criterion has proven useful in identifying parts of the visual system that constitute what is called early vision (Pylyshyn 1999). Cognitive penetrability remains a *necessary but not sufficient condition* for a pattern of observations being due to the architecture of the imagery system.

4. Problem-solving by “mental simulation”: Some examples

The idea that what happens in certain kinds of problem solving can be viewed as off-line *simulation* has had a recent history in connection not only with mental imagery (Currie 1995), but also with other sorts of problems in cognitive science (Klein & Crandall 1995). But even if we grant that the “simulation mode” of reasoning is used in various sorts of problem solving, the question still remains: What does the real work in solving the problem by simulation – a special property of images (i.e., the architecture of the image system), or tacit knowledge?

In what follows I will sketch a number of influential experimental results often cited in support of the picture theory, and compare explanations given in terms of inherent properties of the image with those given in terms of simulation based on tacit knowledge.

4.1. Scanning mental images

Probably the most cited result in the entire repertoire of research motivated by the picture-theory is the image-scanning phenomenon. Not only has this experimental paradigm been used dozens of times, but various arguments about the “metrical” or spatial nature of mental images, as well as arguments about such properties of the mind’s eye as its “visual angle,” rest on this phenomenon. Indeed, it has been referred to as a “window on the mind” (Denis & Kosslyn 1999).

The finding is that it takes longer to “see” a feature in a mental image that is further away from the place in the image an observer was initially focusing upon. So, for example, if you are asked to imagine a dog and inspect its nose, and then to “see” what its tail looks like, it will take you longer than if you were asked to first inspect its hind legs. Here is an actual experiment reported in Kosslyn et al. (1978). Subjects were asked to memorize a map such as the one in Figure 1. They were then asked to imagine the map and to focus their attention on one place – say, the “church.” In a typical experiment (there are many variants of this basic study), the experimenter says the name of a second place (say, “beach” or “tree”) and subjects are asked to examine their image and to press a button as soon as they can “see” the second named place on their image of the map. What many researchers have found consistently is that the further away the second place is from where the subject is initially focused, the longer it takes to “see” the second place in the image.

From this scanning result most researchers have con-




Figure 1. Map to be learned and imaged in one’s “mind’s eye” to study mental scanning

cluded that larger map distances are represented by greater distances in image space. In other words, the conclusion that is drawn from this kind of experiment is that mental images *have* spatial properties – that is, they have spatial magnitudes or distances, as opposed to just encoding such properties in some unspecified manner. This is a strong conclusion about cognitive architecture. It says, in effect, that the symbolic code idea that forms the foundation of computational theories does not apply to mental images. In a symbolic encoding two places can be represented as being further away just the way we do it in language; by saying the places are, say, *n* meters from one another. But the representation of larger distances is not itself in any sense larger.

Is this strong conclusion about the metrical property of mental images warranted? Does the difference in scanning time reveal a property of the architecture, or a property of what is represented? Notice how this distinction exactly parallels the situation in the color-mixing example discussed earlier. There we asked whether a particular observation revealed a property of the architecture, or a property of what people know or believe – a property of the represented situation of which they have tacit knowledge. To answer this question for the scanning experiment we need to determine whether the pattern of increasing reaction time arises from a fixed capacity of the image-encoding or image-examining system, or whether it can be altered by changing subjects’ understanding of the task or the beliefs that they hold about what it would be like to examine a real map; whether it is cognitively penetrable.

This is a question to be settled in the usual way – by careful analyses and experiments. But even before we do the experiment there is reason to suspect that the time-course of

scanning is not a property of the cognitive architecture. Do the following test on yourself. Imagine that there are lights at each of the places on your mental image of the map. Imagine that a light goes on at, say, the beach. Now imagine that this light goes off and one comes on at the lighthouse. Did you need to scan your attention across the image to see the light come on at the lighthouse? Liam Bannon and I repeated the scanning experiment (see the description in Pylyshyn 1981) by showing subjects a real map with lights at the target locations, much as I just described. We allowed the subjects to turn lights on and off. Whenever a light was turned on at one location it was simultaneously extinguished at another location. Then we asked subjects to imagine that very map and to indicate (by pressing a button) when a light was on and they could “see” the illuminated place in their image. The time between button presses was recorded and its correlation to the distances between illuminated places on the map was computed. We found that there was no relation between distance on the imagined map and time. Now, you might think: Of course there was no increase in time with increasing distance, because subjects were not asked to imagine scanning that distance. But that’s just the point: You can imagine scanning over the imagined map if you want to, or you can imagine just hopping from place to place on the imaginary map. If you imagine scanning, you can imagine scanning fast or slow, at a constant speed or at some variable speed, or scanning part way and then turning back or circling around! You can, in fact, do whatever you wish since it is your image.² At least you can do these things to the extent that you can create the phenomenology or the experience of them and provided you are able to generate the relevant measurements, such as the time you estimate it would take to get from point to point.

Whether or not you choose to simulate a certain temporal pattern of events in the course of answering a question may also depend in part on whether simulating that particular pattern seems to be relevant to the task. It is not difficult to set up an experimental situation in which simulating the actual scanning from place to place does not appear so obviously relevant to solving a particular problem. For example, we ran the following experiment that involved extracting information from an image (Pylyshyn 1981). Subjects were asked to memorize a map and to refer to their image of the map in solving the problem. As in the original Kosslyn et al. (1978) studies, subjects had to first focus on one place on their imagined map and then to “look” at a second named place. The experiment differed from the original study, however, in that the task was to indicate the compass direction from the second named place to the previously focused place. This direction-judgment task requires that the subject make a judgment from the perspective of the second place, so it requires focusing at the second place. Yet, in this experiment, the question of how you get from the first place to the second place on the map was far less prominent than it was in the “tell me when you can see *x*” task. In this study we found that the distance between places had no effect on the time taken to make the response. Thus, it seems that the effect of distance on reaction time is cognitively penetrable.³

Not only do observers sometimes move their attention from one imagined object to another without scanning through the space between them, but we have reason to believe that they *cannot* move their attention continuously

through empty imagined space (see sect. 6.4 for a brief description of the relevant study).

4.2. The “size” of mental images

Another study closely related to the mental scanning paradigm, is one where it is found that it takes more time to report some visual detail of an imagined object if the object is imagined to be small, than if it is imagined to be large (e.g., it takes longer to report that a mouse has whiskers if the mouse is imagined as tiny, than if it is imagined as huge). This seems like a good candidate for a tacit knowledge explanation, since when you actually *see* a small object you know that you can make out fewer of its details due to the limited resolution of your eye. So if you are asked to imagine something small, you are likely to imagine it as having fewer visible details than if you are asked to imagine it looming large directly in front of you, *whatever* form of representation that may involve.

The original picture-theory view of this result is problematic in any case. What does it mean for your image to be “larger”? Such a notion is meaningful only if the image *has* a real size or scale. If, as in our null hypothesis, the information in the image is in the form of a symbolic description, then size has no literal meaning. You can *think of* something as large or small, but that does not make *some thing* in your head large or small. On the other hand, *which details are represented* in your imagination does have a literal meaning: You can put more or less detail into your active representation. Inasmuch as the task of imagining the mouse as “small” entails that you imagine it having fewer visible details, the result is predictable without any notion of real scale applying to the image.

The obvious test of this proposal is to apply the criterion of cognitive penetrability. Are there instructions that can ameliorate the effect of the “image size” manipulation, making details easier to report in small images than in large ones, and vice versa? Could you imagine a small but extremely high resolution and detailed view of an object, in contrast to a large but low-resolution or fuzzy view that lacks details? I know of no one who has bothered to carry out an experiment that asks subjects to, say, report details from a large blurry image and then from a small clear one. What if such an experiment were done and it showed that it is quicker to report details from a large blurry object than from a small clear one? The strangeness of such a possibility should alert us to the fact that what is going wrong lies in what it *means* to have a blurred versus a clear image. Such results would be incompatible with what happens in *seeing*. If it took longer to see fine details in a real large object, there would have to be a reason for it, such as, that you were seeing it through a fog or out of focus. Thus, so long as examining a visual image *means* simulating what it is like to see something, the results must be as reported; how could studies involving different sized mental images, or blurred versus clear images, fail to show that they parallel the case of seeing, unless subjects misunderstood the instructions (e.g., they did not understand the meaning of “blurry”)? The same goes for the imagery analogue of any property of seeing of which observers have some tacit knowledge or recollection. Thus, it applies to the findings concerning the acuity profile of imagery, which approximates that of vision (Finke & Kosslyn 1980). Observers do not need to have articulated scientific knowledge of visual

acuity; all they need is to remember roughly how far into the periphery of their visual field things can go before they become hard to see – and it is not surprising that this is easier to do while turning your head (with eyes closed) and pretending to be looking at objects in your periphery (which is how these studies were done).

4.3. Mental “paper folding”

There are many reasons why one might use a “simulation mode” strategy in answering a question; reasons that have nothing to do with the spatial nature of imagery, and sometimes not even with the kind of tacit knowledge available. For example, to answer the question: What is the fourth (or *n*th) letter in the alphabet after “M,” people normally have to go through the alphabetical sequence (and it takes them longer, the larger the value of *n*). Similarly, the findings reported by Shepard and Feng (1972) are easily understood if one considers how the relevant knowledge is organized. In their experiment, subjects are asked to mentally fold pieces of paper, such as shown in Figure 2, and to report whether the arrows marked on the paper would touch one another. Shepard and Feng found that the more folds it would require to actually fold the paper and see whether the arrows coincide, the longer it takes to imagine doing so. From this they concluded that working with images parallels working with real objects.

The question that needs to be asked about this task is the same as the question we asked in connection with the color mixing task: What is responsible for the relation between time taken to answer the question and the number of folds it would have taken in folding real paper? This time the answer is not simply that it depends on tacit knowledge, because in this case it is not just the content of the tacit knowledge that makes the difference. Yes, it is the knowledge that subjects have about paper folding that makes it possible for them to do the task at all. But, in this case, it appears that *imagining* making individual folds is required in order to get the answer. Indeed, it is hard to see how to answer this question without imagining going through the sequence of folds. A plausible explanation for this, which does not appeal to special properties of a mental image system, is that the reason one has to imagine going through a sequence of individual folds is the same as the reason one had to go through a series of letters in the earlier alphabet example. The reason may have to do with how one’s knowledge of the




Figure 2. Two of the figures used in the Shepard and Feng (1972) experiment. The task is to imagine folding the paper (using the dark shaded square as the base) and say whether the arrows in these two figures coincide. The time it takes increases with the number of folds required.

effects of folding is organized. What we know about the effects of paper folding is just this: *we know what happens when we make one fold*. Consequently, to determine what would happen in a task that requires four folds, we have to apply our one-fold-at-a-time knowledge four times. Recall the parallel case with letters: In order to determine what the fourth letter after M is, we have to apply the “next letter” rote knowledge four times. In both cases a person could, in principle, commit to memory such facts as what results from double folds of different types; or which letter of the alphabet occurs exactly *n* letters after a given letter. If that were how paper-folding knowledge was organized, the Shepard and Feng results might not hold. The important point is that once again the result tells us nothing about how the states of the problem are represented – or about any special properties of image representations. They tell us only what knowledge the person has and how it is organized.

The role played by the structure of knowledge is ubiquitous and may account for another common observation about the use of mental imagery in recall. We know that some things are easier to recall than others, and that it is easier to recall some things when the recall is preceded by the recall of other things. Memory is linked in various intricate ways. In order to recall what you did on a certain day it helps to first recall what season that was, what day of the week it was, where you were at the time, and so on. Sheingold and Tenney (1982), Squire and Slater (1975), and others have shown that one’s recall of distant events is far better than one generally believes because once the process of retrieval begins it provides clues for subsequent recollections. The reason for bringing up this fact about recall is that such sequential dependencies are often cited as evidence for the special nature of imagery (Bower & Glass 1976; Paivio 1971). Thus, for example, in order to determine how many windows there are in your home, you probably need to imagine each room in turn and look around to see where the windows are, counting them as you go. In order to recall whether someone you know has a beard (or glasses or red hair), you may have to first recall other aspects of what he or she looks like (that is, recall an image of them). Apart from the phenomenology of recalling an appearance, what is going on is absolutely general to every form of memory retrieval. Memory access is an ill-understood process, but at least it is known that it has sequential dependencies and other sorts of access paths, and that these paths are often dependant on spatial arrangements (which is why the “method of loci” works well as a mnemonic device).

4.4. Mental rotation

One of the earliest and most cited results in the research on manipulating mental images is the “mental rotation” finding. Shepard and Metzler (1971) showed subjects pairs of drawings of three-dimensional figures, such as those illustrated in Figure 3, and asked them to judge whether the two objects depicted in the drawings were identical, except for orientation. Half the cases were mirror reflections of one another (or the 3D equivalent, called enantiomorphs), and therefore could not be brought into correspondence by a rotation. Shepard and Metzler found that the time it took to make the judgment was a linear function of the angular displacement between the pair of objects depicted.

This result has been universally interpreted as showing




Figure 3. Examples similar to those used by Shepard and Metzler (1971) to show “mental rotation.” The time it takes to decide whether two figures are identical except for rotation (a, b), or are mirror images (a, c), increases linearly as the angle between them increases.

that mental images of the objects are “rotated” continuously and at constant speed in the mind, and that this is, in fact, the means by which the comparison is made: We rotate one of the pair of figures until the two are sufficiently in alignment that it is possible to see whether they are the same or different. The phenomenology of the Shepard and Metzler task is clearly that we rotate the figure in making the comparison. I do not question either the phenomenology or the description that what goes on in this task is “mental rotation.” But there is some question about what these results tell us about the nature of mental images. The important question is not whether we can or do imagine rotating a figure, but whether we solve the problem by means of the mental rotation. For mental rotation to be a mechanism by which the solution is arrived at, its utility would have to depend on some intrinsic property of images. As an example, if it were the case that during mental rotation the figure moves *as a rigid form* through a continuum of angles, then mental rotation would be capitalizing on an intrinsic property of the image format.

Contrary to the general assumption, however, figural “rotation” is not a holistic process that operates on an entire figure, while the figure retains its rigid shape. Subjects in the original 3D rotation study (Shepard & Metzler 1971) examined both the target and the comparison figures together. In a subsequent study that monitored eye movements, Just and Carpenter (1976) showed that observers look back and forth between the two figures, checking for distinct features. This point was also made by Hochberg and Gellman (1977) who used simpler 2D figures and found that observers concentrate on significant milestone features when carrying out the task, and that, when such milestone features are available, there is no rotation effect. In studies reported in Pylyshyn (1979b), I showed that the apparent “rate of rotation” depends both on the complexity of the figure and on the complexity of the postrotation comparison task. (I used a task in which observers had to indicate whether a test figure, presented at various orientations, was embedded within the original figure.) The dependence of the rotation speed on such organizational and task factors shows that whatever is going on in this case does not consist in merely “rotating” one figure in a rigid manner into correspondence with the reference figure.

Even if the process of making the comparison in some sense involves the “rotation” of a represented shape, this tells us nothing about the form of the representation and does not support the view that the representation is pictorial. The proposal that a representation maintains its shape

because of the inherent rigidity of the image while it is rotated cannot be literally true, notwithstanding the phenomenology. The *representation* is not literally being rotated; no codes or patterns of codes are being moved in a circular motion. At most, what could be happening is that a representation of a figure is processed in such a way as to produce a representation of a figure at a slightly different orientation, and then this process is iterated (perhaps even continuously). There are probably good reasons, based on computational resource considerations, why the comparison process might proceed by iterating parts of a form over successive small angles (thus causing the comparison time to increase with the angular disparity between the figures) rather than attempt the “rotation” in one step. For example, Marr and Nishihara (1976) hypothesized what they called a primitive SPASAR mechanism, whose function was to compute the rotation of a simple dihedral vertex and determine its orthographic projections in a reference frame (a slightly different version that left out the details of the SPASAR mechanism, was later published in Marr & Nishihara 1978). This was an interesting idea that assumed a limited analogue operation that could be applied to one small feature of a representation at a time. Yet, the Marr and Nishihara proposal did not postulate a pictorial representation, nor did it assume that a rigid configuration was maintained by an image in the course of its “rotation.” It hypothesized a simple primitive operation on parts of a structured representation in a response to a computational complexity issue.

Like the paper folding task discussed earlier, the mental rotation phenomenon is robust and probably not cognitively penetrable, and is not a candidate for a straightforward tacit knowledge explanation (as I tried to make clear in Pylyshyn 1979b). Rather, the most likely explanation is one that appeals to the computational requirements of the task and general architectural (i.e., working memory) constraints, and therefore applies regardless of the form of the representation. No conclusions concerning the format of image representations, or the form of their transformation, follow from the rotation results. Indeed, these findings illustrate how treating the phenomenology as explanatory does not help us to understand why or how the behavior occurs.

5. Are images “depictive”?

5.1. Depiction and mandatory spatial properties of representations

It has frequently been suggested that images differ from structured descriptions in that the former stand in a special relationship to what they represent, a relationship referred to as *depiction*. One way of putting this is to say that in order to depict some state of affairs the representation needs to correspond to the spatial arrangement it represents the way that a picture does. One of the few people who have tried to be explicit about what this means is Stephen Kosslyn,⁴ so I quote him at some length:

A depictive representation is a type of picture, which specifies the locations and values of configurations of points in a space. For example, a drawing of a ball on a box would be a depictive representation. The space in which the points appear need not be physical, such as on this page, but can be like an array in a computer, which specifies spatial relations purely functionally.

That is, the physical locations in the computer of each point in an array are not themselves arranged in an array; it is only by virtue of how this information is “read” and processed that it comes to function as if it were arranged into an array (with some points being close, some far, some falling along a diagonal, and so on). In a depictive representation, each part of an object is represented by a pattern of points, and the spatial relation among these patterns in the functional space correspond to the spatial 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. . . . When a depictive representation is used, not only is the shape of the represented parts immediately available to appropriate processes, but so is the shape of the empty space. . . . Moreover, one cannot represent a shape in a depictive representation without also specifying a size and orientation. (Kosslyn 1994, #880, p. 5)

This quotation introduces a number of issues that need to be examined closely. One idea we can put aside is the claim that depictive representations convey meaning through their resemblance to the objects they depict. This relies on the extremely problematic notion of resemblance, which has been known to be inadequate as a basis for meaning (certainly since Wittgenstein 1953). Resemblance is neither necessary nor sufficient for something to have a particular meaning or reference: Images may resemble what they do not refer to (e.g., an image of John’s twin brother does not refer to John) and they may refer to what they do not resemble (an image of John taken through a distorting lens is still an image of John even though it does not resemble him).

Despite its obvious problems, the notion of resemblance keeps surfacing in discussions of mental images, in a way that reveals how deeply the conscious experience of mental imagery contaminates conceivable theories of mental imagery. For example, Finke (1989) begins with the observation, “People often wonder why mental images resemble the things they depict.” But the statement that images resemble things they depict is just another way of saying that the conscious experience of mental imagery is similar to the conscious experience one would have if one were to see the thing one was imagining. Consider what it would be like if images did not “resemble the things they depict”? It would be absurd if in imagining a table one had an experience that was like that of seeing a dog! Presumably this is because (a) what it means to have a mental image of a chair is that you are having an experience like that of seeing a chair, and (b) what conscious content your image has is something on which you are the final authority. You may be deceived about lots of things concerning your mental image. You may, and typically are, deceived about what sort of thing your image is (that is, what form and substance underlies it), but surely you cannot be deceived about what your mental image looks like, or what it resembles. These are not empirical facts about imagery, they are just claims about what the phrase “mental image” means.

In contrast to the vacuity of the criterion of resemblance, the proposal that images can be decomposed into “parts” with the spatial relations among parts of the image in some way mapping onto the parts and the spatial relationships among the corresponding parts of the world, deserves closer scrutiny although it has not received systematic treatment in the literature. This proposal is based on the assumption that in imagery there is a certain part-to-part homomorphism between the representation and the repre-

sented. Some time ago, Sloman (1971) suggested this as a defining characteristic of analogue representations and it is clearly an important criterion. Although it needs to be spelled out in more detail, this is a reasonable proposal, but it will not yield the conclusion that images are spatial in any sense that bears on the “depiction” story. In fact, it is true of any representational system that is compositional (see sect. 7.1).

Another proposal mentioned in the quotation is that in depictive representations certain aspects are mandatory; so that, for example, if you choose to represent a particular object, you cannot fail to represent its shape, orientation, and size. This claim too has some truth, although the question of which aspects are mandatory, why they are mandatory, and what this tells us about the form of the representation is not so clear. It is a general property of representations that some aspects tend to be encoded (or assigned as default value) if other aspects are. Sometimes this is true by virtue of what it is that you are trying to imagine. For example, you can’t imagine a melody without also imagining each note, and therefore making a commitment as to how many notes it has. This follows from what it means to “imagine a melody,” not from the inherent nature of some particular form of representation. The same is true for other examples of imaginings. When you ask someone to imagine a familiar shape by giving its name, say, the letter “B,” the person will make a commitment to such things as whether it is in upper or lower case. It seems as though you can’t imagine a B without imagining either an upper case “B” or a lower case “b.” But is this not another case of a requirement of the task to “imagine a ‘B’”? In this example, are you not being asked to describe what you would see if you were actually looking at a *token of a particular printed letter*? If you actually saw a token of a B you would see either a lower or an upper case letter, but not both and not neither. If someone claimed to have an image of a B that was noncommittal with respect to its case, you would surely be entitled to say that the person did not have a visual image at all.

In terms of other contents of an image, the situation gets murkier because it becomes less clear what exactly the task of imagining entails. For example, does your image of the letter “B” have to have a color or texture or shading? Must you represent the background against which you are viewing it, the direction of lighting and the shadows it casts? Must you represent it as viewed from a particular point of view? What about its stereoscopic properties; do you represent the changing parallax of its parts as you imagine moving in relation to it? Could you choose to represent any or none of these things? Most of our visual representations, at least in memory, are noncommittal in various respects (for examples see Pylyshyn 1978). In particular, they can be noncommittal in ways that no picture can be noncommittal. Shall we then say that they are not images? How you feel about such questions is more terminological (i.e., what you are disposed to count as an image representation) than empirical. It shows the futility of assuming that mental images are just like pictures. As the graphic artist M.C. Escher once put it,

a mental image is something completely different from a visual image, and however much one exerts oneself, one can never manage to capture the fullness of that perfection which hovers in the mind and which one thinks of, quite falsely, as something that is “seen.” (Escher 1960, p. 7)

5.2. Real versus functional space

Despite the temptation to do so, imagery theorists have been reluctant to claim that images are literally laid out in real space – that is, on a physical surface in the brain. However, because theories of imagery have had to appeal to such notions as distance, shape, size, and so on, some notion of space is always presupposed. Consequently, many writers who see the need for spatial properties speak of a “functional” space, with locations and other spatial properties being defined functionally (e.g., Denis & Kosslyn 1999). The example frequently cited (see the Kosslyn quotation above) is that of a matrix data structure in a computer, which can be viewed as having many of the properties of space without itself being laid out spatially in the physical machine. This is in some ways an attractive idea since it appears to allow us to claim that images have certain spatial properties without being committed to how they are implemented in the brain – so long as the implementation and its accessing operations function the way a real spatial system would function. The hard problem is to give substance to the notion of a functional space that does not reduce it to being either a summary of the data, with no explanatory mechanisms, or a model of real literal space. This problem has been so widely misunderstood that it merits some extended discussion.

Consider first why a matrix data structure might appear to constitute a “functional space.” As typically used, it seems to *have* two (or more) dimensions (since referencing individual cells is typically done by providing two numerical references or “coordinates”), to *have* distances (if we identify distance with the number of cells lying between two places), and to *have* empty spaces (so that it explicitly represents both where there are features and where there are no features). Graphical elements, such as points, contours, and regions can be represented by entering features into the cells at quantized coordinates. There is then a natural sense of the properties of “adjacency,” as well as of places being “between” two specified locations (as well as other simple geometrical properties of sets of features, such as being collinear, forming a triangle, and so on). Because of this, operations such as “scanning” from one feature to another, as well as of “shifting” and “rotating” patterns, can be given natural definitions (see, e.g., Funt 1980). Thus, the format of such a data structure appears to lend itself to being interpreted as “depictive” rather than “language-like,” as noted in the earlier Kosslyn quote.

Notice, however, that all the spatial notions mentioned in the previous paragraph are properties of a way of thinking about, or of interpreting, the data structure; they are not intrinsic properties of the matrix data structure itself. That is, what makes cells in a matrix appear to be locations with properties such as adjacency, between-ness, alignment, distance, and so on, is not any property of the matrix, *nor even of the way that this data structure must be used*. There is no sense in which any pairs of cells is special and so there is no natural sense in which some pairs of cells are “adjacent,” including a sense that derives from how they must be accessed. There are literally no constraints on the order in which cells must be accessed. We can, of course, require that the matrix be accessed in certain ways, and when we model imagery we typically do stipulate that certain pairs of cells be considered “adjacent” and that, in accessing any pair of cells in a serial fashion, certain other cells (the ones

we designate as being “between” the pair) must be visited first and in a certain order (which we call “scanning”). But it is critical to the interpretation of a computational process *as a model of mental imagery* that we understand exactly *why* such constraints hold. If our model of imagery assumed a literal physical surface, then the reason would be clear: *physical laws* require that movement over a surface follow a certain pattern, such as, that the time it takes to get from one place to another is the ratio of the distance traversed to the speed of movement. But in a matrix no such intrinsic constraint exists. Such a constraint must be stipulated as an extrinsic constraint (along with many other constraints, such as those that govern the invariance of adjacency, between-ness, or collinearity, with transformations of scale, orientation, and translation). The spatiality of a matrix, or of any other form of “functional space,” must be stipulated or assumed over and above any intrinsic property of the format of the representation. The crucial fact about extrinsic constraints is that *such constraints are independently stipulated*, and so *could be applied equally to any form of representation*, including a model of imagery that used symbolic expressions or structured descriptions. So far as the notion of functional space is concerned, there is nothing to prevent us from modeling the content of an image as a set of sentence-like expressions in a language of thought. We could then stipulate that in order to go from examining one place (referred to, say, by a unique name) to examining another place (also referred to by a name), you must pass through (or apply an operation to) the places whose names are located between the two names on some list. You might object that this sort of model is ad hoc. It is. But it is no more ad hoc than when the constraints are applied to a matrix formalism. Notice, moreover, that both become completely principled if they are taken to be simulations of a real spatial display.

You might wonder why the matrix feels more natural than other ways of simulating space. The answer may be that a matrix offers a natural model of space because we are used to thinking of and displaying matrices as two-dimensional tables (complete with empty cells) and of viewing the cells as being referenced by names that we think of as pairs of coordinates.⁵ We thus find it easy to switch back and forth between the data-structure view and the (physical) table view. Because of this, it is natural to interpret a matrix as a model of real space and therefore it is easy to make the slip between thinking of it on one hand as merely a “functional space,” and thinking of it, on the other hand, as a stand-in for (or a simulation of) real space – a slip we encounter over and over in theorizing about the nature of mental imagery. As a simulation of real space it is unproblematic so far as the sorts of problems discussed here are concerned. But we must recognize that in this case we *are assuming that images are written on a literal spatial medium*, which we happen to be simulating by a matrix (for reasons of convenience). In fact, in Kosslyn et al. (1979) this view is made explicit when the authors invoke what they call the “cathode-ray tube model.” In that case it is the literal space that has the explanatory force, notwithstanding the fact that, as a practical matter, it is being simulated on a digital computer.

The point is that there is no such thing as a “functional space” apart from the set of extrinsic stipulations or constraints we choose to impose on such things as how symbolic names (e.g., matrix coordinates) map onto places in a

physical display, and how distances and geometrical predicates are to be interpreted over the data structure. What we have, rather, is one of two things: either a real physical space, with its (approximately) Euclidean properties, or a symbolic model of such a space.⁶ Anything else is merely metaphoric and not explanatory. It allows one to think of an image as spatial without the attending disadvantages of having made an untenable assumption about the architecture of mental imagery.

The real scientific question is not how we can model space in a theory of mental imagery. Rather, it is whether there is any sense in which the architecture of mental imagery incorporates the geometry of real space. Only after we have answered this empirical question can we know whether one *should* model properties of space in modeling imagery. My purpose in belaboring the distinction between intrinsic and extrinsic constraints, and what is being presupposed when we talk of “functional space,” is simply to set the stage for the real issues, which are empirical. I have already described some of the relevant empirical findings in connection with the mental scanning and have suggested that the same is likely to be true for other findings that imply that images have metrical properties. The cognitive penetrability of such phenomena suggests that the mind does not work as though the imagery architecture imposes constraints like those you would expect of a real spatial display. It appears that we are not required to scan through adjacent places in getting from one place to another in an image – we can get there as quickly or as slowly as we wish, with or without visiting intermediate filled or empty places (assuming that visiting empty places is even possible – see sect. 6.4).

5.3. Projected mental images: Inheriting spatial properties from real space

In most imagery studies subjects are asked to imagine something while looking at a scene; thus, at least in some phenomenological sense, superimposing or projecting an image onto the perceived world. Yet, it has been amply demonstrated (O’Regan & Lévy-Schoen 1983) that true superposition of visual percepts does not occur when visual displays are presented in sequence, or across saccades. So what happens when a mental image (whether constructed or derived from memory) is superimposed over a scene? In many of these cases (e.g., Farah 1989; Hayes 1973; Podgorny & Shepard 1978) a plausible answer is that one allocates attention to the scene according to a pattern that corresponds roughly to the projected image. Alternatively, and more plausibly, one simply *thinks* of imagined objects as being located at places actually occupied by certain perceived ones. Thinking that something is at a certain location need not entail projecting an imagined shape onto some background. It might require nothing more than allocating attention to a particular object in a scene and thinking of that object as having a certain property. It is no more than thinking “this (e.g., referring to a bit of texture) is where I imagine feature F to be located.” The capacity for this sort of “demonstrative reference” has been investigated extensively and discussed by Pylyshyn (2000; 2001c).

Consider, for example, the study reported by Podgorny and Shepard (1978). In their vision control condition, experimenters asked subjects to indicate as fast as possible whether a dot appeared on, or beside, a simple figure, such




Figure 4. Observers were shown a figure (display 1) which they then had to retain as an image and indicate whether the dot (display 2) occurred on or off the imagined figure. The pattern of reaction times was found to be similar to that observed when the figure was actually present. (Based on Podgorny & Shepard 1978.)

as the letter F. The pattern of reaction times was recorded (e.g., times were found to be shorter when the dot was ON the figure than OFF, and shorter when it was at a vertex rather than mid-stroke, etc.). This pattern was then compared with the pattern obtained when the figure was merely imagined to be on the grid while a real dot was presented at corresponding locations (as shown in Fig. 4). In the image condition, the pattern of reaction times obtained was very similar to the one obtained from the corresponding real display. This was interpreted as showing that in both vision and projected imagery, subjects perceived a similar visual pattern. But a more parsimonious account is that in imagining the figure in this task, subjects merely attended to the rows and columns in which the imagined figure would have appeared. We know that people can indeed direct their attention to several objects such as rows or columns or cells in a display, or even to conform their attention to a particular shape. Focusing attention in this way is all that is needed in order to generate the observed pattern of reaction times. In fact, using displays similar to those used in the Podgorny and Shepard (1978) study, but examining the threshold for detecting spots of light, Farah (1989) showed that the instruction to simply *attend* to certain letter-shaped regions was more effective in enhancing detection in those regions than instructions to superimpose an image over the region.

A similar story applies to other tasks that involve responding to image properties when images are superimposed over a perceived scene (e.g., Hayes 1973). If, for example, you imagine the map used to study mental scanning (shown in Fig. 1) superimposed over one of the walls in the room, you can use the visual features of the wall to anchor various objects in the imagined map. You can think a thought which might be paraphrased as “the church is located where *this* (speck) is on the wall, the beach is beside *that* (corner) . . . ,” where each of the locative terms “this” and “that” picks out an object in the visual field and binds it to terms in the thought. Once the appropriate items are bound, “scanning the image” is accomplished by scanning between the selected items in the actual visual display. Thus, the increase in the time it takes to scan between items that are further apart on the imagined map is easily explained, since it involves scanning greater distances in the real scene. In general, such cases of superposition allow many of the spatial properties of the real scene (e.g., properties expressed by Euclidean and metrical axioms) to be inherited by the combined image-percept. For example, if image features A, B, and C are imagined to be collinear and they are bound to three actual collinear features in some

visible scene,⁷ then the fact that feature B is *between* A and C can be *visually* read off the perceived scene. The mechanism for indexing imagined objects to visual features, called *visual indexes*, has been described extensively elsewhere (e.g., Pylyshyn 1994b; 1998; 2000; 2001c). Such a binding is an example of the use of indexicals in image representations, and illustrates one way in which symbols underlying imagery can be what Barsalou (1999) calls perceptual symbols. This also appears to be what Glenberg and Robertson (2000) have in mind when they speak of image symbols as being grounded or embodied (this sort of grounding is discussed extensively in Pylyshyn 2001c).

Note that this vision-plus-indexes story is far from being equivalent to superimposing an image over the scene, because it assumes no pictorial properties of the “superimposed” image, only the binding of imagined objects/locations to real perceived ones. Moreover, because the relevant information involves only sparse spatial locations and not other detailed visual properties, the memory demand is minimal. So this process might conceivably be carried out even for a short time after the scene has been removed, as is the case when subjects are asked to close their eyes when they recall the image (there is evidence that short term recall of low-level iconic information can persist for several minutes; see Ishai & Sagi 1995). Thus this sort of story might provide another mechanism to explain phenomena such as mental scanning, even when scanning is carried out with eyes closed.

5.4. Visuomotor interaction with images

Another reason to think of mental images as being spatial is that they appear to connect with spatial aspects of the motor system in a way that is similar to how visually perceived space connects with the motor system. For example, it seems that you can imagine a scene and point to some object in the imagined scene, which suggests that the image is in some sense spatial. This apparent coordination with movements is surely one major reason to think of images as spatial, since motor control commands must be issued in spatial coordinates. If we could move our eyes to a place in an image or point to a feature of the image we would be exhibiting some aspect of the image’s spatiality.⁸

This raises the empirical question of whether images engage the motor system the way vision does. In a series of ingenious experiments, Finke (1979) showed that adaptation to displacing prisms could be obtained using only the imagined location of the observer’s hand as feedback. These studies are of special interest because they illustrate the way in which projected images can work like real percepts. Finke asked subjects to imagine that their (hidden) hand was in a certain specified location. The sequence of locations where he asked them to imagine their hand actually corresponded to the errors of displacement made by another subject who had worn displacing prisms. Finke found that both the pattern of adaptation and the pattern of after-effects shown by the subjects who were asked to imagine the location of their hand was similar to that exhibited by subjects who actually wore displacing prisms and so could see the displaced location of their hand.

It is known that adaptation can occur in response to cognitive factors, and indeed even to verbally presented error information (Kelso et al. 1975; Uhlarik 1973), though in that case the adaptation occurs more slowly and transfers

completely to the nonadapted hand. Yet, Finke found that in the case of imagined hand position, the adaptation, though significantly lower in magnitude, followed the pattern observed with the usual visual feedback of hand position. Moreover, when subjects were told that their hand was actually not where they imagined it to be, the adaptation effect was nonetheless governed by the imagined location, rather than by where they were told their hand was, and followed the same pattern as that observed with both imagery and actual visually presented error information.⁹ It thus seems that the visuomotor system may be involved when adapting to imagined hand positions. The question is: What is the nature of this involvement?

Exactly what causes prism adaptation has been the subject of some debate (Howard 1982), but it is generally accepted that an important factor is the discrepancy between the seen position and the felt position of the hand (or the discrepancy between visual and kinesthetic or proprioceptive location information). Significantly, such discordance does not require that the visual system recover any visual property of the hand *other than its location*. Indeed, in some studies of adaptation, subjects viewed a point source of light attached to their hand rather than the hand itself (Mather & Lackner 1977), with little difference in adaptation. But exactly where the subject attends is important (Canon 1970; 1971). In some cases, even an immobile hand can elicit adaptation provided that the subject visually attends to it (Mather & Lackner 1981). In Finke’s experiments, subjects focused their gaze towards a particular location, where they were, in effect, told to pretend (incorrectly) that their hand was located, thus focusing attention on the discordance between this imagined location of their hand and their kinesthetic and proprioceptive sense of the position of their arm.¹⁰ Thus, the imagery condition in these studies provides all that is needed for adaptation – without requiring any assumptions about the nature of imagery.

It seems that there are a number of imagery-motor phenomena that depend only on orienting one’s gaze or one’s focal attention to certain perceived locations. The Finke study of adaptation of reaching is a plausible example of this sort of phenomenon, as is the Tlauka and McKenna (1998) study of S-R compatibility for manually responding to information in images. None of these results require that imagery feed into the visuomotor system, let alone that images be spatial or depictive. Indeed, these cases involve *actual* visual perception of location (i.e., there really are visible elements at the relevant locations in real space that are being attended). Since the adaptation phenomena (as well as the S-R compatibility phenomena) require only location information and no pictorial information (e.g., shape, color, etc.), they do not in any way implicate the “depictive” character of mental images.

When we look at cases in which images are not projected onto a perceived scene (say, they are imagined in the dark or with eyes closed), or in which more than just the imagined location of an object is relevant to the motor action, we find that images do not interact with the perceptual motor system in a way that is characteristic of visual interaction with that system. To show this we need to look at certain signature properties of the visual control of movements, rather than at cases where the control may actually be mediated only by spatial attention. One clear example of strictly visual control of motor action is smooth pursuit.

People can track the motion of slowly moving objects with a characteristic sort of eye movement called *smooth pursuit*. There are also reports that under certain circumstances people can track the voluntary (and perhaps even involuntary) movement of their hand in the dark (Mather & Lackner 1980). They can also track the motion of objects that are partially hidden from view (Steinbach 1976), and even of induced (apparent) motion of a point produced by a moving frame surrounding the point (Wyatt & Pola 1979). In other words, they can smoothly pursue inputs generated by the early vision system. Yet, what people cannot do is smoothly pursue the movement of imagined objects. In fact, it appears to be impossible to voluntarily initiate smooth pursuit tracking without a moving stimulus (Kowler 1990).

Another example of a characteristic visually guided control is reaching to grasp a visible object. Although we can reach out to grasp imagined objects, when we do so we are essentially *pantomiming* a reaching movement, rather than engaging the visuomotor system. Visually guided reaching exhibits certain quite specific trajectory properties not shared by pantomimed reaching (Goodale et al. 1994). For example, the time and magnitude of peak velocity, the maximum height of the hand, and the maximum grip aperture, are all significantly different when reaching to imagined than to perceived objects. Reaching and grasping gestures towards imagined objects exhibit the distinctive pattern that is observed when subjects are asked to pantomime a reaching and grasping motion. There is considerable evidence that the visuomotor system is itself encapsulated (Milner & Goodale 1995) and, like the visual system, is able to respond only to information arriving from the eyes, which often includes visual information that is not available to consciousness. As with the visual system, only certain limited kinds of modulations of its characteristic behavior can be imposed by cognition. When we examine signature properties of the encapsulated visuomotor system, we find that mental images do not engage it the way that visual inputs do.

6. Are images “seen” by the visual system?

One of the most actively pursued questions in contemporary imagery research has been the question of whether mental imagery uses the visual system. Intuitively the idea that imagery involves vision is extremely appealing since the experience of imagery is phenomenally very like the experience of seeing – indeed, there have been (disputed) claims that when real perception is faint because of impoverished stimuli, vision and imagery may be indistinguishable (Perky 1910). But, from the perspective of the present thesis, there is a more interesting reason for asking whether the visual system is involved in examining images. If vision is used to interpret mental images, it might support the idea that images are things that can be *seen*, thus lending credence to the intuitive view of images as pictorial (or depictive). The question of the overlap between imagery and vision has been investigated with particular zeal within the cognitive neuroscience community. We shall examine the neuroscience findings in section 7. For the present I want to consider some of the psychological evidence that suggests an overlap between imagery and visual perception, and to ask whether this evidence supports the view that im-

ages are pictorial objects that are “seen” by the visual system.

6.1. The experience of seeing and of imagining

It may well be that the most persuasive reason for believing that mental imagery involves the visual system is the subjective one: Mental imagery is accompanied by a subjective experience that is very similar to that of seeing. As I remarked at the beginning of this article, this sort of phenomenal experience is very difficult to ignore. Yet it is quite possible that both vision and imagery lead to the same kind of experience because the same symbolic, rather than pictorial, form of representation, underwrites them both. The experience might correspond to a fairly high level of the analysis of the visual information, say, at the point where the stimulus is recognized as something familiar (e.g., Crick & Koch 1995, and Stoerig 1996, suggest that the locus of our awareness occurs higher than primary visual cortex). In that case, even though imagery and vision shared common mechanisms and forms of representation, one could not infer that the form was depictive or pictorial. At most one might conclude, as does Barsalou (1999), that the form of representation underlying images, while symbolic in nature, is also modality-specific, inasmuch as it consists of some subset of the neural activity that is associated with the corresponding visual perception. This alternative is compatible with my proposal (in the next section) that the representation underlying vision and visual imagery may use the same modality-specific symbolic vocabulary.

6.2. Interference between imaging and visual perception

One of the earliest findings that persuaded people that images involve the visual system was that the task of examining images could be disrupted by a subsidiary visual or spatial task. For example, Brooks (1968) showed that reporting spatial properties from images is more susceptible to interference when the response must be given by a spatial method (e.g., pointing) than by a verbal one (i.e., speaking). When subjects were asked to describe the shape of the letter **F** by providing a list of the right and left turns one would have to take in traveling around its periphery, their performance was worse when they had to point to the left or right (or to left- and right-pointing arrows) than when they had to say the words “left” and “right.” Segal and Fusella (1969; 1970) subsequently confirmed the greater interference between perception and imagery in various same-modality tasks and also showed that both sensitivity and response bias measures (derived from Signal Detection Theory) were affected. Segal and Fusella concluded that “imagery functions as an internal signal which is confused with the external signal” (p. 458). This may be the correct conclusion to draw; but it does not show that either of the inputs is pictorial. All that it implies is that the same type of representational *contents* are involved, or to put it another way, that the same concepts are deployed. For the sake of argument, think of the representations in these studies as being in a common language of thought: What, in that case, do the representations of visual patterns have in common with mental images of visual patterns? One obvious commonality is that they are both about the appearance of visual patterns. Like sentences about visual appearances, they all involve the use of concepts such as “bright,” “red,” “right

angle,” “parallel to,” and so on. It is not surprising that two responses requiring the same modality-specific conceptual vocabulary would interfere. Thus, it may be that visual percepts and visual images interact because both consist of symbolic representations that use some of the same proprietary spatial or modality-specific vocabulary. That the linguistic output in the Brooks study is not as disruptive as pointing, may simply show that spatial concepts are not relevant to *articulating* the words “left” or “right” once they have been selected for uttering, whereas these concepts *are* relevant to issuing the motor commands to move left or right.

6.3. Visual illusions induced by superimposing mental images

Other studies that are cited in support of the view that images are interpreted by the visual system are ones showing that projecting images of certain patterns onto displays creates some of the well-known illusions, such as the Müller-Lyer illusion, the Pogendorf illusion or the Herring illusion, or even the remarkable McCollough effect.¹¹ For example, Bernbaum and Chung (1981) showed subjects displays such as those in the top part of Figure 5. Subjects were asked to imagine lines connecting the endpoints of the visible line to either the outside or the inside pairs of dots in this display (when the endpoints are connected to the inside pair of dots they produce outward-pointing arrows, and when they are connected to the outside pair of dots they produce inward pointing arrows, as in the original Müller-Lyer illusion). Bernbaum and Chung (also Ohkuma 1986) found that imagining the arrows also produced the illusion, with the inward pointing arrows leading to the perception of a longer line than the outward pointing arrows. For the sake of argument, let us take these results as valid, notwithstanding the obvious susceptibility of such findings to experimenter demand effects – see note 11.

Before one can interpret such findings one needs to understand why the illusion occurs in the visual case. Explanations for the Müller-Lyer and similar illusions tend to fall into one of two categories. They either appeal to the detailed shapes of contours involved and to the assumption that these shapes lead to erroneous interpretations of the pattern in terms of 3D shapes, or they appeal to some general characteristics of the 2D envelope created by the dis-




Figure 5. Figures used to induce the Müller-Lyer illusion from images. Imagine the end points being connected to the inner or the outer pairs of dots in the top figure (Bernbaum & Chung 1981), or selectively look at the inward or outward arrows in the bottom figure (based on Goryo et al. 1984).

play and the consequent distribution of attention or eye movements. The former type of explanation, which includes the “inappropriate constancy scaling” theory attributed to Richard Gregory (1965), has not fared well in general since the illusion can be produced by a wide variety of types of line-endings (see the review in Nijhawan 1991). The latter type of explanation, which attributes the illusion to the way attention is allocated and to mechanisms involved in preparing eye-movement, has been more successful. For example, one theory (Virsu 1971) appeals to the tendency to move one’s eyes to the center of gravity of a figure. The involvement of eye movements in the Müller-Lyer illusion has also been confirmed by Bolles (1969), Coren (1986), Festinger et al. (1968), Hoenig (1972), and Virsu (1971). Another example of the envelope type of theory is the framing theory (Brigell et al. 1977; Davies & Spencer 1977), which uses the ratio of overall figure length to shaft length as predictor. Such envelope-based theories have generally fared better than shape-based theories, not only on the Müller-Lyer illusion, but also for most cases in which there are context effects on judgments of linear extent. What is important about this for our purposes is that these explanations do not appeal to pattern-perception mechanisms and therefore are compatible with attention-based explanations of the illusions.

Further evidence that attention can play a central role in these illusions (as well as in other visual illusions induced by mental imagery, e.g., Wallace 1984a; 1984b) comes from studies that actually manipulate attention focus. For example, it has been shown (Goryo et al. 1984) that if both sets of inducing elements (the outward and inward arrowheads) were present (as in the bottom part of Fig. 5), observers could selectively attend to one or the other and obtain the illusion appropriate to the one to which they attended. This is very similar to the effect demonstrated by Bernbaum and Chung (1981), but without requiring that any image be superimposed on the line. Coren and Porac (1983) also confirmed that attention alone could create, eliminate, or even reverse the Müller-Lyer illusion. Attention-mediation was also shown explicitly in the case of an ambiguous motion percept (Watanabe & Shimojo 1998). This is in keeping with the evidence we considered in section 5.3 showing that in many cases in which mental images are interpreted as having visual effects, the effect can be explained by appeal to the attention-focusing role that imagery plays in connection with visual perception. Finally, the relevance of the imagined induction of the Müller-Lyer and similar illusions to the picture theory is further cast into doubt when one recognizes that such illusions, like many other imagery-based phenomena, also appear in congenitally blind people (Patterson & Deffenbacher 1972).

6.4. Imagined versus perceived motion

Gilden et al. (1995) used visual motion adaptation to study whether the visual system is involved in imagery. This study is of special interest to us since motion adaptation is known to be retinotopic, and therefore occurs in the early visual system. When a region of the visual field receives extensive motion stimulation, an object presented in that region is seen to move in the opposite direction to the inducing movement (this is called the “waterfall illusion”) and a moving object is seen as moving more slowly. Gilden et al. designed their study with the intention of showing that the

motion of an imagined object is affected by the aftereffect of a moving field. They had subjects gaze for 150 seconds at a square window on a screen containing a uniformly moving random texture. Then they showed subjects a point moving towards that window and disappearing behind what appeared to be an opaque surface, and they asked subjects to imagine the point continuing to move across the previously stimulated region and to report when the point would emerge at the other side of the surface. Gildea et al. did find an effect of motion adaptation on imagined motion, but it was not exactly the effect they had expected. They found that when the point was imagined as moving in the same direction as that of the inducing motion field (i.e., against the motion aftereffect), it appeared to slow down (it took longer to reach the other side of the region). However, when the point was imagined as moving in the opposite direction to the inducing motion field (i.e., in the same direction as the motion aftereffect), the point appeared to speed up (it reached the other side in a shorter time).

The latter effect is *not* what happens with real moving points. In visual motion adaptation, motion appears to slow down no matter which direction the inducing motion field moves, presumably because all motion sensitive receptors have been fatigued. But, as Gildea et al. recognized, the effect they observed is exactly what one would expect if, rather than imagining the point moving continuously across the screen, subjects imagined the point as being located at a series of static locations along the imagined path. This suggests a quite different mechanism underlying imagined motion when the latter is generated as the extrapolation of perceived motion. We know that people are very good at computing time-to-contact of a uniformly moving object at a specified location (e.g., DeLucia & Liddell 1998). What may be going on in imagined motion is that people may simply pick out one or more marked places (e.g., elements of texture) along the path, using the visual indexing mechanism discussed earlier, and then compute the time-to-contact for each of these places.

We explicitly tested this idea (Pylyshyn & Cohen 1999) by asking subjects to extrapolate the motion of a small square, which disappeared by occlusion behind an apparent opaque surface. They were asked to imagine the smooth motion of the square in a dark room. At some unpredictable time in the course of this motion the square would reappear, as though coming out through a crack in the opaque surface, and then recede back through another crack, and subjects had to indicate whether this reappearance occurred earlier or later than when their imagined square reached that crack. This task was carried out in several different conditions. In one condition, the location of the “cracks” where the square would appear and disappear was unknown (i.e., the cracks were invisible). In another condition, the location at which the square was to appear was known in advance: it was indicated by a small rectangular figure that served as a “window” through which, at the appropriate time, subjects would briefly view the square that was moving behind the surface (the way the squares appeared and disappeared in the window condition was identical to that in the no-window condition except that the outline of the window was not visible in the latter case). And finally, in one set of conditions the imagined square moved through total darkness, whereas in the other set of conditions the path was marked by a sparse set of dots that could be used as reference points to compute time-to-contact. As

expected, the ability to estimate where the imagined square was at various times (measured in terms of decision time) was significantly improved when the location was specified in advance, and also when there were visible markers along the path of imagined motion.

Both of these findings confirm the suggestion that what subjects are doing when they report “imagining the smooth motion of a square,” is selecting places at which to compute time-to-contact and then merely thinking that the imaginary square is at those places at the estimated times. According to this view, subjects are thinking the thought “now it is *here*” repeatedly for different visible objects (picked out by the visual indexing mechanism mentioned earlier), and synchronized to the independently computed arrival times. This way of describing what is happening requires neither the assumption that the visual system is involved (other than the attentional indexing mechanism), nor the assumption that an imagined square is actually moving through some mental space and occupying each successive position along a real spatial path. Indeed, there is no need to posit any sort of space except the visible one that serves as input to the time-to-contact computation.

6.5. Extracting novel information from images: Visual (re)perception or inference?

One of the alleged purposes of mental images is that they can be examined in order to visually discover new properties or new interpretations or reconstructions. It would therefore seem important to ask whether there is any evidence for such visual reconstructions. This question turns out to be more difficult to answer than one might have expected, for it is clear that by examining images one can draw *some* conclusions that were not explicitly present in, say, the verbal description under which it was imagined. So if I ask you to imagine a square and then to imagine drawing in both diagonals, it does not seem surprising if you can tell that the diagonals cross or that they form an “X” shape. Since this inference is very simple it does not seem *prima facie* to qualify as an example showing that images are interpreted *visually*. On the other hand, suppose I ask you to imagine two parallelograms, one directly above the other on a page, and then to connect each vertex of the top one to the corresponding vertex of the bottom one. What do you see? As you keep watching, what happens in your image? When presented visually, this figure leads to certain phenomena that do not appear in mental imagery. The signature properties of spontaneous perception of certain line drawings as depicting three-dimensional objects and spontaneous reversals of ambiguous figures do not appear in this mental image (which happens to be a Necker Cube).

But what counts, in general, as a visual interpretation as opposed to an inference? I doubt that this question can be answered without a sharper sense of what is meant by the term “visual.” Since the everyday (pretheoretical) sense of the notion of “vision” clearly involves most, if not all, of cognition, the question of the involvement of vision in reconstructions cannot be pursued without a more restricted sense of what counts as a visual phenomenon. Clearly, deciding whether two crossed lines form an “X” is not one of these phenomena, nor is judging that when a D is rotated counterclockwise by 90 degrees and placed on top of a J the result looks like an umbrella. You don’t need to use the early visual system in deciding that. All you need is an elemen-

tary inference based on what makes something look like an umbrella (e.g., it has an upwardly convex curved top attached below to a central vertical stroke – with or without a curved handle at the bottom). Thus, these the sorts of examples (which were used in Finke et al. 1989), cannot decide the question of whether images are *visually* (re)interpreted.

Hochberg (1968) suggested a few signature properties of vision, including spontaneous interpretation of certain line drawings as three-dimensional objects, spontaneous reversal of ambiguous shapes, and spontaneous interpretation of certain sequences as apparent motion. Although these criteria cannot always be used to decide whether a particular interpretation is visual, they do indicate the sort of constructs that appear to be constitutive of early vision. For the time being, all we can do is ask whether the candidate reconstrual meets this sort of intuitive condition. A more thorough analysis would attempt to look at converging evidence concerning the level of the visual system implicated.

The clearest evidence I am aware of that bears on the question of whether images are subject to *visual* reconstruals, is provided by studies carried out by Peter Slezak (1991; 1992; 1995). Slezak asked subjects to memorize pictures such as those in Figure 6. Then he asked them to mentally rotate the images clockwise by 90 degrees and to report what they looked like. None of his subjects was able to report the appearance (of the mentally rotated shapes) that they could easily report by rotating the actual pictures. The problem was not with their recall or even their ability to rotate the simple images; it was with their ability to recognize the rotated image in their mind's eye. We know this because subjects were able to draw the figures from memory, and when they rotated those, they did see the other construals. What is special about these examples is that the resultant appearance is so obvious – it comes as an “aha!” experience when carried out by real rotation. Unlike the figures used by Finke et al. (1989), these shapes were not familiar and their appearance after the transformation (by rotation) could not be easily inferred from their representation.

A related question that can be asked is whether images can be ambiguous, since this also concerns the question of whether images can be visually reinterpreted. This case, however, presents some additional methodological problems. Not all ambiguities contained in pictures are visual ambiguities, just as not all reinterpretations are visual reconstruals. For example, the sorts of visual puns embodied in some cartoons or, most characteristically, in so-called “doodles” (see <http://www.doodles.com> for examples), rely on ambiguities; but clearly not on ones that are based in part on different visual organizations being produced by



Figure 6. Orientation-dependent figures used by Slezak (1991, 1995). To try these out, memorize the shape of one or more of the figures, then close your eyes and imagine them rotated clockwise by 90 degrees (or even do it while viewing the figures). What do you see? Now try it by actually rotating the page.

early visual processes. By contrast, the reversal of figures such as the classical Necker Cube is at least in part the result of a reorganization that takes place in early vision. Do such reorganizations occur with visual images? In order to answer that question we would have to control for certain obvious alternative explanations of any reports of apparent reinterpretations. For example, if a mental image appeared to reverse, it might be because the observer knew of the two possible interpretations and simply replaced one of its interpretations with the other. This is the view that many writers have taken in the past (Casey 1976; Fodor 1981).

Chambers and Reisberg (1985) were the first to put the question of possible ambiguous mental images to an empirical test. They reported that no reversals or reinterpretations of any kind took place with mental images. Since that study was reported there have been a series of studies and arguments concerning whether images could be visually (re)interpreted. Reisberg and Chambers (1991) and Reisberg and Morris (1985) used a variety of standard reversal figures and confirmed the Chambers and Reisberg finding that mental images of these figures could not reverse. Finke et al. (1989) took issue with these findings, citing their own experiments involving operations over images (mentioned above); but, as I suggested, it is dubious that the reinterpretation of their simple familiar figures should be counted as a *visual* reinterpretation. Even if these could be considered so, the more serious problem is to explain why clear cases of visual interpretations, such as those studied by Chambers and Reisberg, do not occur with images.

A more direct and extensive exploration of whether mental images can be ambiguous was undertaken by Peterson (1993; Peterson et al. 1992), who argued that certain kinds of reconstruals of mental images do take place. Peterson first distinguished different types of image reinterpretations. In particular, she distinguished what she calls *reference-frame realignments* (in which one or more global directions are reassigned in the image, as in the Necker cube or rabbit-duck ambiguous figures), from what she calls *reconstruals* (in which reinterpreting the figure involves assigning new meaning to its parts, as in the wife/mother-in-law or snail/elephant reversing figures). We will refer to the latter as *part-based reconstruals*, to differentiate them from other kinds of reconstruals (since their defining characteristic is that their parts take on a different meaning). A third type, figure-ground reversal (as in the Rubin vases), was acknowledged to occur rarely if ever with mental images (a finding that was also confirmed by Slezak 1995, using quite different displays). Among her findings, Peterson showed that reference-frame realignments do not occur in mental images unless they are cued by either explicit hints, or implicit demonstration figures, whereas some part-based reconstruals occurred with 30% to 65% of the subjects.

Recall that our primary concern is not whether any reinterpretations occur with mental images. The possibility of some reinterpretation depends upon what information or content-cues are contained in the image, which is orthogonal to the question of the mechanisms used in processing it. What we are concerned with is whether the format of images is such that their interpretation and/or reinterpretation involves the specifically visual (i.e., the early vision) system, as opposed to the general inference system. The crucial question, therefore, is how Peterson's findings on reinterpreting mental images compare with the reinterpretations observed with ambiguous visual stimuli. The answer

appears to be that even when reinterpretations do occur with mental images, they are qualitatively different from those that occur with visual stimuli. For example, Peterson (1993) showed that, whereas reference-frame reversals are dominant in vision, they are rare in mental imagery; while the converse is true for part-based reconstructions. Also, the particular reconstructions observed with mental images tend to be different from those observed with the corresponding visual stimuli. Visual reconstructions tend to fall into major binary categories – in the case of the figures used by Peterson et al. these are the duck-rabbit, or the snail-elephant categories. On the other hand, in the imagery case subjects provided a large number of other interpretations (which, at least to this observer, did not seem to be clear cases of distinctly different appearances – certainly not as clear as the cases of the Necker Cube reversal or even the reconstructions observable when the shapes in Fig. 6 are rotated). The number of subjects showing part-based reconstructions with mental images dropped by half when only the particular interpretations observed in the visual case were counted. Reinterpretation of mental images is also highly sensitive to hints and strategies, whereas there is reason to doubt that the early stages of vision are sensitive to such cognitive influences (Pylyshyn 1999).

The reason for these differences between imagery and vision is not clear, but they add credence to the suggestion that what is going on in the mental image reconstructions is not a perceptual (re)interpretation of a generated picture, but something else. Perhaps what is going on is inference and memory retrieval based on shape properties – the sort of process that goes on in the decision stage of vision, after early vision has generated shape-descriptions. This is the stage at which beliefs about the perceived world are established so we expect it to depend on inferences from prior knowledge and expectations, like all other cases of belief fixation. It seems quite likely that parts of the highly ambiguous (though not clearly bistable) figures used by Peterson et al. might serve as cues for inferring or guessing at the identity of the whole figure (for illustrations of these figures, see Peterson 1993). Alternatively, as suggested earlier, several possible forms might be computed by early vision (while the figures were viewed) and stored, and then, during the image-recall phase, a selection might be made from among them based on a search for meaningful familiar shapes in long-term memory. While in some sense all of these are reinterpretations of the mental images, they do not all qualify as the sort of visual “reconstructions” of images that show that mental images are pictorial entities whose distinct perceptual organization (and reorganization) is determined by the early vision system. Indeed, they seem more like the kind of interpretations one gets from Rorschach inkblots.

7. Can evidence from neuroscience settle the question?

7.1. Searching for the “mind’s eye” and the “image” in the brain

The neuroscience research concerned with mental imagery has been devoted primarily to attempting to show that the early visual system is involved in this process. The involvement of visual mechanisms in mental imagery is of interest to the picture theorists primarily because of the possibility

that the particular role played by the visual system in processing mental images will vindicate a version of the picture theory, by showing that imagery does indeed make use of a special sort of spatial display (this is explicitly the claim in Kosslyn 1994). The question that naturally arises is whether we can make a case for this view by obtaining evidence concerning which areas of the brain are involved in mental imagery and in visual perception. Before examining the evidence, it is worth reiterating that the question of the involvement of the visual system and the question of the form of mental images are largely independent questions. It is logically possible for the visual system to be involved in both vision and mental imagery and yet in neither case generate picture-like depictive representations. Similarly, it is possible for representations to be topographically organized in some way and yet have nothing to do with visual perception, nor with any depictive character of the representation.

In a certain sense the physical instantiation of any adequate cognitive representation *must* be topographically organized. Fodor and I (Fodor & Pylyshyn 1988) have argued that any form of representation that is adequate as a basis for cognition must be compositional, in the sense that the content of a complex representation must derive from the content of its parts and the rules by which the complex is put together (i.e., the way the meaning of sentences is compositional and depends on the meaning of its parts, together with how they are syntactically put together). The physical instantiation of any representation that meets the requirement of compositionality will itself be compositional (see Pylyshyn 1984, pp. 54–69; 1991b). In the case of symbolic representations, parts of expressions are mapped recursively onto parts of physical states and syntactic relations are mapped onto physical relations. As a result, there is a very real sense in which the criteria in the Kosslyn quotation at the beginning of section 5.1 are met by *any* adequately expressive physical symbol system, not just a depictive one. Note that in a digital computer, data-structure representations are compositional and topographically distributed, and yet are generally not thought to be depictive, whereas when they are supposed to be depictive, as when they encode images (say as JPEG or TIFF files), their topographical distribution generally does not mirror the physical layout of the picture. Consequently, the question of the spatial distribution of images, the question of whether they are depictive, and the question of whether they are connected with vision are logically independent questions.

Notwithstanding the logical independence of the question of the format of images and the question of the involvement of the visual system, the following line of reasoning continues to hold sway in the neuroscience literature (Kosslyn et al. 1995; 1999a). Primary visual cortex (Area 17) appears to be organized retinotopically (at least in the monkey brain). So if Area 17 is activated when subjects generate mental images, then it must be that higher-level cortical centers are generating retinotopically-organized activity in the visual system. In other words, during mental imagery a spatial pattern of activity is generated in the same parts of the visual system where such activity occurs in vision. From which it follows (so the argument goes) that a spatial or “depictive” form of activity is laid out in the cortex, much as it is on the retina, and this activity is then interpreted (or “perceived”) by the visual system. This line of reasoning is very much in keeping with the views held by proponents of the subjectively satisfying picture theory. It is no surprise then

that those who hold the picture view welcome any evidence of the involvement of early vision in mental imagery.

With this as background, we can ask whether there is evidence for the involvement of early, topographically organized areas of vision cortex in mental imagery and, if so, whether the involvement is of the right kind. Some evidence has been reported showing that mental imagery involves activity in areas of striate cortex associated with vision. Most of this evidence comes from studies using neural imaging to monitor regional cerebral blood flow.¹² While some neural imaging studies report activity in topographically organized cortical areas (Kosslyn et al. 1995; 1999a), most have reported that only later visual areas, the so-called visual association areas, are active in mental imagery (Charlot et al. 1992; Cocude et al. 1999; D'Esposito et al. 1997; Fletcher et al. 1996; Goldenberg et al. 1995; Howard et al. 1998; Mellet et al. 1996; 1998; Roland & Gulyas 1994b; 1995; Silbersweig & Stern 1998); but see the review in Farah (1995b) and the some of the published debate on this topic (Farah 1994; Roland & Gulyas 1994a, 1994b). Other evidence comes from clinical cases of brain damage and is even less univocal in supporting the involvement in mental imagery of the earliest, topographically organized areas of visual cortex (Roland & Gulyas 1994b). There is some reason to think that the activity associated with mental imagery occurs at many loci, including higher-levels of the visual stream (Mellet et al. 1998).

Despite the weight placed on neural imaging studies by proponents of the picture theory, the involvement of visual areas of cortex – even of topographically organized areas of early vision – in mental imagery would not itself support a cortical display view of imagery. In order to support such a view, it is important not only that such topographically organized areas be involved in imagery, but also that their involvement be of the right sort – that the way their topographical organization is involved reflects the spatial properties of the image, particularly as the latter is experienced and as it is assumed to function in accounting for the many imagery findings in the literature. Very few neuroscience studies meet this criterion, even when they show that early visual areas are activated during mental imagery. That is because such studies generally do not provide any evidence concerning *how* spatial information is mapped in the visual cortex. One of the few examples of a study that does address this question was reported in Kosslyn et al. (1995). Unlike most neural imaging studies, this one relates a specific spatial property of a phenomenal mental image (its size) to a *pattern* of neural activity. Thus it behooves us to look at the findings in detail.

The Kosslyn et al. (1995) study showed that “smaller” mental images (mental images that the observer subjectively experiences as occupying a smaller portion of the available “mental display”) are associated with more activity in the posterior part of the medial occipital region, while “larger” images are associated with more activity in the anterior parts of the region. Since this pattern is similar to the pattern of activation produced by small and large retinal images, respectively, this finding was taken to support the claim that not only does activation of visual cortical areas occur during mental imagery, but also that the *form* of the activation is the same as that which derives from vision. This, in turn, is interpreted as showing that imagery creates a cortical display which maps represented space onto cortical space. Because of this, Kosslyn et al. (1995, p. 496) feel

entitled to conclude that the findings “indicate that visual mental imagery involves ‘depictive’ representations, not solely language-like descriptions.”

But notice that even if the cortical activity monitored by PET scans corresponded to a mental image, the evidence only showed that a larger mental image involves activity that is located where parts of larger retinal images tend to project – that is, at best they show that larger mental images involve activity in *different locations* in the cortex than the activity associated with smaller mental images. We have good reason to believe that the reason larger retinal images activate the regions they do (in the case of vision) is because of the way that the visual pathway projects from the periphery of the retina to the occipital cortex (Fox et al. 1986). This pattern of activation does *not* map the size of the image onto a metrical spatial property of its cortical representation. In particular, it is important that the data do not show that image size is mapped onto the *size of the active cortical region*, as would be required by the cortical display view, and as would be required if images had the property that they “preserve metrical information,” as claimed by Kosslyn and others. It is also not the pattern that is required in order to account for the imagery data reviewed earlier. For example, the explanation of why it takes less time to notice features in a large image is supposed to be that it is easier to discern details in a *large cortical display*, not that the image is located in the more anterior part of the medial occipital cortex. The property of being located in one part of the visual cortex rather than another simply does not bear on any of the behavioral evidence regarding the spatial nature of mental images discussed earlier. Consequently, the finding cannot be interpreted as supporting the cortical display view of mental imagery, nor does it in any way help to make the case for a literal-space view or for the picture theory.

Those of us who eschew dualism naturally assume that *something* different happens in the brain when a different phenomenal experience occurs. Consequently, something different must occur in the brain when a larger image is experienced (this is called the supervenience assumption and few scientists would dispute it). The point has never been to question materialism, but only to question the claim that the content of an image maps onto the brain in a way that helps explain the imagery results (e.g., mental scanning times, image-size effects, etc.) and perhaps even the subjective content of mental images. Discovering that a larger phenomenal image mapped onto a larger region of brain activity might have provided some support for this view, since it might at least have suggested a possible account for such findings as that it takes longer to scan larger image distances and that it takes longer to see details in smaller images. But finding that a larger mental image merely activated a different area of the brain is no help in this regard. Incidentally, while the data tend to confirm that the pattern of cortical activity changes in similar ways for similar differences in perceived and imagined patterns, this does not support the existence of a cortical display in *either* imagery or vision. And that is as we would expect, given that there is no evidence that an extended display is involved in visual processing beyond the retina and its primary projection (but see Note 14). The hope of enlisting neuroscience to provide support for a picture theory of mental imagery by showing an overlap between vision and imagery rests, in the first instance, on the acceptance of a false theory of vision.

A somewhat different kind of evidence for the neural basis of image size was reported by Farah et al. (1992), based on a clinical case. Farah et al. reported that a patient who developed tunnel vision after unilateral occipital lobectomy also developed tunnel imagery. Once again this may well show that the two cases of tunnel vision have a common underlying neural basis, and that this basis may be connected with the visual system. But it does not show that this basis has anything to do with a topographical mapping of the spatial property of images onto spatial properties of a neural display. I might also point out that other explanations for the Farah et al. finding are possible which do not assume that the surgery resulted in damage to the imagery system. If, as suggested earlier, many of the properties of mental images actually arise from the implicit task requirement of simulating what things would look like in the corresponding perceptual situation, then a plausible explanation for the Farah et al. finding is that the patient was simply demonstrating a tacit knowledge of what the post-surgery world looked like to her. In the Farah et al. (1992) study, the patient had nearly a year of post-surgery recovery time before the imagery testing took place. During this time she would certainly have become familiar with what things *looked like* to her, and would therefore have been in a position to simulate her visual experience by demonstrating the relevant phenomena when asked to image certain things (e.g., to answer appropriately when asked at what distance an image of a horse would overflow her image). This would not be a case of the patient being disingenuous or being influenced by what the experimenters expected to find, which Farah et al. were at pains to deny, but of doing her best to carry out the task required of her, namely, to “imagine how it would look.”¹³

7.2. What would it mean if all the neuroscience claims turned out to be true?

Results such as those of Kosslyn et al. and Farah et al. have been widely interpreted as showing that retinotopic picture-like displays are generated on the surface of the visual cortex during imagery, and that it is *by means of this spatial display* that images are processed, patterns perceived from mental images and the results of mental imagery experiments produced. In other words, these results have been taken to support the view that mental images *are literally two-dimensional displays* projected onto primary visual cortex. I have already suggested some reasons why the neuroscience evidence does not warrant such a strong conclusion (and that a weaker “functional space” assumption is incapable of explaining the results of mental imagery studies without extrinsic constraints that are compatible with, and could be applied to, any form of representation). In addition, it should be remembered that standing against this interpretation of the neuroscience findings, is a large body of behavioral evidence that cannot be ignored. If we are to take seriously the conclusions suggested by the researchers who use neuroscience evidence to argue for the picture-theory (or the cortical-display theory), we need to understand the role that could be played by a literal picture being projected onto the visual cortex.

Here is a summary of some reasons to doubt the assumption that a picture is projected onto the visual cortex when we entertain mental images.

- (1) There is a great deal of evidence that the capacity for

visual imagery is independent of the capacity for visual perception. It is hard to reconcile the view that it is the topographical form of activation in early visual areas that is responsible for the visual image results discussed earlier, given the evidence for the dissociation between imagery and visual perception. For example, there is convergent evidence for the dissociation between the capacity for visual imagery and such visual deficits as cortical blindness (Chatterjee & Southwood 1995; Dalman et al. 1997; Goldenberg et al. 1995; Shuren et al. 1996), dyschromatopsia (Bartolomeo et al. 1997; De Vreese 1991; Howard et al. 1998), visual agnosia (Behrmann et al. 1992.; 1994; Jankowiak et al. 1992; Servos & Goodale 1995), and visual neglect (Beschin et al. 1997). The independence of imagery and vision is also supported by a wide range of both brain-damage and neuroimaging data, and the dissociation has been shown in both directions. The case for independence is made all the stronger by the evidence that blind people show virtually all the skills and psychophysical phenomena associated with mental imagery (including the reaction-time data discussed in sect. 3) (Barolo et al. 1990; Carpenter & Eisenberg 1978; Cornoldi et al. 1979; 1993; Craig 1973; Dauterman 1973; Dodds 1983; Easton & Bentzen 1987; Hampson & Duffy 1984; Hans 1974; Heller & Kennedy 1990; Johnson 1980; Jonides et al. 1975; Kerr 1983; Marmor & Zaback 1976; Zimler & Keenan 1983). Blind people may even report a comparable phenomenology concerning object shape as that of sighted people.

While there have been attempts to explain these dissociations by attributing some of the lack of overlap to an “image generation” phase that is presumably involved only in imagery (see the recent review in Behrmann 2000), this image-generation proposal does not account for much of the evidence for the independence of imagery and vision; in particular, it cannot explain how one can have spared imagery in the presence of such visual impairments as total cortical blindness. It has also been suggested that what characterizes patients who show a deficit on certain imagery-generation tasks (e.g., imagining the color of an object) is that they lack the relevant knowledge of the appearance of objects (Goldenberg 1992; Goldenberg & Artner 1991). Consequently, insofar as blind people know (in a factual way) what objects are like (including aspects of their “appearance” – such as their size, shape, and orientation), it is not surprising that they should exhibit some of the same psychophysical behaviors in relation to these properties.

Of course, it is also possible that many cortically blind people have deficits in only some parts of their visual system, and in particular, that they have damage to the more peripheral parts – those that are closer to sensory neurons. Thus, it might be that they still have the use of other parts of their visual system where the input is from a more central cortical locus. While this is certainly a possibility, it is not compatible with the view that in both vision and imagery, images are projected onto the primary visual cortex, since cortical blindness invariably involves damage to the visual cortex. A more plausible alternative may be that imagery studies do not tap a specifically visual capacity, but a more general spatial capacity.

There is good reason to believe that blind subjects have normal spatial abilities – and, indeed, blind children show the same spontaneous acquisition and use of a spatial vocabulary as do sighted children (Landau & Gleitman 1985).

Thus, mental imagery might involve a spatial mechanism rather than a visual one. This possibility, however, provides no comfort to the picture-theorists and it also leaves open the question of why these mechanisms connect with the motor system in a different manner than when they are visually stimulated (see sect. 5.4).

(2) The conclusion that many people have drawn from the neural imaging evidence cited earlier, as well as from the retinotopic nature of the areas that are activated, is that images are two-dimensional retinotopic displays. But that can't be literally the case. If mental images are depictive, they would have to be three-dimensional, inasmuch as the phenomenology is that of seeing a three-dimensional scene. Moreover, similar mental scanning results are obtained in depth as in 2D (Pinker 1980), and the phenomenon of "mental rotation" is indifferent as to whether rotation occurs in the plane of the display or in depth (Shepard & Metzler 1971). Neither can the retinotopic "display" in the visual cortex literally be three-dimensional. The spatial properties of the perceived world are not reflected in a volumetric topographical organization in the brain: as one penetrates deeper into the columnar structure of the cortical surface one does not find a representation of the third dimension of the scene. Furthermore, images represent other properties besides the spatial ones. For example, they represent color and luminance and motion. Are these also to be found displayed on the surface of the visual cortex? If not, how do we reconcile the apparently direct spatial mapping of 2D spatial properties with a completely different form of mapping for depth and for other contents of images of which we are equally vividly aware?

(3) The cortical display view of mental imagery assumes not only that mental images consist in the activation of a pattern that is the same as the pattern activated by the corresponding visual percept, but also that such a pattern mimics the retinotopic projection of a corresponding visual scene. In some versions, it even assumes that the pattern displayed in the cortex is a spatial extension of such a retinal projection, which incorporates a larger region of the scene than that covered by the retina, thereby explaining the subjective impression we have of a stable panoramic view of the world as we move our eyes around. Among the arguments put forward for the existence of an inner display is that it is needed to explain stability of the percept over eye movements and the invariance of recognition with translations of the retinal image (Kosslyn 1994, Ch. 4). It has also been suggested that the display is needed to account for the completion of the apparently detailed visual percept in the face of highly incomplete and partial sensory data (Kosslyn & Sussman 1995). The assumption behind these arguments is that incomplete sensory data are augmented in a visual display before being given over to the visual system responsible for recognition (and which presumably leads to our conscious awareness). While it may be that neural processes are responsible for certain cases of "filling in" phenomena, it is also clear that they do not do so by completing a partially filled display. (For a sophisticated discussion of the issue of filling-in, which makes it clear that cases of neural completion do not imply "analytical isomorphism," see Pessoa et al. 1998.)

Notwithstanding the fact that the early part of the visual cortex appears to be organized retinotopically, it is highly unlikely that this retinotopic organization serves to shield the inner eye from the incompleteness and instability of the

incoming sensory data, and thereby gives rise to such properties as the invariance of recognition over different retinal locations. There is every reason to believe that vision does not achieve stability and completeness, despite rapidly changing and highly partial information from the sensors, by accumulating the information in a spatially extended internal display. The fact that we sometimes feel we are examining an internal display in vision is simply a mistaken inference. Even if we had such a display we would not *see* it; we see the world and it is the world we see that appears to us in a certain way. The evidence clearly shows that the assumption that visual stability and saccadic integration is mediated by an inner-display is untenable (Blackmore et al. 1995; Irwin 1991; McConkie & Currie 1996; O'Regan 1992), since information from successive fixations cannot be superimposed in a central image as required by this view. Recent work on *change blindness* also shows that the visual system encodes surprisingly little information about a scene between fixations, unless attention has been drawn to it (Rensink 2000a; Rensink et al. 1997; Simons 1996), so there is no detailed pictorial display of any kind in vision, let alone a panoramic one.

(4) Although we can reach for imagined objects, there are significant differences between the way our motor system interacts with vision compared with the way it interacts with mental imagery (Goodale et al. 1994), as we saw in section 5.4. Such differences provide strong reasons to doubt that imagery provides input into the dorsal stream of the early vision system where the visuomotor control process begins, as it would if it were a retinotopic cortical projection.

(5) Accessing information from a mental image is very different from accessing information from a scene. To take just one simple example, we can move our gaze as well as make covert attention movements relatively freely about a scene, but not on a mental image. Try writing down a 3 x 3 matrix of random letters and read them in various orders. Now imagine the matrix and try doing the same with it. Or, for that matter, try spelling a familiar word backwards by imagining it written. Unlike the 2D matrix, some orders (e.g., the diagonal from the bottom left to the top right cell) are extremely difficult to scan on the image. If one scans one's image the way it is alleged one does in the mental scanning experiments, there is no reason why one should not be able to scan the matrix freely. Of course, one can always account for these phenomena by positing various properties specific to images generated by cognitive processes, as opposed to ones we retain from short-term visual memory. For example, one might assume that there is a limit on the number of elements that can be generated at one time, or one might assume that elements decay. But such assumptions are completely ad hoc. Visual information does not appear to fade as fast and in the same way from images held in short-term visual memory (Ishai & Sagi 1995), nor does it appear to fade in the case of images used to investigate mental scanning phenomena (which are much more complex, as shown in Fig. 1). Moreover, the hypothesized fading rates of different parts of an image have to be tuned post hoc to account for the fact that it is the *conceptual* as opposed to the *graphical* complexity of the image which determines how the image can be read and manipulated (i.e., to account for the fact that what one sees the image *as*, *how* one interprets it, rather than its geometry, is what determines its apparent fading). For example, it is the conceptual complexity of images that matters in determin-

ing the difficulty of an image superposition task (Palmer 1977), or in determining how quickly figures are “mentally rotated” (Pylyshyn 1979).

(6) The central role of conceptual, as opposed to graphical properties of an image, alluded to above, is an extremely general and important property of images. It relates to the question of how images are distorted or transformed over time, to how mental images can or cannot be (re)interpreted, and to how they can fail to be determinate in ways that no picture can fail to be determinate (Pylyshyn 1973; 1978; 1984). For example, no picture can fail to have a size or shape, or can fail to indicate which of two adjacent items is to the left and which to the right, or can fail to have exactly n objects (for some n), whereas mental images can be indeterminate in many ways. Not surprisingly, there are many ways of patching up a picture theory to accommodate such findings. For example, one can add assumptions about how images are tagged as having certain properties (perhaps including the property of not being based on real perception), and how they have to be incrementally refreshed from non-image information stored in memory, and so on, thus providing a way to bring in conceptual complexity and indeterminacy through the image generation function. With each of these accommodations, however, one gives the actual image less and less of an explanatory role until eventually one reaches the point where the display becomes a mere shadow of the mechanism that does its work elsewhere, as when the behavior of an animated computer display is determined by an extrinsic encoding of the principles that govern the animation, rather than by intrinsic properties of the display itself.

(7) The visual appearance of information projected onto a retinotopic display is very different from the appearance of information in a mental image. Images on the retina, and presumably on the retinotopically-mapped visual cortex, are subject to Emmert’s law: Retinotopic images superimposed onto a visual scene change their apparent size depending on the distance of the background against which they are viewed. Mental images imagined over a perceived scene do not change their size as the background recedes, providing strong evidence that they are not actually projected onto the retinotopic layers of the cortex.

(8) Images do not have the signature properties of early vision (such as the properties discussed in Hochberg 1968). If we create mental images from descriptions we do not find such phenomena as spontaneous interpretation of certain 2D shapes as representing 3D objects, spontaneous reversals of bistable figures, amodal completion or subjective contours (Slezak 1995), visual illusions, as well as the incremental construction of visual interpretations and reinterpretations over time, as different aspects are noticed. There is even evidence (discussed in sect. 6.4) that such early vision phenomena as motion aftereffects do not affect imagined motion the same way that they affect real perceived motion.

7.3. Is the “mind’s eye” just like a real eye?

Here is another way to think about the question of whether mental images could plausibly consist of patterns projected onto the cortex. Suppose it turns out that when we enter a mental image there is an actual copy of that very image (say, in the form of neural activity) on the surface of the primary visual cortex (or, for that matter, on the retina it-

self; the conceptual issue would be the same in either case). What would that tell us about the nature and role of mental images in cognition? We have known at least since Descartes that there is an image on our retina in visual perception, and perhaps there is also some transformed version of this image on our cortex, yet knowing this has not made us any wiser about how visual perception works. Indeed, ruminating on the existence of an image has only raised such problems as why we do not see the world as upside down, given that the image on the retina is upside down. The temptation to assume a literal picture observed through a “mind’s eye” may be very strong, but it leads us at every turn into blind alleys.

For example, some of the psychophysical evidence that is cited in support of a picture theory of mental imagery suggests a similarity between the mind’s eye and the real eye that is so remarkable that it ought to be an embarrassment to picture-theories. It not only suggests that the visual system is involved in imagery and that it examines a pictorial display, but it appears to attribute to the “mind’s eye” many of the properties of our own eyes. For example, it seems that the mind’s eye has a visual angle like that of a real eye (Kosslyn 1978), and that it has a field of resolution which is roughly the same as our eyes; it drops off with eccentricity according to the same function and inscribes a similar elliptical acuity profile as that of our eye (Finke & Kurtzman 1981a). It even appears that the “mind’s eye” exhibits the “oblique effect” in which the discriminability of closely-spaced horizontal and vertical lines is superior to that of oblique lines (Kosslyn et al. 1999b). Since, in the case of the eye, such properties arise from the structure of our retina and of its projection onto the visual cortex, it would appear to suggest that the mind’s eye is similarly constructed. Does the mind’s eye then have the same color profile as that of our eyes – and perhaps a blind spot as well? Does it exhibit after-images? And would you be surprised if experiments showed that it did?

Of course, the observed parallels could be just coincidence, or it could be that the distribution of neurons and connections in the visual cortex has come to reflect the type of information it receives from the eye. But it is also possible that such phenomena reflect what people have implicitly come to *know* about how things appear to them, a knowledge which the experiments invite them to use in simulating what would happen in a visual situation that parallels the imagined one. Such a possibility is made all the more plausible in view of the fact that the instructions in these imagery experiments explicitly ask observers to “imagine” a certain visual situation – that is, to imagine that they are in a certain visual circumstances and to imagine what it would look like to see things, say, in their peripheral vision. (I have often wondered whether people who wear thick-framed glasses would have a smaller field of vision in their mind’s eye.)

The picture that we are being presented, of a mind’s eye gazing upon a display projected onto the visual cortex, is one that should arouse our suspicion. It comes uncomfortably close to the idea that properties of the external world, as well as of the process of vision (including the resolution pattern of the retina and the necessity of moving one’s eyes around the display to foveate features of interest), are internalized in the imagery system. If such properties were built into the architecture, our imagery would not be as plastic and cognitively penetrable as it is. If the “mind’s eye”

really had to move around in its socket, we would not be able to jump from place to place in extracting information from our image the way we can. And if images really were pictures on the cortex, the necessity of a homunculus to interpret them would not have been discharged, notwithstanding claims that such a system had been implemented on a computer.

Computer implementation does not guarantee that what is said about the system, viewed as a model of the mind/brain, is true. Nor does it guarantee that the theory it implements is free of the assumption that there is an intelligent agent in one of the boxes. As Slezak (1995) has pointed out, labels on boxes in a computational model are not merely mnemonic; the choice of a label often constitutes a substantive claim that must be independently justified. Labeling a box as, say, “attention,” (as is done in the model described in Kosslyn et al. 1979) may well introduce a homunculus into the theory, despite the fact that the system is implemented as a running program which generates some of the correct predictions in a very limited domain. That’s because the label implies that the performance of the system will continue to mirror human performance in a much broader domain; it implies that the system can be scaled up in ways that are consistent with the assigned label.

7.4. What has recent neuroscience evidence done for the “imagery debate”?

Where, then, does the “imagery debate” stand at present? In the first place, although many investigators (including Kosslyn 1994, Ch. 1) write as though recent neuroscience evidence supersedes all previous behavioral evidence, nothing could be further from the truth. It was behavior (and phenomenological) considerations that raised the puzzle about mental imagery in the first place, and that suggested the picture theory. And it is a careful consideration of that evidence and its alternative interpretations that has cast doubt on the picture theory. Consequently, even if real colored stereo pictures were found on the visual cortex, the problems raised thus far in this article would remain, and would continue to stand as evidence that such cortical pictures were not serving the function attributed to them. For example, the fact that phenomena such as mental scanning are cognitively penetrable is strong evidence that whatever is displayed on the cortex is not what is responsible for the patterns of behavior observed in mental imagery studies. The mere fact that the data are biological does not give them a privileged status in deciding the truth of a psychological theory, especially one whose conceptual foundations are already shaky.

As I suggested near the beginning of this article, where the “imagery debate” stands today depends on what you think the debate was about. If it was supposed to be about whether reasoning by using mental imagery is *somehow* different from reasoning without it, nobody can doubt that. If it was about whether in some sense imagery involves the visual system, the answer there too must be affirmative, since imagery involves similar experiences to those produced by (and, as far as we know, *only* by) activity in some part of the visual system (though not in V1, according to Crick & Koch 1995). The real question is: *in what way* is the visual system involved and what does that tell us about the properties of mental imagery, and about how the mind generates and

uses images? It is much too early and much too simplistic to claim that the way the vision system is deployed in visual imagery is by allowing us to *look at* a reconstructed retinotopic input of the sort that comes from the eye (or at least at some locally-affine mapping of this input).

Is the debate, as Kosslyn (1994) claims, about whether images are depictive as opposed to descriptive? That all depends on what you mean by “depictive.” Is any representation of geometrical, spatial, metrical, or visual properties depictive? If that makes it depictive, then any description of how something looks is thereby depictive. Does being depictive require that the representation be organized spatially? As I suggested, that depends on what restrictions are placed on “being organized spatially”; most forms of representation, including symbol structures, use different spatial locations to distinguish among represented individuals. Does being depictive require that images “preserve metrical spatial information,” as has been claimed (Kosslyn et al. 1978)? Again, that depends on what it means to “preserve” metrical space. If it means that the image must *represent* metrical spatial information, then any form of representation will have to do that, to the extent that it can be shown that people do encode and recall such information. But any system of numerals, as well any analogue medium, can represent magnitudes in a useful way. If the claim that images preserve metrical spatial information means that imagery *uses spatial magnitudes to represent spatial magnitudes*, then this is a form of the literal picture theory, which I have argued is not supported by the evidence.

The neuroscience evidence we briefly looked at, while interesting in its own right, does not appear capable of resolving the issue about the nature of mental images, largely because the questions have not been formulated appropriately and the options are not well understood – with the single exception of the literal cortical display theory which turns out to be empirically inadequate in many different ways. One major problem with providing a satisfactory theory of mental imagery is that we are not only attempting to account for certain behavioral and neuroscience findings, but we are attempting to do so in a way that remains faithful to certain intuitions and subjective experiences. It is not obvious that all these constraints can be satisfied simultaneously. There is no a priori reason why an adequate theory of mental imagery will map onto conscious experience in any direct and satisfactory way. Indeed, if the experience in other sciences and in other parts of cognitive science is any indication, the eventual theory will not do justice to the content of our subjective experience and we will simply have to live with that fact, the way physics has had to live with the fact that the mystery of action-at-a-distance does not have a reductive explanation.

7.5. Is the “picture theorist” a straw man?

The typical response to arguments such as those raised in this section is that it takes the picture theory too literally and nobody really believes that there is an actual 2D display in the brain. For example, Denis and Kosslyn (1999) maintain that “No claim was made that visual images themselves have spatial extent, or that they occupy metrically defined portions of the brain,” and Kosslyn (1994, p. 329) admits that “images contain ‘previously digested’ information.” But if that is the case, how does one explain the increase in time to scan greater image distances or to report details in

smaller images? An explanation of these phenomena that appeals to a depictive representation *requires a literal sense of "spatial extent,"* otherwise the explanation does not distinguish the depictive story from what I have called the null hypothesis (see the discussion of the "functional space" alternative in sect. 5.2). If one denies the literal view of a cortical display, how does one interpret the claim that activation of topographically organized areas of the visual cortex during imagery establishes that images are "depictive"? If one were looking in the brain for evidence of a "functional space," what exactly would one look for? It is because picture theorists are searching for a literal 2D display that the research has focused on showing imagery-related activity in cortical Area 17.

The view that is favored by picture theorists is clearly illustrated by the importance that has been attached to the finding described in Tootell et al. (1982). In this study, macaques were trained to stare at the center of a pattern of flashing lights, while the monkeys were injected with radioactively tagged 2-deoxydextroglucose (2-DG), whose absorption is related to metabolic activity. Then the doomed animal was sacrificed and a record of 2-DG absorption in its cortex was developed. This record showed a retinotopic pattern in V1, which corresponded closely to the pattern of lights. In other words, it showed a *picture* in the visual cortex of the pattern that the monkey had received on its retina, written in the ink of metabolic activity. This led some people to conclude that we now know that a picture in the primary visual cortex appears during visual perception and is the basis for visual perception.¹⁴ Although no such maps have been found for imagery, there can be no doubt that this is what the picture-theorists believe *is* there and is responsible for both the imagery experience and the empirical findings reported when mental images are being used. People who have accepted this line of argument are well represented in the imagery debate: They are not "straw men"!

The problem is that while the literal picture-theory or cortical display theory is what provides the explanatory force and the intuitive appeal, it is always the picture *metaphor* that people retreat to in the face of the implausibility of the literal version of the picture theory. This is the strategy of claiming a decisive advantage for the depictive theory because it has the properties cited in the quotation in section 5.1 (e.g., it resembles what it represents), it is located in the topographically organized areas of visual cortex, it "preserves metrical information," and so on; then, in the face of its implausibility, systematically retreating away from the part of the claim that is doing the work – the literal spatial layout.

8. Conclusion: What is special about mental imagery?

The theme that has run through this essay is that we have thus far not been given adequate reasons to reject the null hypothesis and to accept that what goes on in mental imagery is in any way like examining a picture.¹⁵ Yet, the conclusion that reasoning using imagery is the same as reasoning that is not accompanied by the experience of "seeing in the mind's eye" is surely premature. It suggests that the fact that we have certain phenomenal experiences is irrelevant to understanding the nature of images or that the image ex-

perience is "epiphenomenal,"¹⁶ neither of which is warranted, even if at the present time we have no adequate understanding of what role the conscious experience of image content might play (see sect. 6.1). So what exactly are we entitled to conclude? What I have argued here is primarily that we are not entitled to take the tempting road of assuming that the content of our experience reflects in any direct way the nature of our cognitive information processing activity (in other words, one ought to guard scrupulously against what Pessoa et al. [1998] call the "analytical isomorphism" assumption). Thus, one ought to start off with strong skepticism about the idea that images are like pictures.

The search for a system of representation that retains some of the attractive features of pictures and yet can serve as the basis for reasoning has been the holy grail of many research programs, both in cognitive science and in artificial intelligence. The hope has been that one might develop a coherent proposal that captures some of what is special about reasoning with images, without succumbing to the Cartesian Theater trap. One way to approach this problem is to consider the most general constraints or boundary conditions that have to be met by a system of imagistic reasoning. Even if we do not have a detailed theory of the form of representation underlying imagistic thoughts, we do know some of the conditions it must meet. Starting by setting conditions on an adequate theory is not a new strategy. It was extremely successful in linguistics (Chomsky 1957a; 1957b) and is routine in theoretical physics ("thought experiments" can be viewed as explorations of what is entailed by such general conditions). Newell and Simon (1976) have also remarked on the importance of such orienting perspectives (which they called "laws of qualitative structure") for scientific progress.

As an example of such constraints, Fodor and I have argued (Fodor & Pylyshyn 1988) that in order to be adequate as vehicles of reasoning, such representations must meet the conditions of productivity, compositionality, and systematicity. Representations underlying imagery should meet additional conditions as well. Some years ago I proposed several possible constraints that are specific to mental images (Pylyshyn 1978). These have nothing to do with how images appear. Rather, they focus on the idea that mental images represent potentially visible token individuals or small sets of individuals. Because they represent individuals, they do not explicitly encode such set-properties as the cardinality of the set of individuals (they do not explicitly encode facts such as that there are eight boxes, or universally quantified propositions such as "all Xs are Y"). Because they represent individuals, they in effect assert the *presence* of some individuals or properties, and not their *absence* (e.g., they cannot represent a prepositional content such as "there is no X" or "it is not the case that P"). In addition, the *content* of images tends to involve visual rather than abstract properties.

The only theory I am aware of that shares some of the formal properties listed above, is a system of formal reasoning developed by Levesque and Brachman (1985). Levesque discovered a fundamental trade-off between the expressive power of a system of representation and the complexity of drawing inferences in that system. In Levesque (1986), he describes an expressively weaker form of logical representation (which he calls a "vivid representation") that allows inferences to be drawn essentially by

pattern matching. As in my earlier speculation about what is special about mental imagery, representations in this system do not permit the direct expression of negation (e.g., the only way that they can represent the proposition “there are no red squares” is by representing a scene that contains no red squares), or disjunction (e.g., they can only represent the proposition “the squares are either red or large” by allowing two possible representations, one with red squares and one with large squares), and they do not allow universal quantification (e.g., they can only represent the proposition “all squares are red” by explicitly representing each square, however many there are, and showing each as red). A vivid representation can only express the fact that there are five objects in a scene by representing each of the objects (of which there would be five in all). Levesque then proves some remarkable complexity properties for databases consisting of such vivid representations. Even though this work was not directly motivated by the imagery debate, it has the virtue of meeting the boundary conditions on an adequate system of representation; it does not postulate properties that we know are inadequate for the representation of knowledge, such as literally spatial displays.

ACKNOWLEDGMENTS

I wish to thank Jerry Fodor, Ned Block, and Peter Slezak for useful exchanges, and reviewers Michael McClosky, David Marks, Art Glenberg, and Mel Goodale, for providing helpful comments on an earlier draft. This work was supported by NIH Grant 1R01-MH60924.

NOTES

1. Some of these demonstrations can be viewed by downloading Quicktime® animations from: <http://www.cs.ubc.ca/~rensink/flicker/download/index.html>

2. When we first carried out these studies we were criticized (quite rightly, in my view) on the grounds that it was obvious that you did not have to scan your image if you did not want to, and if you did, you could do so according to whatever temporal pattern you chose. It still seems to me that the studies we carried out only demonstrate the obvious. That being the case, one might wonder what the great fuss was (and is) about over the scanning phenomenon (as well as the image size phenomenon described below); why dozens of studies have been done on it; and why it is interpreted as showing anything about the nature of mind as opposed to choices that subjects make.

3. People have suggested that one can accommodate this result by noting that the observed phenomenon depends on both the form of the image and on the particular processes that use it, so that the differences in the process can account for the different result obtained with different tasks (e.g., in this case attention might be moved by a “jump” operation). In that case the assumption that there is a depictive representation that “preserves metrical distance” does not play a role. The problem then becomes to specify the conditions under which the spatial character of the image does or does not play a role. A plausible answer is that scanning results are obtained when the subject thinks that imagining scanning a display is part of the task at hand. Of course, these are also the conditions under which the subject understands the task to be the simulation of visual scanning and, thus, recreating the time-distance scanning effect.

4. I don’t mean to pick on Stephen Kosslyn, who (along with Allan Paivio and Roger Shepard) has done a great deal to promote the scientific study of mental imagery. I focus on Kosslyn’s work here because he has provided what is currently the most highly developed and detailed theory of mental imagery, and has tried to be particularly explicit about his assumptions; and also because his work has been extremely influential in shaping psychologists’

views about the nature of mental imagery. In that respect, his views can be taken as the received view in much of the field.

5. If we look closely at what goes on in a computer implementation of a matrix, we see even more clearly that some of the space-like features are only in the mind of the user. For example, the matrix is said to contain a representation of empty space, in that the cells between features are actually represented explicitly. Whether registers are actually reserved in a computer for such empty cells is a matter of implementation and need not always be the case – indeed, it often is not the case in efficient implementations of sparse matrices. Moreover, since an empty place is just a variable with no value (or a default zero value), any form of representation can make the assumption that there are names for unfilled places. In fact, you don’t even have to assume that such place names exist prior to an inquiry being made about their contents; names can be (and are) created on the fly as needed (e.g., using LISP’s Gensym function). The same goes for the apparent pairs of numbers we think of as matrix coordinates; these are mapped onto individual names before being used to retrieve cell contents. The point is not just that the implementation betrays the assumption that such properties are inherent, it is also that how a matrix functions can be just as naturally viewed nonspatially since a matrix is not required by any computational constraints to have the properties assumed in a table display.

6. There is one other way of interpreting a functional space such as associated with a matrix. Rather than viewing it as a model of (real, physical) space, it might be thought of as a model of a (real, physical) analogue system that itself is an approximate analogue of space. In order for there to be an analogue model of space in the brain, however, there would have to be a system of brain properties that instantiate at least a local approximation of the Euclidean axioms. It would not do to just have an analogue representation of some metrical properties such as distance, – which itself is an eminently reasonable assumption, but we would need a whole system of such physically instantiated analogue properties in the brain. As in any analogue model, there would have to be a well-defined homomorphism from a set of spatial properties to a set of analogue properties: it would have to be possible to define predicates like between, adjacent, collinear, and so on, as well as the operation of moving-through these analogue dimensions at a specified speed. As far as I know, nobody has seriously developed a proposal for such an analogue representation (although the work of Nicod 1970 might be viewed as a step towards such a goal). But from the perspective of the present thesis this alternative suffers from the same deficiency that a literal spatial proposal does: It fails to account for the cognitive penetrability of the empirical phenomena that are cited in support of the picture theory of mental imagery.

7. There are almost always some visual elements, such as “textons” (Julesz 1981), that can serve as indexed objects. It is dubious whether places unoccupied by any visible feature can be indexed (which is why vision in a featureless environment is so unstable; Avant 1965). The one possibility, suggested by the work of Taylor (1961), is that locations clearly definable in terms of nearby visible objects can be indexed. Taylor showed that observers can encode locations more accurately when they are easily specified in relation to some visible anchors (such as being the “midpoint” between two visible objects).

8. Brandt and Stark (1997) reported that the sequence of eye movements observed when inspecting a mental image inscribe similar scan paths to those observed when inspecting an actual display. But since the experiment was not carried out in total darkness, the eye movements could have been made to the faint visible cues, rather than to image features (see sect. 5.3). In addition, since moving one’s eye does not result in viewing different parts of an inner picture, the contents of the display could not provide the feedback required to control a sequence of eye movement beyond the initial ballistic saccade, so that in any case the sequence of eye movements during imagery is likely due to a different mechanism than the one that controls the sequence of eye movements in vision.

9. Finke also found a significant effect of “vividness” of imagery, as determined from the Vividness of Visual Imagery Questionnaire (VVIQ) (Marks 1973), with the adaptation effect being much higher for vivid imagers. It is not clear how to interpret such a result, however, given that subjects who were high in vividness also had significantly higher adaptation scores in the control condition where there was no feedback (visual or imaginal) about errors of movement. These findings (along with the connection between vividness and hypnotic suggestibility reported by Crawford 1996; Glisky et al. 1995; Kunzendorf et al. 1996) increase the likelihood that experimental demand effects may be involved in the performance patterns of high-scoring subjects.

10. The conditions under which one gets more or less adaptation are discussed in Howard (1982). The most important requirement is that the discordant information be salient for the subject and that it be interpreted as a discordance between two measures of the position of the same limb. Thus, anything that focuses more attention on the discordance and produces greater conviction that something is awry helps strengthen the adaptation effect. Thus it is not surprising that merely telling subjects where their hand is does not produce the same degree of adaptation as asking them to pretend that it actually is at a particular location, which is what imagery instructions do.

11. Many of these studies have serious methodological problems, which we will not discuss here in detail. For example, a number of investigators have raised questions about many of these illusions (Predebon & Wenderoth 1985; Reisberg & Morris 1985) where the likelihood of experimenter demand is high. The usual precautions against experimenter influence on this highly subjective measure were not taken (e.g., the experiments were not done using a double-blind procedure). The most remarkable of the illusions, the orientation-contingent color aftereffect, known as the McCollough effect, is perhaps less likely to lead to an experimenter-demand effect since not many people know of the phenomenon. Yet, Finke and Schmidt (1977) reported that this effect is obtained when part of the input (a grid of lines) is merely imagined over the top of a visible colored background. But the Finke finding has been subject to a variety of interpretations as well as to criticisms on methodological grounds (Broerse & Crassini 1981; 1984; Harris 1982; Kunen & May 1980; 1981; Zhou & May 1993) so will not be reviewed here. Finke himself (Finke 1989) appears to accept that the mechanism for the effect may be that of classical conditioning rather than a specifically visual mechanism.

12. In a recent paper, Kosslyn et al. (1999a) also claimed that if area 17 is temporarily impaired using repetitive transcranial magnetic stimulation (rTMS), performance on an imagery task is adversely affected (relative to the condition when they do not receive rTMS), suggesting that the activation of area 17 may be not only co-relational, but may also play a causal role. However, this result must be treated as highly provisional since the nature and scope of the disruption produced by the rTMS is not well established and the study in question lacks the appropriate controls for this critical question; in particular, there is no control condition measuring the decrement in performance for comparable tasks that do not involve imagery.

13. A possible control for this explanation would be to study patients whose loss of peripheral vision and delay in testing followed roughly the same pattern as Farah’s patient but in which the damage was purely retinal. The expectation is that under the same instructional conditions such patients would also exhibit tunnel imagery, even though there was presumably no relevant cortical damage involved. Another control would be to test the patient immediately after surgery, before she learned how the visual world looked to her in her post-surgical condition.

14. That a topographic display is involved in vision is hardly surprising, since we know that vision begins with retinal images. But before either retinal or cortical patterns become available to cognition as percepts, they have to be interpreted; this is what vision is for, it is not for turning one retinotopic pattern into another. The original motivation for hypothesizing a visual image was to ac-

count for the completeness and spatially extended nature of visual perception despite the incompleteness of retinal information, the stability of the perceptual world despite eye movements, and the robustness of recognition despite differences in the size and location of objects on the retina (these are among the reasons Kosslyn [1994] gives for needing a display that also serves as the screen for mental images). It is now pretty clear that there are no visual images serving these purposes in vision (see, e.g., O’Regan & Noë 2001).

15. This paper has been concerned primarily with the picture-theory account of imagery because it appears to be the overwhelmingly dominant one. I have not attempted to review other options, such as those proposed by Barsalou (1999) and Thomas (1999). In any case, these other approaches are not claimed to provide an explanation of the many experimental findings sketched in this paper.

16. The claim that images are epiphenomenal or have no causal role rests on an ambiguity about what one means by an image. As Block has correctly pointed out (Block 1981), the appeal to epiphenomenalism is either just another way of stating the disagreement about the nature of a mental image, or else it simply confuses the functional or theoretical construct “mental image” with the experience of having a mental image.

Open Peer Commentary

Commentary submitted by the qualified professional readership of this journal will be considered for publication in a later issue as Continuing Commentary on this article. Integrative overviews and syntheses are especially encouraged.

Depicting second-order isomorphism and “depictive” representations

Hedy Amiri and Chad J. Marsolek

Department of Psychology, University of Minnesota, Minneapolis, MN 55455.
amir0010@tc.umn.edu chad.j.marsolek-1@umn.edu
<http://levels.psych.umn.edu>

Abstract: According to Pylyshyn, depictive representations can be explanatory only if a certain kind of first-order isomorphism exists between the mental representations and real-world displays. What about a system with second-order isomorphism (similarities between different mental representations corresponding with similarities between different real-world displays)? Such a system may help to address whether “depictive” representations contribute to the visual nature of imagery.

Pylyshyn argues that depictive representations can be explanatory in a theory of imagery only if the physical-spatial extent in the underlying representational medium is used to represent information about spatial extent in the real-world display that is being imagined. We suggest that a different perspective involving second-order isomorphism deserves comment within this part of the imagery debate; a less restrictive notion of veridical representation may help to clarify fundamental issues.

Shepard and Chipman (1970) used the term first-order isomorphism to describe the situation in which a similarity relation exists between an internal representation and the individual real-world object being represented. From this perspective, Pylyshyn’s claim essentially is that depictive representations can be explanatory only if there is a certain kind of first-order isomorphism between the mental representations and the real-world displays; a first-order isomorphism must exist between the physical-spatial

extent in the representational medium and the spatial information in the displays being imagined. An alternative to first-order isomorphism, according to Shepard and Chipman, is second-order isomorphism, in which a similarity relation exists between the similarities among internal representations and the corresponding similarities among multiple real-world objects being represented. Indeed, Shepard and Chipman argued that it is second-order isomorphism that should be sought in theories of imagery. A system with second-order isomorphism may provide an understanding of visual imagery in which depictive representations of the first order need not be posited; however, important and inherent aspects of the system may be that “depictive” representations of the second order are posited and the visual nature of imagery phenomena may be attributable at least in part to the relevant cognitive architecture.

To be concrete about such a system, the retina and perhaps early visual areas presumably represent two-dimensional perceptual views of shapes in a manner that is at least close to first-order isomorphic. However, the ultimate representation underlying shape recognition resides in high-level visual areas (i.e., inferior-temporal cortex for nonhuman primates; occipital-temporal areas in humans) of the sort that are activated very frequently in neuroimaging studies of visual mental imagery (see sect. 7.1). Presumably, multiple transformations from the initial retinotopic representation to the high-level representation take place, but the focus here is on how to understand such high-level representations. Edelman (1998) put forth a mathematical formulation of the kinds of transformations that may be performed on visual inputs, which allow a mapping from distinct points in a high-dimensional input space to distinct points in a (lower-dimensional) representation space. The mapping is a composite of four functions, involving the object's geometrical coordinates, the object's viewing conditions, a convolution of the image through a number of filters and application of a nonlinear function, and a dimensionality reduction. Most important for present concerns, the net result of this mapping is that two points that were near each other before the mapping will be near each other after the mapping. In this way, what is represented is the similarity between shapes, not the geometry of the shapes themselves. In other words, the representations are second-order isomorphic, and the veridicality is with respect to a correspondence between original (represented) and internal (representing) similarities among shapes. Thus, by Pylyshyn's relatively strong definition, the representations are not depictive. But, is there a way in which such representations, if used in visual imagery, nonetheless reflect an important aspect of the cognitive architecture?

Within such a system, mental imagery tasks involving visual shapes could be explained by systematic movements through locations in the space of second-order isomorphic representations. Smooth, continuous movements through such a space could be used to accomplish mental rotation and perhaps mental scanning. Given that nearby locations in the representation space correspond to nearby locations in what would be perceived in the original shape space, the assumption is that, to get from one location to a distant location in representation space, the intermediate locations must be worked through not unlike the way an attractor neural network passes through a series of similar activation states when settling into an attractor. Of course, this could account for chronometric results in rotation studies. It is also possible that such movements through the representation space may be accompanied by sequential top-down activations of early visual areas due to the reciprocal excitatory connectivity of low- and high-level visual areas (Felleman & Van Essen 1991). But, by this account, we should note that such activation of early visual areas may not be sufficient or even necessary for visual imagery (instead, it could reflect involuntary/automatic top-down activation that is used normally in perceptual recognition but is not doing the work of underlying imagery effects).

If imagery is accomplished in such a manner, an inherent and important aspect of imagery may be that it reflects what Pylyshyn

labels the “cognitive architecture.” The cognitive architecture cannot be “directly altered by changes in knowledge, goals, utilities, or any other representations (e.g., fears, hopes, fantasies, etc.)” (sect. 3, para. 2). If mental rotation occurs as it does due to systematic movement through similar locations in a space representing visual similarities, then it occurs as it does in part due to well established mappings between views of real-world shapes and the higher-level mental representations, where a distinct point in the former maps to a distinct point in the latter. It is the form of this representation that would dictate the “visual” nature of the phenomenology and the observable visual imagery effects. This would be a stable aspect of the architecture; changing it would entail impairing visual recognition abilities and changing the cognitive architecture. There is also a sense in which the high-level representations would not be depictive but would have properties that make them “pseudo-depictive.” They would not use spatial extent across cortex to represent information about spatial extent in the image. However, they may be “pseudo-depictive” in that they would use an isomorphic representation system (albeit second-order) to represent similarities between original images that do use spatial extent to represent spatial information. The possibility of such second-order “depictive” representations may be important for addressing whether “depictive” representations contribute to the visual nature of imagery.

ACKNOWLEDGMENT

This work was supported by NIH grant MH60442.

Visual imagery is not always like visual perception

Martha E. Arterberry,^a Catherine Craver-Lemley,^b and Adam Reeves^c

^aDepartment of Psychology, Gettysburg College, Gettysburg, PA 17325;

^bDepartment of Psychology, Elizabethtown College, Elizabethtown, PA 17022; ^cDepartment of Psychology, Northeastern University, Boston, MA 02115 USA. arterber@gettysburg.edu lemleyce@acad.etown.edu reeves@neu.edu

Abstract: The “Perky effect” is the interference of visual imagery with vision. Studies of this effect show that visual imagery has more than symbolic properties, but these properties differ both spatially (including “pictorially”) and temporally from those of vision. We therefore reject both the literal picture-in-the-head view and the entirely symbolic view.

Pylyshyn repeatedly draws parallels between imagery and visual perception to support his argument that imagery is not pictorial. He suggests that processes that use the same symbolic vocabulary can explain the similarities between imagery and perception. This very well may be the case when one considers visual imagery and visual perception in high-level tasks. We offer one more piece of evidence for this view. In Craver-Lemley et al. (1999), participants were asked to imagine a geometrical figure, such as a blue open triangle, and then they were presented with a display containing two or more geometrical figures (Fig. 1A). Participants were asked to report on the features of one of the physically presented objects, following the methodology of Treisman and Schmidt (1982). Craver-Lemley et al. found that features of the imagined figure were mistakenly conjoined with features of the physically presented figures, and the effect was not due to participants hearing some features spoken aloud when they were given instructions regarding what to imagine (e.g., “a blue, open triangle”). These findings suggest that imagery influences perception at the level of visual processing at which features are combined, and they are consistent with Pylyshyn's contention that the same processes may underlie the representation of real and imagined objects. Moreover, we can speculate with him that these processes are symbolic, in the sense that subjects may encode a list of features (red,

square, filled, etc.) and that this list is tagged in some way to indicate their collocation on the same object. If imagined and real features are encoded, stored, or represented in the same way, then at recall they may be confused, resulting in illusory conjunctions between percepts and images.

The argument of a common symbolic vocabulary is less convincing when one considers other tasks, particularly lower-level tasks such as acuity or target detection. There are several findings with visual imagery that do not mirror those with physically presented stimuli. For example, in a signal-detection paradigm based on Segal and Fusella (1970), participants presented with an acuity target (vertical line offsets, Fig. 1B) in a region where they imagined four vertical lines showed strong interference (a loss of 0.8 d' units). Similar interference was found when imagining four horizontal lines (Craver-Lemley & Reeves 1987, Fig. 1C). However, with physically presented lines, only the vertically oriented lines interfered with acuity. Again, Craver-Lemley and Reeves found that imagined lines interfered with the target at spatial extents in which physical lines have no effect, and that imagined lines still interfered for up to five seconds after the subjects stopped imagining them, unlike real lines which only interfere (as masks) for one or two tenths of a second after presentation. A lack of correspondence between physical and imagined stimuli was discovered under conditions of induced-depth as well. Imagining four vertical lines or a solid bar in front of a line target in an induced depth display (Fig. 1D) interfered with acuity but imagining the four lines or a solid bar behind the target location did not (Fig. 1E, Craver-Lemley et al. 1997). In contrast, physically presented bars interfered with acuity regardless of whether they were located in front of or behind the target. And finally, imagining a solid bar interfered with target detection (an asterisk) only when the target overlapped the image location (Craver-Lemley & Arterberry 2001). In this case, the target and the image had no features in common except spatial location.

We explained many of these results by postulating that the visual system suppresses competing (local) visual input from the visual field in order to facilitate entertaining a visual (mental) image. But why might imagery and perception compete? We do not

know, but an interesting suggestion is from Sartre (1948), who, having elegantly dismantled various picture-in-the-head views, concluded, "The image and the perception, far from being two elementary psychological factors of similar quality which simply enter into different combinations, represent the two main irreducible attitudes of consciousness. It follows that they exclude each other" (p. 153). Whether or not this is so, the properties of the interference effect have fairly clear implications for the spatial nature of visual imagery. We note here that Segal and Fusella (1970) also demonstrated that complex auditory images (e.g., of bells) interfered with auditory detection (of tones), but no one has followed this up with simpler images. Our notion that perceptual systems suppress inputs in order to make room for images (or Sartre's notion of exclusion) implies that auditory imagery is similar to visual in this respect. Thus, images of pure notes should interfere chiefly with neighboring frequencies, should do so over a broader spectrum than real-tone maskers, and perhaps should interfere for a longer period of time.

Thus, we have several examples where we do not find complete concordance of effects with real and imagined stimuli; a visual image does not always mimic the effects of physical stimuli. The difference in interference effects between imagined and real stimuli described in the above examples cannot be accounted for by attentional factors, as shown by dual-task attentional manipulations (see Craver-Lemley & Reeves 1992). Nor can they be explained easily by tacit knowledge. Many of the interference effects are contrary to expectations based on experiences with real stimuli (e.g., as mentioned, real horizontal lines do not interfere with Vernier acuity; if subjects knew this to be the case, then imagined horizontal lines should not interfere with Vernier acuity either, but they do). Finally, we are sure that expectations are irrelevant. Craver-Lemley and Reeves (1992) told half the participants in one experiment that imagined vertical lines would facilitate performance, and half that the image would impair it. The participants all believed the cover story and all showed *interference* despite their different expectations.

Pylyshyn states, "It may be that visual percepts and visual images interact because both consist of symbolic representations that use some of the same proprietary spatial or modality-specific vocabulary" (sect. 6.2). We accept that symbol interference may happen at a higher level of representation responsible for image-stimulus illusory conjunctions, but we think the application of "symbolic representations" to spatial and temporal contiguity in the interference effect is stretching matters. Surely the spatial properties of interference point to a pictorial component of visual mental imagery.




Figure 1 (Arterberry et al.). A. Stimulus used to study illusory conjunctions between physical and imagined geometric figures. The hashed lines represent the item the participant imagined. B. Vernier acuity target and a four-vertical-line image. C. Vernier acuity target and a four-horizontal-line image. In D and E the bar image was positioned either in front of (D) or behind (E) the acuity target in an induced-depth display.

Can we change our vantage point to explore imaginal neglect?

Paolo Bartolomeo^a and Sylvie Chokron^b
^aINSERM Unit 324, Centre Paul Broca, F-75014 Paris, France; ^bLPE-CNRS UMR 5105, UPMF, BP 47 38000, Grenoble, France. paolo@broca.inserm.fr
chokron@upmf-grenoble.fr
<http://www.upmf-grenoble.fr/upmf/RECHERCHE/lpe/index.html>

Abstract: Right brain-damaged patients with unilateral neglect, who ignore left-sided visual events, may also omit left-sided details when describing known places from memory. Modulating the orienting of visual attention may ameliorate imaginal neglect. A first step toward explaining these phenomena might be to postulate that space-related imagery is a cognitive activity involving attentional and intentional aspects.

Patients with lesions in the posterior part of the right hemisphere may ignore events on their left side, a condition known as unilateral neglect. Neglect patients are often completely unaware of their disorder (they are said to be "anosognosic"), and extremely unwilling to acknowledge it. A large amount of neuropsychologi-

cal evidence (reviewed in Bartolomeo & Chokron 2002) suggests that left-sided stimuli fail to exert their normal attraction on neglect patients' attention. Thus, a basic mechanism of left neglect could be a deficit of exogenous, or stimulus-related, orienting of attention toward left-sided targets. In partial disagreement with this interpretation, it has been shown that neglect can occur not only in vision, but also in the absence of any physical object in the patient's visual field. For example, when asked to imagine and describe from memory familiar surroundings from a determined vantage point, neglect patients can omit left-sided details, only to later describe these same details when invited to assume the opposite point of view (Bisiach et al. 1981; Bisiach & Luzzatti 1978). In these studies, imaginal neglect co-occurred with visual neglect. This association has often been interpreted as supporting pictorial models of visual mental imagery (Bisiach & Berti 1990; Kosslyn 1994). Neglect patients would avoid mentioning left-sided imagined details because they would lack the left half of a (spatially organized) mental representation (Bisiach & Luzzatti 1978). It would indeed be difficult to contend that neglect patients have a (however tacit) knowledge of their visual exploratory bias, and would consequently reproduce in imaginal tasks a neglect behavior of which they are, as a rule, completely unaware (Bisiach & Berti 1990). It is, of course, also hard to see how a propositional code compatible with Pylyshyn's "null hypothesis" could have such spatial or directional properties to account for imaginal neglect.

On the other hand, the accumulation of neuropsychological evidence of multiple dissociations between imagery and perceptual abilities in brain-damaged patients (recently reviewed in Bartolomeo 2002), has proved devastating for models of mental imagery based on a functional and anatomical equivalence between these abilities, like Kosslyn's pictorial model. Some of these dissociations are not only functional, but seem to have also an anatomical basis. While occipital damage can determine perceptual deficits, it seems neither necessary, nor sufficient to produce imagery deficits. On the other hand, rather extensive damage of the left temporal lobe seems necessary in order to produce visual imagery deficits for object shape or color (Bartolomeo 2002), as well as for orthographic material (Bartolomeo et al. 2002). Although dissociations have been described between visual and imaginal neglect (see Bartolomeo & Chokron 2001 for a recent review), no such anatomical segregation apparently emerged. Apart from occasional case descriptions of imaginal neglect after right frontal (Guariglia et al. 1993) or thalamic damage (Ortigue et al. 2001), most cases of imaginal neglect result from lesions in the right temporal-parietal cortex, which is the same anatomical correlate of visual neglect (Vallar 1993).

To explore the relationships between visual and imaginal neglect, we assessed them in 30 right- and 30 left-brain-damaged patients, and found imaginal neglect only in right-brain-damaged patients (Bartolomeo et al. 1994). Imaginal neglect always co-occurred with visual neglect,¹ and scores measuring the lateral bias in the two types of tasks positively correlated, thus suggesting that the two disorders share some common underlying mechanism. Additional evidence confirming a relationship between visual and imaginal neglect comes from the outcome of maneuvers known to modulate visual neglect. When a patient had his eyes and head physically turned toward the left side, his descriptions from memory included more left-sided details (Meador et al. 1987). Similar results were obtained by irrigating patient's left ear with cold water (Rode & Perenin 1994), a vestibular stimulation likely to induce a leftward orienting of attention (Gainotti 1993). Imaginal neglect was also reduced by introducing a short adaptation period to a prismatic rightward shift of the visual field to the right (Rode et al. 2001), another maneuver known to ameliorate visual neglect (Rossetti et al. 1998). Thus, sensory-motor procedures can influence imaginal neglect.² It has been proposed that at least some of these procedures act by facilitating leftward orienting of attention (Chokron & Bartolomeo 1999; Gainotti 1993).

If so, one could surmise that neglect patients' visual attention can be laterally biased during place description, thus producing

signs of imaginal neglect. In section 5.4 of his target article, Pylyshyn suggests that visuo-motor effects on imagery might depend on orienting one's gaze or attention on *real*, as opposed to imagined, locations. This interesting possibility, which would be coherent with what we know about the neglect patients' tendency to be attracted by right, non-neglected, visual targets (Gainotti et al. 1991), could perhaps help explain imaginal neglect. During place description, patients' attention could be attracted by right-sided visual details, and this could in some way influence their performance in imaginal tasks. However, this account does not hold, at least for the studies of the Lyon group, in which patients kept their eyes closed during the imaginal tasks (Gilles Rode, personal communication). If there is an asymmetry of attentional shifts in imaginal neglect, then, it would be rather akin to analogous biases that neglect patients show in situations where no external stimulus is present, as, for example, in the disappearance of leftward REMs during sleep (Doricchi et al. 1993). An implication of this possibility, and one which is relevant to the "imagery debate," is that orienting of attention can influence space-related imagery. Although visual images are certainly not "seen" by the visual system, the phenomenon of imaginal neglect is consistent with the possibility that visual imagery involves some of the attentional-exploratory mechanisms that are employed in visual behavior (Thomas 1999). According to a recent proposal (O'Regan & Noë 2001), these motor processes are actually responsible for the "visual" character of visual experience. The "perceptual" aspects of visual mental images might thus result not from the construction of putative "quasi-perceptual" representations, but from the engagement of attentional and intentional aspects of perception in imaginal activity.

ACKNOWLEDGMENT

We thank Nigel Thomas for very helpful discussion on visual mental imagery.

NOTES

1. In fact, about two thirds of left neglect patients showed definite signs of neglect only in visual tasks, and not in imaginal tasks, probably because right-sided visual details exerted a powerful attraction on patients' attention (see Gainotti et al. 1991). However, when imaginal neglect was present, it was always associated with visual neglect.

2. Conversely, a purely imaginal training can ameliorate visual neglect (Smania et al. 1997).

Spatial models of imagery for remembered scenes are more likely to advance (neuro)science than symbolic ones

Neil Burgess

Institute of Cognitive Neuroscience and Department of Anatomy, University College London, London, WC1E 6BT, United Kingdom.

n.burgess@ucl.ac.uk <http://www.icn.ucl.ac.uk/members/Burge12/>

Abstract: Hemispatial neglect in imagery implies a spatially organised representation. Reaction times in memory for arrays of locations from shifted viewpoints indicate processes analogous to actual bodily movement through space. Behavioral data indicate a privileged role for this process in memory. A proposed spatial mechanism makes contact with direct recordings of the representations of location and orientation in the mammalian brain.

Pylyshyn's target article omits some of the evidence for the spatial organisation of visual imagery to be found in studies of memory for spatial scenes or arrays of objects. While not conclusive, this evidence may be instructive in escaping some of the logical caveats raised by Pylyshyn, and extending the discussion of the functional space in which retrieval products from memory are processed. Although other caveats will be found regarding these data, interpreting them in terms of their mapping onto space and our phys-

ical movements within it will take us closer to understanding the relevant neural mechanisms. Thus, since science advances by a process in which one flawed but partially explanatory theory replaces another flawed but slightly-less explanatory theory, the spatial interpretation appeals to me as a neuroscientist. The evidence discussed here concerns (1) the spatial organisation of hemi-spatial neglect in imagery; (2) reaction time and performance data in memory for spatial locations; (3) the neuronal mechanisms suggested by single unit recordings in animals.

In patients with hemi-spatial neglect, damage to the internal image or to the means of accessing it occurs preferentially to the side contra-lateral to the lesion. How could this be unless the internal image itself were spatially organised? Pylshyn (sect. 7.1) discusses Farah et al.'s (1992) patient who shows tunnel vision and also similar tunnel imagery. He argues that this patient has simply learned to simulate her impaired visual perception in imagery, that is, that this may not relate to the "cognitive architecture" of imagery. Can this objection be applied to hemi-spatial neglect in imagery? The majority of patients showing hemispatial neglect in imagery also show a similar perceptual neglect (Bisiach et al. 1979; 1981), indicating significant overlap between the architecture of the two systems. However, the caveat that imagery might imitate perception is ruled out by the (albeit much rarer) case of patients showing relatively pure imaginal neglect (Beschlin et al. 1997; Guariglia et al. 1993), and even imaginal neglect on one side and perceptual neglect on the other (Beschlin et al. 2000).

The second piece of evidence concerns memory for the locations of objects in an array following a change in viewpoint. In these experiments, reaction times show a linear dependence on the size of the change in the subject's location or orientation between presentation and retrieval (Diwadkar & McNamara 1997). Related imagery experiments require the subject, previously shown an array of locations, to point in the direction a location would have following a (imagined) rotation or translation of the subject. These experiments show a similar dependence of reaction time on the size of the rotation or translation between the subject's current position and the position from which they should imagine pointing (Easton & Sholl 1995). These tasks probably differ from those involving single objects (e.g., Shepherd & Metzler 1971) in being solved by imagined movement of viewpoint as opposed to an equivalent imagined movement of the array (for which RTs and performance are worse). Only when a single object need be considered can imagined rotation of the array produce performance approaching that for imagined movement of viewpoint (Wraga et al. 2000).

The same caveats apply to the interpretation of viewpoint manipulation data that Pylshyn raises against mental rotation of single objects. However, in this case, there is independent evidence that our "cognitive architecture" is specifically adapted to accommodate the effects of physical movement through the environment compared to an equivalent movement of the array (Simons & Wang 1998; Wang & Simons 1999). In these experiments, subjects' recognition memory for an array of objects on a circular table top is better after the subject had moved around the table to a new viewpoint than after an equivalent rotation of the table top. Since this effect is also observed in the dark (using phosphorescent objects) and in purely visual virtual reality (Christou & Bulthoff 1999), the facilitation appears to apply to any processes corresponding to movement of viewpoint within the subject's mental model of the world.

What are the neural bases of these processes? A patient with focal damage to both hippocampi is specifically impaired at shifted view recognition of two or more object-locations compared to fixed-view recognition, or shifted view recognition of a single object-location (King et al. 2002). The neural bases of self-location and orientation have been well examined in the rat. "Place cells" in the hippocampus encode the animal's current location in the environment (O'Keefe 1976; Wilson & McNaughton 1993) while "head-direction cells" nearby in the presubiculum (also mammillary bodies and anterior thalamus) encode its current orientation

(Taube et al. 1990). Additionally, cells in the connected area 7a of monkey parietal cortex represent stimulus locations in frames of reference relative to eye, head, and trunk, and allow translation between these frames and the environmental frame (Andersen et al. 1985; Pouget & Sejnowski 1999; Snyder et al. 1998). Viewpoint-dependent retrieval of remembered places can be modelled as an interaction between the parietal, place and head-direction systems, possibly accounting for their involvement in episodic memory (Burgess et al. 2001). Interestingly, current models of the place and head-direction systems see each representation as a "continuous attractor" (Zhang 1996), in which the represented location or direction can shift under internal dynamics, but at a fixed speed (determined by the effective asymmetry of the connections between cells). This mechanism, applied to the viewpoint-dependent retrieval model, could explain the reaction time data, providing an explanatory model linking cells to spatial memory and imagery. Symbolic accounts seem less well formed to address these types of data.

Pictures, propositions, and primitives in the head

Anjan Chatterjee

Department of Neurology and Center for Cognitive Neuroscience, University of Pennsylvania, Philadelphia, PA 19104. anjan@mail.med.upenn.edu
<http://ccn.upenn.edu/people/anjan.html>

Abstract: Data from neuropsychology do not support the idea that the primary visual cortex necessarily displays internal visual images. However, the choice of formats used in human cognition is not restricted to depictive or descriptive representations. Nested between pictures and propositions, primitive spatial schemas with simple analog features extracted from pictorial scenes may play a subtle but wide role in cognition.

The hypothesis that the primary visual cortex serves as a plasma screen for our subjective experience of visual images falters when faced with neuropsychological evidence. We reported that two of three patients with cortical blindness imaged very well (Chatterjee & Southwood 1995). The third had a lesion involving the left temporal-occipital junction, an area implicated in the generation of such images. Patients with cortical blindness and relatively spared imagery are not exceedingly rare (Goldenberg et al. 1995). Butter and colleagues (Butter et al. 1997) contend that such patients do not have complete damage to primary visual cortex implying that islands of preserved tissue can display internally generated images. This position is peculiar. The idea that preserved islands of primary visual cortex support visual imagery but not vision undermines the original point of a close functional homology between imagery and perception in early visual cortex.

To oppose the notion of pictures displayed in primary visual cortex, however, sidesteps the deeper issue of whether human cognition involves more than one representational format. If one accepts that much of perception has an analog organization and much of language does not, and further, that one uses language to communicate information gleaned from perception, then how does one get from perception to language? One possibility is that primitive spatial schemas lie between perceptual and linguistic representations and play a role in human cognition (Chatterjee 2001). This proposal is sketchy at best since it is informed by empirical observations that are at present limited. However, see Talmy (2000) for a related and developed discussion of schemas born of a separate theoretical tradition.

The general idea is that schemas retain some analog features of perception and incorporate discreteness found in symbol systems. From pictures, simple geometric features such as points, planes, and vectors are distilled, while the details and sensorial richness of perception are discarded. The schemas are discrete in that their features are distinct, generalizable, and easily categorized. As an

example of such a schema, events with participants involved in actions are conceived with agents on the left of the recipient of actions, and with the actions moving from left to right. If this schema underlies event concepts, then we should see its traces when subjects engage these representations.

The following observations support this hypothesis.

1. When normal subjects make semantic associations of pictures of actions as opposed to objects, elementary visual motion areas MT/MST are activated preferentially, even though the pictures are not moving (Kable et al. 2002). Although these data do not establish that action representations are schematic, they demonstrate that the neural mediation of perception and cognition overlap.

2. Some aphasic individuals have a grammatical disturbance called a thematic role assignment deficit (Chatterjee & Maher 2000). They are unable to determine who is doing what to whom when matching sentences to pictures. We described such a patient who could not map thematic roles from sentences to pictures. However, rather than respond randomly, he was accurate with simple active sentences only on pictures with the agent on the left (Chatterjee et al. 1995a; Maher et al. 1995). We proposed that his language deficit uncovered a primitive representation of such events.

3. Normal subjects demonstrate spatial biases when conceiving action events. When asked to draw events with two participants described by sentences, right-handed subjects draw the agent on the left and the recipient of the action on the right. When asked to draw either the agent or the recipient alone, they draw agents to the left of where they draw recipients (Chatterjee et al. 1995b). They also draw action trajectories, such as the path described by the phrase “dog running” from left to right. In sentence-picture matching experiments, they respond faster if pictures have the agent on the left and if the action proceeds from left to right (Chatterjee et al. 1999).

This schema appears in other cognitive domains such as memory and even aesthetics. Representational momentum alluded to by Pylyshyn is not symmetric. Subjects recall action pictures as having proceeded further along if depicted going left to right than right to left (Halpern & Kelly 1993). Children, including Israeli and Arab children, draw depth relations in pictures with the nearer and more salient object on the left (Braine et al. 1993). Subjects judge pictures with implied left to right motion as more pleasing than images without such implied motion (Christman & Pinger 1997; Freimuth & Wapner 1979).

Curiously, portrait paintings in the western canon seem subject to this schema. Portraits are usually painted from an angle so that more of one side of the face is depicted (McManus & Humphrey 1973). The gender, social status, and personal characteristics of the subjects influence these orientation biases. Hypotheses invoking motor tendencies in right-handed artists, effects of maternal imprinting, hemispheric specialization for emotional expression or facial perception have been proposed, but they do not adequately explain the observations (Chatterjee 2002). If agents are implicitly conceived on the left, displaying more of the right side of their face, the bias to show the right versus the left side of the face reflects the extent to which portrait subjects are considered agents or recipients of actions. For example, arguably in the fifteenth century women were considered less “agent-like” than men. A bias to show the left more than right cheek was far greater for portraits of woman than men in that period. This bias diminished over time, and is absent in portraits of women sovereigns (Grusser et al. 1988). The notion of agency in the spatial schema offers a parsimonious explanation for these painting tendencies (Chatterjee 2002).

Whether the features of this schema are determined by properties of hemispheric specialization or by cultural reading habits is not clear. Even if fashioned by cultural habits, their influence across such disparate domains makes them unlikely to be epiphenomenal. Note that this schema does not seem subject to “tacit knowledge.” Regardless of how narcissistic we are, humans prob-

ably don't believe that actions in the world proceed from left to right anchored to our specific viewpoint.

In summary, primitive schemas may be nestled between depictive and descriptive mental representations. This claim does not address directly the picture in the head proposal, but does address a question that seems to be at stake: Does human cognition engage different forms of representation including some with analog properties?

The nature of mental imagery: How null is the “null hypothesis”?

Gianfranco Dalla Barba,^a Victor Rosenthal,^a and Yves-Marie Visetti^b

^aINSERM Unit 324, Centre Paul Broca, 75014 Paris, France; ^bLATTICE-CNRS-ENS, 92120 Montrouge, France. dallabarba@broca.inserm.fr vicorro@broca.inserm.fr yves-marie.visetti@ens.fr

Abstract: Is mental imagery pictorial? In Pylyshyn's view no empirical data provides convincing support to the “pictorial” hypothesis of mental imagery. Phenomenology, Pylyshyn says, is deeply deceiving and offers no explanation of why and how mental imagery occurs. We suggest that Pylyshyn mistakes phenomenology for what it never pretended to be. Phenomenological evidence, if properly considered, shows that mental imagery may indeed be pictorial, though not in the way that mimics visual perception. Moreover, Pylyshyn claims that the “pictorial hypothesis” is flawed because the interpretation of “picture-like” objects in mental imagery takes a *homunculus*. However, the same point can be objected to Pylyshyn's own conclusion: if imagistic reasoning involves the same mechanisms and the same forms of representation as those that are involved in general reasoning, if they operate on symbol-based representations of the kind recommended by Pylyshyn (1984) and Fodor (1975), don't we need a *phenomenological homunculus* to tell an imagined bear from the real one?

The central argument in Pylyshyn's essay is that mental imagery does not need to be pictorial and that there is not enough evidence to reject the “null hypothesis,” that “the process of imagistic reasoning involves the same mechanisms and the same forms of representation as are involved in general reasoning” (sect. 1.2, para. 2). In addition, he claims that we are deeply deceived by our subjective experience of mental imagery and that phenomenology “as explanatory does not help to understand why and how the behavior [mental imagery] occurs” (sect. 4.4, last para.). The author is probably right saying that phenomenology does not explain why and how mental imagery occurs, but phenomenology has never aimed at causal explanation. Phenomenology is descriptive, it tells us what imagery *is*, what it is like to entertain a mental image, and what purpose it serves in lived human experience. No scientific perspective, whether neurobiological or cognitive, can thus abstain from phenomenological description. The latter may even serve an explanatory purpose telling us, for instance, how mental imagery differs from visual perception. And indeed what phenomenological evidence tells us is that mental imagery may be pictorial, though not in the format of visual perception.

It has been argued (Dalla Barba 2002), following Jean-Paul Sartre's work (Sartre 1940), that mental imagery is a specific form of consciousness, a subjective experience that cannot be mistaken for the experience of *perceiving* something. “Imaginative consciousness” reflects a particular way of addressing the world. The object of perception and that of imagination may be identical: a pack of cigarettes, for example. But consciousness places itself in a relationship with the object in two different modes. The image entertained in consciousness is not a faded reproduction of a real object, as has long been thought, but reflects an original relationship between consciousness and the object. Imaginative consciousness and visual perception differ in nature and serve different goals. The latter yields partial, progressive presentation of the object while the object of imagination is present in its entirety.

Furthermore, while visual perception aims at exploring the outside world, the aim of imaginative consciousness is neither to know what is there, nor to explore it. Knowing and imagining also rely on a different relationship to the world. The relationship that perceptual consciousness entertains with its object is that of “realising”: we see the object of perception as real. The relationship of imaginative consciousness to its object is, on the contrary, that of “nonrealising”: we imagine the object as unreal, as *imagined*.

Note also that the object of visual perception continually reveals and hides itself. The more I observe this pack of cigarettes the more I notice details which escaped me before. In contrast, the object present in imaginative consciousness reveals nothing, for whatever is in the image is *already revealed*. “I can keep looking at an image for as long as I wish: I will never find anything but what I put there,” Sartre (1940) says. This is to say that in visual perception the object constantly goes beyond consciousness, while in imagination the object is nothing but consciousness. Thus, for example, while my visual perception can betray me, my imaginative consciousness cannot. I may happen to mistake this pack of cigarettes for a box of sweets. No such error can be made by imaginative consciousness: if I imagine a box of sweets, it is the box of sweets that I imagine, not a pack of cigarettes which erroneously appears to me as a box of sweets. Sartre called this quasi-observation our *disposition* towards the imagined object. We may inspect, scrutinize our imagination, but this will teach us nothing new, add nothing to what is already there. The world of mental images is a world where nothing *happens*. “I can make one or other object evolve into an image as I wish, I can rotate a cube, make a plant grow, a horse run, there will never be the slightest difference between the object and consciousness. Not an instant of surprise” (Sartre 1940).

The foregoing discussion clearly provides an illustration of what phenomenology can teach us about mental imagery and about the ways it differs from visual perception. It also shows that mental imagery may be pictorial, at least in the sense that in mental imagery “we see something.” It is not phenomenology that ever suggested that when we close our eyes and imagine a scene we engage in an act of visual perception minus the stimulus on the retina.

Where do we stand then? Pylyshyn is certainly right when he says that the available empirical data provide no convincing support to the “pictorial hypothesis” of mental imagery. He may even be right when he claims that cognitive neuroscience is irrelevant to this issue. There is, however, phenomenological evidence which shows that mental imagery may be pictorial, though clearly not in the way that mimics visual perception. The “pictorial hypothesis” is, in Pylyshyn’s view, flawed because the interpretation of “pictorial-like” objects in mental imagery takes a “*homunculus*,” an unconscious subject who “sees” mental images and interprets them as mental images. However, the same point can be objected to Pylyshyn’s own conclusion: If imagistic reasoning involves the same mechanisms and the same forms of representation as those that are involved in general reasoning, if they operate on symbol-based representations of the kind recommended by Pylyshyn (1984) and Fodor (1975), don’t we need a *phenomenological homunculus* to tell an imagined bear from the real bear? Better trust yourself.

Mental imagery: In search of my theory

Edward de Haan and André Aleman

Neuropsychology Section, Helmholtz Research Institute Utrecht and University Medical Centre Utrecht, 3584 CS Utrecht, The Netherlands.
e.dehaan@fss.uu.nl <http://www.fys.ruu.nl/~wwwfm/>

Abstract: We argue that the field has moved forward from the old debate about “analogical” versus “symbolic” processing. First, it is questionable that there is a strong a priori argument for assuming a common processing mode. Second, we explore the possibility that imagery is not a unitary mental function. Finally, we discuss the empirical basis of the involvement of primary areas.

So here we are, apparently some thirty years after the “Imagery Debate” commenced, and one of the original contenders stands up to deny his rival victory. It transpires that we have not even been able to decide on the playing field: the task was to defy the null hypothesis. In essence, the target article by Pylyshyn on the psychological and neuro-anatomical basis of imagery is timely. However, some of us who are interested in mental imagery do not subscribe to the view that there is a line in the sand, and that our research was aimed at deciding at which side the truth falls. It is our conviction that we have passed this stage and that recent research allows us to formulate new theoretical ideas concerning how we are able to mentally imagine the outside world. In this commentary, we would like to focus on three important issues. First, we question the suggestion that there is a strong a priori argument for assuming a central or common processing mode. Second, we want to explore the possibility that imagery is not a unitary mental function, and finally, we would like to discuss the empirical basis of the involvement of primary visual cortical areas in visual imagery.

First, to be controversial, let us start with the observation that imagery is a mental function that is close to perception. Indeed, it has been suggested that mental imagery is a trial run of what it would be like to experience such a perception. (An instructive observation is that some neurosurgeons practice complex operations via imagery; James Reason, personal communication.) In adaptive terms, it would be useful to develop this ability in order to train and sharpen the system in the comfort of safety. Thus, if it is reasonable to assume that imagery constitutes a (perceptual) ability that we share with other species, it follows that the mental processes that are responsible for imagery are related to the perceptual mechanisms. Certainly, most animals do not appear to have access to general – perception-independent – representations for reasoning.

Another reason to question the claim for single mental mode representation coding (the null hypothesis) is the observation that memory appears to have a category specific build-up. It has been demonstrated that our knowledge of the outside world is organised according to different categories and that these different categories are instantiated in different brain areas (Tranel et al. 1997). These observations are in line with the proposals of a distributed organisation of our memory, with the hippocampus as a central reference system that re-activates primary areas (Rolls 2000). Thus, whatever LISP and Logic gave to the world, there appears to be – at least – as much circumstantial evidence (and that is really what we are talking about here) to assume that representational coding reflects sensory processing.

Second, the target article appears to neglect a body of neuropsychological literature indicating selective deficits within the realm of imagery. An interesting double dissociation put forward by Ziyah Mehta (Mehta et al. 1992; Mehta & Newcombe 1996) is the distinction between the ability to image a verbal code, that is, text and pictorial information. Both patients experienced severe problems in visualising the appearance of common stimuli, but the type of stimulus was crucial. One of the patients, MS, who also suffered from severe object agnosia, was very poor at mental imaging objects but he could comment on the visual appearance of

words. In contrast, the patient SM showed the opposite pattern. Other evidence for the multi-dimensionality of mental imagery comes from specific face imagery deficits (Young et al. 1994), and studies in which attention-based and memory-based imagery processing were clearly distinct from each other on empirical grounds (Kosslyn 1994).

Third, it has been suggested that a dissociation between depictive and spatial images might explain discrepancies in the neuroimaging literature (Kosslyn et al. 1999a). Specifically, according to these theorists, tasks involving imagery of object forms at a high resolution would involve the primary cortex, whereas imagery tasks that require a spatial decision would rely on the parietal cortex. Indeed, the very vivid imagery associated with synesthesia, has been shown to activate area V1 in the absence of visual stimulation (Aleman et al. 2001a). Transcranial magnetic stimulation, being a “virtual lesion method,” can provide more decisive answers to these questions than functional brain imaging, where one does not know whether activation is essential for task performance. To our knowledge, two transcranial magnetic stimulation (TMS) studies of visual mental imagery have been conducted, at different laboratories, one using a depictive task (Kosslyn et al. 1999a) and one using a spatial task (Aleman et al. 2002). Consistent with the prediction above, repetitive TMS (rTMS) targeted at the primary visual cortex disrupted performance in depictive imagery, but not in spatial imagery. In the latter study a significant effect was observed of rTMS over the right posterior parietal cortex.

From the position considered above, interesting new questions can be generated, such as the one formulated by Trojano et al. (2000), “Do spatial operations on mental images and those on visually presented material share the same neural substrate?” This approach contrasts conditions that target the same cognitive process but differ in whether the stimuli on which the mental operations are performed have a top-down versus bottom-up genesis. A similar approach was recently taken in a study reporting evidence that congenitally totally blind people are able to perform not only “visuospatial imagery” tasks, but also tasks involving imagery of the visual shape of objects (Aleman et al. 2001b). This finding is at odds with both the pictorialist position and the explanation in terms of demand characteristics forwarded by Pylyshyn (i.e., that tacit knowledge is used to simulate what would happen in a visual situation).

Moving beyond the classical imagery debate, contemporary research should focus on common pathways in which representations of different sensory origin converge, the interaction between top-down and bottom-up processing, and the integration of attention, perception, and performance for which the brain was probably built in the first place.

Does your brain use the images in it, and if so, how?

Daniel C. Dennett

Center for Cognitive Studies, Tufts University, Medford MA 02155.
ddennett@tufts.edu <http://ase.tufts.edu/cogstud/>

Abstract: The presence of spatial patterns of activity in the brain is suggestive of image-exploiting processes in vision and mental imagery, but not conclusive. Only behavioral evidence can confirm or disconfirm hypotheses about whether, and how, the brain uses images in its information-processing, and the arguments based on such evidence are still inconclusive.

Nobody denies that when we engage in mental imagery we *seem* to be making pictures in our heads – in some sense. The question is: Are we really? That is, do the processes occurring in our brains have any of the properties of pictures? More pointedly, do those processes *exploit* any of the properties of pictures? When you make a long-distance telephone call, there is a zigzag pattern of activity running through various media from you to your listener

across the country, but if the curves and loops and angles happened to spell out “Happy Birthday” (as seen from space), this would be an image on the surface of the planet that was not exploited in any way by your information-transmission, even if it was a birthday greeting.

Consider the mindless doodling that involves filling in all the closed loops on a printed page with your pen. The word “doodling” would get ink on every letter except the “lin” group, wouldn’t it? That process depends only on the image properties of the text, and not at all on the meaning, or even on the identity of the symbols. You can perform it just as easily on printed text in any language. Now consider spell-checking. You can’t spell-check a bit-map picture of a page of text. You have to run it through OCR (optical character recognition) first, changing the categories from shapes of black-on-white to strings of alphabet characters. Is the resulting data structure an image or not? Since there is no canonical and agreed-upon list of image-exploiting *processes*, this is an ill-posed question. In some ways it is (like) an image and in some ways it is not. The processes that can extract depth in a random dot stereogram are, like the doodling process, strongly imagistic – in one sense: they are totally dependent on the topographical properties of the pattern of stimulation, and not at all on the content thereby represented (there being none until after the shape in depth has been extracted). The processes that can “rotate images” à la Shepard and Metzler are, in contrast, strongly dependent on previously extracted content (try rotating half a random-dot stereogram), so they are a bit more like spell-checking, a bit less a matter of “brute” image-processing.

What sometimes looks like deliberate shifting of the goal-posts in the long-running debate over imagery is better seen, I think, as the gradual clarification of the ill-defined question above. But confusion and talking-past-each-other persists. As Pylyshyn stresses, the evidence from neuroimaging studies is, so far, almost irrelevant to the points of contention. The presence of readable images of activity in the brain is *suggestive* of image-exploiting processes, probably a practical necessity for such processes to occur, but not conclusive. As Kosslyn (1994, p. 80) notes, in the long passage quoted by Pylyshyn, the issue is about a *functional* space, not necessarily a physical space. It’s like computer graphics. As long as the data structures consist of properly addressed registers over which the operations are defined, the activity can be arbitrarily scattered around in space in the computer’s memory without hindering the image-exploitation that is going on.

I explain this to my students with a little thought experiment: Dismantle a mosaic, tile by tile, numbering each space consecutively, line by line, and mailing the tiles individually to friends all over the world, writing the address to which each tile is mailed after the number on a long list (a list, *not* a map). The mosaic on the floor is gone, the physical image destroyed. Then ask yourself questions such as: “Are there any strings of four black tiles in a horizontal row surrounded by white?” This can be answered, laboriously, by asking all your pen pals to send a “Yes” message if they have been sent a black tile, and then analyzing the list to see if any four consecutive “Yes” answers show up, and then calculating which pen pals (the “neighbors” of those four, wherever in the world they are) to ask if their tiles are white. The physical image on the floor has been destroyed, but the information in it is all available for image-exploiting processes to work on. You might *happen* to mail neighboring tiles in the mosaic to friends who lived near each other, but the system’s operation does not at all depend on this coincidence, however convenient it might be in practical terms. The brain, needing to work fast with rather slow connective fibers, probably preserves as much geographical correspondence as possible – the retinotopic maps – for exploitation in such inquiry-processes.

Probably the brain’s preservation of topological relationships is no mere byproduct of thrift-in-wiring, but also an enabling condition for image-exploiting processes of information-extraction. Now, can we prove it? That is what Pylyshyn has been asking all these years, and as he says, the answers so far largely fail to come

to grips with the logical requirements for a positive answer. In particular, one cannot establish that mental imagery involves image-exploiting processes by showing that it utilizes the brain's vision systems, because it has yet to be established when and how *vision* utilizes image-processing! Vision isn't television. The product of vision is not a picture on the screen in the Cartesian Theater (Dennett 1991). The fleeting retinal images punctuated by saccades are the first images, and they are not the last, as Julesz (1971) demonstrated by showing perception of depth in random dot stereograms that requires image-processing after the optic chiasma. But which subsequent cortical processes also exploit any of the informational properties of images? The *eventual* "products" of vision are such things as guided hand and finger motions, involuntary ducking, exclamations of surprise, triggering of ancient memories, sexual arousal, . . . and none of these is imagistic in any sense, so assuming that the events in their proximal causal ancestry are imagistic is rather like assuming that power from a hydroelectric plant is apt to be wetter and less radioactive than power from a nuclear plant. The raw retinal data are cooked in many ways between eyeball and verbal report (for instance). How cooked are the processes involved in (deliberate or voluntary) mental imagery? We don't know yet, though investigations are gradually peeling away the alternatives.

As Pilyshyn says, behavioral evidence – patterns of ease and difficulty, timing and vulnerability to disruption, and the like – is the only evidence that can show that, and how, the brain *uses* images in its activities. To organize the evidence, we use the heterophenomenological method (Dennett 1982; 1991) to provide a neutral catalog of *how it seems to subjects* under many varied conditions, and then our task is to devise and confirm theories that predict and explain all that seeming.

Interpreting the neuroscience of imagery

Ian Gold

Philosophy and Bioethics, Monash University, Clayton, VIC 3800, Australia.

ian.gold@arts.monash.edu.au

<http://www.arts.monash.edu.au/phil/department/gold/>

Abstract: Pilyshyn rightly argues that the neuroscientific data supporting the involvement of the visual system in mental imagery is largely irrelevant to the question of the format of imagistic representation. The purpose of this commentary is to support this claim with a further argument.

Pilyshyn's paper provides a review of some of the neuroscientific evidence concerning the overlap between vision and imagery. The findings are of considerable interest on their own, but, as Pilyshyn argues, they are of special interest to proponents of the pictorial view. If – the implicit argument goes – vision and mental imagery overlap, then imagery, like vision, must be depictive. Pilyshyn's response to this argument, however, is exactly right: Even if the evidence supports the overlap between vision and imagery, the argument depends on the assumption that the format of visual representation is depictive. But it is quite possible that both vision and imagery share an underlying form of representation that is non-pictorial.

This seems a very obvious point, and the fact that it is regularly overlooked suggests that it never occurs to most investigators to doubt that visual representation is pictorial. What makes this unargued for assumption so compelling? I will suggest that it is based on the error, familiar to philosophers, of confusing the properties of the represented thing with the properties of the representing medium. This error is one that is easy to make, and the natural tendency to make it is, I suggest, the reason why the assumption that vision is pictorial is rarely questioned.

The error of interest is one that arises in the context of "qualia," the putative properties of an experience that give the experience its felt qualities. One reason for thinking there are such things is

that this is what our experience apparently teaches us. Reflect on your experience of a strawberry and you will simply *observe* the "feel" (i.e., felt-quality) associated with its red color. The "feel" is not part of the strawberry; as a "feel," it must be in the mind. It seems to follow that the experience of the strawberry must have a property in virtue of which the experience of the strawberry has its characteristic "feel."

This argument moves from a feature of experiences – that they have felt-qualities – to a claim about the ontology of experience – that experiences have special properties. But, as has often been pointed out,¹ this argument depends on a confusion between what is a property of an experience and what is a property represented *by* the experience. When one is asked to describe the property of the experience of the strawberry that is its red feel, one inevitably ends up describing a feature of the strawberry itself and not the experience. Experience, it is sometimes said, is "diaphanous": one sees through it to the object or property the experience is representing. The experience itself has no properties accessible to the experienter. Therefore, the inference from the phenomenology of experience to the ontology of qualia depends on a confusion between properties of objects represented – the redness of a strawberry, for example – and properties of the experiences that do the representing – the putative redness of the experience itself. Once the confusion is pointed out, the argument for the existence of qualia – *this* argument, at any rate – can be seen to be flawed.

I believe that this same mistake is the one that leads to the tendency to assume that visual experience must be pictorial. The mistaken inference can be reconstructed as follows: visual experience is picture-like: it represents objects as having spatial properties (among other things); therefore, the underlying representational form of vision is depictive. When put this way, the mistake is obvious. From the fact that visual experience represents objects as having spatial properties, one is supposed to draw the inference that the representing experience also has spatial properties. But this is no more valid than the inference from the fact that the sentence "The truck is larger than the car" represents spatial properties, to the claim that language is a spatial medium.

Once the error is made, however, it provides the needed assumption for the implicit argument above that moves from the overlap between the neuroscience of vision and of imagery to the claim that the representational form of imagery is depictive. In fact, however, the spatial nature of the objects represented visually says nothing whatsoever about the properties of the representing medium of visual experience. The similarity between vision and imagery, therefore, entails nothing about the properties of the representing medium of imagery.

One might be inclined to object that while it might be correct to deny that the representational medium of vision is depictive, one cannot make the same claim about imagery. The reason is that while there can be a confusion between the properties of the objects that visual experience represents and the properties of visual experience itself, the same conflation cannot arise in the case of imagery because imagery need not represent objects at all. If I visually image the strawberry I have just eaten, how can I confuse the properties of my experience with the properties of the thing represented by that experience? There is no thing whose properties can be confused with the properties of my imagistic experience.

Here, however, the similarity of visual experience and imagistic experience works *against* the picture theorist. Consider the following case (see Lewis 1980): I am looking at a strawberry on the table in front of me. Unbeknownst to me, a mad scientist is manipulating my brain as well as the objects in front of me. At the very same instant, he destroys the strawberry and preserves the state of my visual system so that I experience the scene as unchanged. At one moment, my visual experience represents an object; at the next it is a vivid image or hallucination. It makes little sense, however, to think that I have access to information about the properties of the representing medium after the strawberry is destroyed, when I fail to have this while the strawberry is in exist-

tence. The existence of the strawberry has no bearing on my access to my own mental states. If visual phenomenology reveals nothing about the ontology representation, there is no reason to think that imagery does.

NOTE

1. The modern classic is Harman (1990).

Loss of visual imagery: Neuropsychological evidence in search for a theory

Georg Goldenberg

Neuropsychological Department, Munich Bogenhausen Hospital, D 81925 Munich, Germany. Georg.Goldenberg@extern.lrz-muenchen.de

Abstract: Observations on patients who lost visual imagery after brain damage call into question the notion that the knowledge subserving visual imagery is “tacit.” Dissociations between deficient imagery and preserved recognition of objects suggest that imagery is exclusively based on explicit knowledge, whereas retrieval of “tacit” visual knowledge is bound to the presence of the object and the task of recognizing it.

Pylyshyn concludes that neuropsychological evidence does not support the contention that mental images are based on retinotopically organized neural representations. This argument is convincing but does not exhaust the contribution of neuropsychology to the theory of mental images. Observations of patients who lost visual imagery after brain damage call for a refinement or revision of the “tacit knowledge” hypothesis, too.

There are at least five visual categories for which imagery can be selectively affected by brain damage: shapes and colors of common objects, shapes of faces, shapes of letters, and topographical relationships (review in Goldenberg 1993). A patient who is unable to answer imagery questions about the shape of the ears or the length of the tail of animals (Kosslyn 1983) may do perfectly well on imagery questions like those shown in Figure 4 of the target article concerning the shape of letters (Goldenberg 1992). Such dissociations can hardly be explained by damage to a visual buffer or any other structure subserving generation of visual images independent of their content. The more likely hypothesis that these patients have lost knowledge of the visual appearance of only one category of things calls for a theory of that knowledge. How is it organized that it can break down for only one category? How is it related to knowledge of non-visual properties?

Another challenge to the tacit knowledge hypothesis is constituted by patients with loss of visual imagery and preserved visual recognition (Basso et al. 1980; Farah et al. 1988; Goldenberg 1992). The proposal that these patients have preserved knowledge of the visual appearance of objects, but are unable to employ an “image generation process” transforming knowledge into mental pictures (Farah 1984), has been criticized on two grounds. First, as already mentioned, the imagery deficit can be restricted to only certain categories of things. Second, it has been shown that these patients make errors when they are shown pictorial versions of imagery questions, although in this condition the crucial images are before them and need not be generated before the “mind’s eye.” For example, when shown images of bears with rounded and with pointed ears they cannot decide which of them is correct. Obviously, they lack knowledge of the shape of the bear’s ears. Nonetheless they readily recognize that these are bears. Their visual recognition must have access to knowledge of the global shape and the characteristic features of bears to distinguish them from lions or dogs. But they are completely unable to imagine the visual appearance of a bear and not just that of its ears! The knowledge they use in recognition cannot be used for imagery.

Based on these lines of evidence I proposed that there are two kinds of knowledge of the visual appearance of things (Goldenberg 1992; 1998; Goldenberg & Artner 1991). Knowledge used in

recognition is restricted to those features which permit a reliable identification of an object under varying circumstances. It neglects details like the shape of the bear’s ears. There is a second store of visual knowledge within semantic memory. This knowledge includes information on features not necessary for recognition in addition to those used for recognition. The source of semantic visual knowledge may be an active interest in the visual appearance of objects, possibly enhanced by the high value given to visual arts in our culture and education (Armstrong 1996; Farah 1995a). The crucial point of this hypothesis is that knowledge used for visual recognition is completely embedded in visual recognition and cannot be used for any other purpose. Visual imagery is based exclusively on visual knowledge within semantic memory. If only this knowledge is lost, patients are unable to imagine the visual appearance of objects although the knowledge embedded in recognition enables them to recognize the same objects.

This hypothesis calls into question the idea that visual imagery is based on “tacit” knowledge. Pylyshyn states that “knowledge is called ‘tacit’ because it is not always explicitly available for . . . answering questions” (target article, sect. 3.1). Presumably “not always” means that retrieval is bound to a certain context or task. This applies to the knowledge used for recognition: Its retrieval is bound to the presence of the object and to the task of recognizing it. By contrast, you can form mental visual images in the absence of the object and in response to a wide variety of questions (or just for fun), that is, in principle, always!

I propose that visual imagery is equivalent to the explicit recall of semantic knowledge of the visual appearance of things. This position is not meant to be a theory of imagery, but a request for such a theory. An adequate theory of imagery should explain how such knowledge is acquired, how it is organized, and how it differs from knowledge of other properties of things. It seems to me that imagery is still in search of a theory.

You are about to see pictorial representations!

Frédéric Gosselin^a and Philippe G. Schyns^b

^aDépartement de Psychologie, Université de Montréal, Montréal QC, H3C 3J7, Canada; ^bDepartment of Psychology, University of Glasgow, Glasgow G12 8QB, United Kingdom. frederic.gosselin@umontreal.ca philippe@psy.gla.ac.uk <http://mapageweb.umontreal.ca/gosseliff/cv.html>

Abstract: Pylyshyn argues against representations with pictorial properties that would be superimposed on a scene. We present evidence against this view, and a new method to depict pictorial properties. We propose a continuum between the top-down generation of internal signals (imagery) and the bottom-up signals from the outside world. Along the continuum, superstitious perceptions provide a method to tackle representational issues.

In a memorable courtship scene from the movie “A beautiful mind,” John Nash asks his future wife to think of an object. “Anything!” he says. She chooses an umbrella. He then turns toward the starry sky and, connecting some stars one by one with his finger, shows her a sparse, but nonetheless recognizable, umbrella. You might not be capable of performing this feat on demand, but you have surely seen sparse versions of objects or scenes in the sky or elsewhere at one time or another. On a continuum extending from pure top-down mental images (internal signals) to strong bottom-up signals, these extremely sparse objects (we call them *superstitious perceptions* in reference to Skinner’s celebrated 1948 article) are closer to mental images than extraneous signals. More importantly, we will demonstrate that they provide a powerful analytic tool to address the issue of internal representations.

We have recently produced a situation similar to the “umbrella in the stars” in our laboratories (Gosselin & Schyns, in press). In

one experiment, we instructed two observers (MJ and NL) to detect the presence of a letter “S” (for Superstitious) inserted in white noise (black and white pixels peppered across the image field). The observers were instructed that the letter “S” was black on a white background, filled the image, and was present on 50% of 20,000 trials. No more detail was given regarding the attributes of the letter. Unbeknownst to the observers, each trial only consisted in the presentation of a 50 x 50 pixels white noise image (see our Fig. 1a, for one example) with a black-pixel density of 50%. Crucially, no bottom-up external signal (i.e., an “S”) was ever presented.

At first, the observers found the task rather difficult, but, soon, they said, they responded with ease. In fact, observer NL said that after about 1,000 trials the “S” popped out when it was present. In any case, the observers detected an “S” in noise on 46% (NL) and 11% (MJ) of the trials, respectively.

What did the observers respond to? As already stated, no external signal was ever presented, and the observer only saw white noise. One possibility is that observers generated an internal signal via imagery, and tried to superimpose this signal onto the incoming white noise. Sometimes, this internal signal will be weakly correlated (here, a correlation smaller than .026) with the external white noise and the observer will detect the letter corresponding to his or her imagined signal. On the remaining trials, the mismatch will simply be too large and the observer will reject the noise as being what it is – noise. However, and this is important to stress, the observer must first generate an internal signal via imagery to be able to perform this detection task, and attempt to superimpose this internal signal to external noise. What is the internal signal of the observer? We will contend that whatever it is, it represents pictorial properties of the imagined letter.

From Wiener (1958), we know that systematic responses of a black box to white noise can be used to analyze its behavior. We are thus looking for a systematic correlation between the noise fields (x_i) and the detection responses (y). This is what *reverse correlation* does (see also Ahumada & Lovell 1971). The first Wiener kernel (the linear component) is equal to $k^{-1} \sum_t y(t) x_t$, where k is a constant and t is variables indexing all the trials. Leaving aside k , this amounts to subtracting the sum of all the noise fields that led to a rejection response from the sum of all the other noise fields (see Fig. 1a, NL and MJ). For each observer, we best-fitted a Gaussian density function (see Fig. 1b, the solid lines) to the energy distribution of his or her first order Wiener kernel (Fig. 1b, the open circles). This kernel (called the “classification image”) represents the template of information that drives the detection of the target “S” letter for this observer. In other words, the first order kernel provides a first approximation of the representation of the imagined internal signal for the letter “S.” To better visualize this representation, we sought an information peak in a spectral analysis of the kernel, and filtered out all spatial frequencies one standard deviation away from the mean (i.e., keeping a bandwidth of 0–3 cycles per letter). The outcomes are black “S”s on a white background filling the image (see Fig. 1c, NL and MJ).

The first order kernel predicts the detection response from each pixel, individually. However, it is likely that observers used higher order relationships between the elements of the internal signal – for example, combinations of two pixels. The second Wiener kernel examines what these second order relationships are. It is equal to $k^{-2} \sum_t y(t) x_t x_t$. Leaving aside k , this is equivalent to subtracting the sum of all the autocorrelations of the other noise fields (i.e., the outer product of each noise field vector with itself) that led to a rejection response from the sum of all the autocorrelations of the other noise fields. Figure 1f (NL and MJ) shows the regions of the second-order kernels that are statistically significant ($p < .01$). The number of significant regions far exceeds what would be expected by chance for both observers (937 pixels for NL, $p < .01$, and 1,318 pixels for MJ, $p < .01$), revealing that the imagined internal signal imagined *did* include nonlinear relationships.

What conclusions can be drawn from this study? We have in-




Figure 1 (Gosselin & Schyns). Adapted from Gosselin and Schyns (in press, Experiment 1). (a) Raw first order Wiener kernels. (b) Distributions of the average squared amplitude energy for different spatial frequencies (collapsed across all orientations) of (a) (expected energy = constant). The solid lines are the best Gaussian fits. (c) (a) filtered with a smooth Butterfield low-pass. We squeezed pixel intensities within two standard deviations from the mean. (d) Best matches between (c) and 11,284 letters. (e) Raw second order Wiener kernels. (f) Statistically significant ($p < .01$) pixels of (e).

duced superstitious perceptions of an “S” by instructing observers to detect this letter in noise. Unknown to them, the stimuli never comprised the letter, but only white noise. If the observers had been performing only according to an external signal (i.e., in a bottom-up manner), their kernels should have had the same properties as averaged white noise – that is, zero energy across all spatial frequencies. However, there was a marked peak of energy between 1 and 3 cycles per letter that could only arise from top-down influences arising from an internally generated signal – that is, a mental image. Further analyses revealed the properties of the internal signal driving the detection behavior. With white noise as inputs, the revealed letter could only depict the observer’s imagined letter “S.”

Is the internal signal pictorial in nature? Functionally, yes, because, if not from a matched internal signal, where else would the pictorial properties present in the kernels come from? Does this imply that the observers actually used an image of a “S” from their memory? Not necessarily, but they had to have knowledge of *all* the pictorial characteristics of an “S,” functionally isomorphic to an actual image of an “S.” We believe that you have just seen representations with pictorial properties!

Mental imagery during sleep

Claude Gottesmann

Laboratoire de Neurobiologie Comportementale, Faculté des Sciences,
Université de Nice-Sophia Antipolis, 06108 Nice cedex 2, France.

Gottesma@unice.fr <http://www.unice.fr/neurobiol>

Abstract: The descriptive “null” hypothesis is strengthened by the fact that during dreaming sleep stage, the primary visual cortex is deactivated as compared with other sleep stages.

The topic of this research field seems to be posited in terms of the thesis and antithesis. Mental imagery involves either inspection of a brain-induced picture-like object (depictive theory) (Kosslyn et al. 1999a) or “the same processes as that of reasoning in general, except that the content or subject matter of thoughts experienced as images includes information about how things would look” (target article, Abstract). Are these two theories inevitably antinomical, and necessarily exclusive? When the author writes of “image-based thinking” rather than what could have been “thinking-based imagery” in view of his theory, one realises the complexity of the subject. I would like to make some comments on mental imagery during sleep to possibly fuel the discussion.

The brain is never silent psychologically. Not only during waking, but also during each sleep stage, the brain gives rise to psychological content, the nature of which varies at different stages. Already at sleep onset there are hypnagogic hallucinations. These consist of “floating sensations, flashing lights, lantern slide phenomena, fleeting progressions of thoughts and images” (Foulkes 1962), during which the sleeper is rather passive and an onlooker. The visual system is involved although there is no clearly organized mental background. These visual phenomena are difficult to interpret in respect of the two hypotheses of mental imagery. During the true slow wave sleep which follows, stages II, III, IV, the main results show moderate “thought-like” contents (Fosse et al. 2001; Foulkes 1962) which follow the rules of the secondary process (Freud 1875/1975) involving the principle of reality (Freud 1911) as during waking. Although imagery contents have been described by several authors (Bosinelli 1995; Cavallero et al. 1992; Foulkes 1962; Tracy & Tracy 1974), recent results suggest that genuine dreams, with their characteristic visual components, can only occur against the physiological background of rapid eye movement (REM) sleep (Gottesmann 1999; Hobson et al. 2000; Nielsen 2000; Takeuchi et al. 2001), even if this stage is “covered,” that is, it does not show all its electroencephalographic (EEG) and peripheral characteristics (Nielsen 2000). Night terrors (different from nightmares) and somnambulism attacks also occur during slow wave sleep, all these data showing non-visual mental activity during sleep.

REM sleep is the real dreaming sleep stage. It is characterized by activation of the majority of brain structures involved in mentation. Several points related to the author’s target article have to be emphasized. Already, old results showed that the well-known rapid eye movements are the more numerous as the sleeper is active in the dream (Dement & Wolpert 1958); this could be related to the activation of the cortical saccadic eye movement system (Hong et al. 1995) and could explain the occasionally observed relationship between eye movements and dreaming content (Dement & Kleitman 1957; Herman et al. 1984; Roffwarg et al. 1962). It is difficult to determine whether this is the result of scanning a visual scene or of movements generated by theoretical thinking. However, although intuitively it would seem to suggest the first process, cerebral blood flow is decreased in the striate visual cortex during REM sleep eye movements while it is increased in extrastriate areas (Braun et al. 1998). Otherwise, there are pontogeniculo-occipital phasic (PGO)-like waves in humans (McCarley et al. 1983; Miyauchi et al. 1987) specifically related to the eye movements generated by pontine structures (Vanni-Mercier et al. 1996); these spikes were hypothesized to be the instigators of dreams (Hobson & McCarley 1977; McCarley et al. 1983; Miyau-

chi et al. 1987). Since the spikes end mainly in the visual cortex, this would suggest direct visual activation processes which would argue in favour of the depictive approach. Steriade et al. (1989) even stated that, since prior to REM sleep entrance there are very high amplitude isolated spikes (without eye movements), vivid imagery may occur during these short periods. Nevertheless, verbal reports obtained after awakening from this period do not reveal visual contents but “a feeling of indefinable discomfort, anxious perplexity and harrowing worry” (Lairy et al. 1968, p. 279). Also Larson and Foulkes (1969) showed that mental contents during this stage of sleep “are inconsistent with the hypothesis of an intensification of mental activity or cerebral vigilance” (p. 552). Moreover, in addition to these old data, which are the only ones available, the time scale of dreaming is ill-matched to PGO-wave lengths with a maximum of 100 milliseconds, unless we accept that the successive spikes are responsible for rapid changes of dream content, a hypothesis which currently seems unlikely (Gottesmann 2000). Their only probable consequence is to induce a transient higher activation of the posterior cortex areas (Satoh 1971). Coupled with the eye movement data they do not therefore support the depictive theory.

Nevertheless, all these results related to REM sleep do not exclude the depictive theory of mental imagery, since hallucinatory activity is the main characteristic of REM sleep mental activity (Fosse et al. 2001; Hobson et al. 1998). However, there is a major point which seems to contradict it: the primary visual cortex is deactivated throughout this sleep stage (Braun et al. 1998), a result which moreover reinforces Llinas and Ribary’s (1993) convergent results showing an absence of reset of gamma EEG activity (centered on 40 Hz) under sensory stimuli, unlike waking. Indirectly, it also partly confirms Crick and Koch’s (1995) hypothesis of dream forgetting. Whether picture-like objects would be elaborated, then this should occur in visual associative areas which are activated (Braun et al. 1998; Lövblad et al. 1999; Madsen et al. 1991). Today, it has not been definitively established whether picture-like depictive representations are possible in the extrastriate visual cortex but many authors suggest its participation in mental imagery (Braun et al. 1998; Mellet et al. 1998). Moreover, as postulated by the psychoanalytic model, although dreams give rise to vivid and generally precise imagery, these images (“manifest content”) are only the disguised visual display of “latent” psychological contents. For Freud (1900), the spatiality of the oneiric content is therefore no more than an auxiliary representation. This theory corroborates the descriptive “null” hypothesis. However, the results of sleep mechanisms still do not contribute decisive information to the delicate problem of mental imagery.

Functional versus real space: Is pictorialism hopeless?

Verena Gottschling

Philosophisches Seminar, University of Mainz, 55099 Mainz, Germany.
verena@gottschling-net.de

Abstract: Pylyshyn raises hot topics like the number and kinds of pictorialist theories there are and their explanatory power. Pylyshyn states that pictorialists have only two possibilities – they can posit either “only functional” images or “really spatial” images – and that neither of these possibilities is convincing or sufficient in explanatory power for empirical and theoretical reasons. Is pictorialism, in principle, untenable?

In Pylyshyn’s challenging target article there are two issues that I want to focus on: (1) The thesis that if depictive theories propose only functional space, there is no explanatory power; and (2) the thesis that if depictive theories propose a literal sense of “spatial extent,” there is no explanatory power and there are logical and conceptual problems.

Pylyshyn presents a challenging argument¹ for (1): Taking up

Kosslyn's (1981) example of a matrix to explain functional space, he argues that a matrix is in fact no example for functional space because the use of spatial notions is only found in our description of the matrix; the matrix itself is organized differently (sect. 5.2). There are no intrinsic constraints on how the cells in a matrix must be processed. And if we use extrinsic constraints there is no advantage at all because extrinsic constraints can be applied to any form of representation. That is why the explanatory power lies in the extrinsic constraints, not in the format of the representation.

What does this argument show? It only shows that the normal intuitive understanding of a matrix as an example of functional space doesn't work. However, it does not show that the idea of purely functional space does not work. Yes, we need intrinsic constraints to explain spatial characteristics of mental imagery and to explain the results of imagery studies. But the central question is: Are there no other possibilities that can be used to characterize functional space? Even if the current proposal does not work, we cannot conclude that the whole idea of functional space as such has no explanatory power. Thesis (2) presupposes that there is *no* interpretation of the functional thesis that can explain the data via intrinsic properties of the representing relations (sect. 7.2). However, I think that a possible alternative is hinted at in Note 6 of Pylshyn's target article: what we need is the representation of spatial relations via nonspatial relations. And, in addition, these relations have to satisfy the same inherent constraints as the spatial relations. A clearer understanding of what that could mean can be found in Palmer (1978) and Rehkaemper's (1991) concept of natural isomorphism.

The imagery debate would not be solved for Pylshyn even if we do find "real colored stereo pictures" on the visual cortex (sect. 7.4, para. 1). Why? Because, alongside several empirical arguments, there seem to be a number of important conceptual and systematic issues, as follows:

The point at issue for Pylshyn is whether early vision participates in imagery (sect. 6.5). The key notion is cognitive penetrability: Cognitive penetrability is the criterion for differentiating functional architecture and the cognitive system.² In contrast to high-level vision, early vision is not and cannot be cognitively penetrable, because there is no top-down processing from the cognitive system (see Pylshyn 1999, sect. 1.1). Early states of vision are therefore not sensitive to cognitive influences, that is, they are not cognitively penetrable. But imagery is cognitively penetrable, so early states of vision cannot be involved in imagery. Images cannot be located in early vision.

Kosslyn admits the cognitive penetrability of imagery. In his theory there are many subsystems on different levels of processing (e.g., early and higher level vision). However, even if the whole process is cognitively penetrable, and this process involves activities at many levels, we cannot conclude *that every part of this process has to have this property*. It is important to be aware that (at least Kosslyn's) pictorialism is a hierarchical theory. Images are subordinated to descriptive representations. It is even more complicated if you look closely: Sometimes an image is taken to be the *conjunction* of a quasi-pictorial component (in Kosslyn's terms; 1981, p. 213) and a descriptive component, stored in long-term memory, that is, only one part of an image is pictorial. Most people (like Pylshyn) mean by "image" only the alleged pictorial part of the representation.

Pylshyn's argument only works if we use "image" in the second meaning. But even then there is a close connection with the corresponding knowledge in long term memory, which is usually thought to be descriptive. For Pylshyn, bringing in conceptual complexity is the first step in the direction "where one gives the actual image less and less of an explanatory role" (sect. 7.2, para. 10). But that does not imply that every proposal of this kind has *no* explanatory power at all. If including conceptual information in a theory of imagery has this consequence, then no hierarchical depictive theory of imagery is in fact possible. But as far as I know all proponents of hierarchical pictorialism admit that they need

conceptual information. The whole idea of generating an image in short term memory from stored descriptive information takes that for granted.

Another important problem for the real display view is the homunculus problem, or in Dennett's terms the "Cartesian Theater" (Dennett 1997, p. 83). That is, the notion that there is a central place in the brain where experiences were *first* presented, and *then* analyzed and interpreted. In the reinterpretation debate this has remained the central problem too. In my opinion, the description of the process is itself wrong. Even in the primarily visual cortex there is no pure presentation of unanalyzed information without analysis. Analysis is an all-pervading functional process.

Further, pictorialists must show that the processes in the visual system interpret images in a way that *depends* on their retinotopic shape (sect. 7.1). That is a hard nut to crack. It means that pure neuroscientific findings can never solve the imagery debate, because all they can find are activities in brain areas. And these activities can be epiphenomenal.

To sum up: To defend thesis (1) we need more than the trivial discovery that the intuitive interpretation of functional space does not work. There is no principled reason why a proposal like the one suggested could not work. Central for thesis (2) is Pylshyn's criterion of cognitive penetrability and his conception of the basic assumptions purportedly shared by all true cognitive theories. In this conception there is no place for images. So proponents of real display pictorialism need to propose an alternative conception. This suggests to me that the situation for pictorialists of both kinds is not hopeless, rather, it remains very challenging.

ACKNOWLEDGMENT

Thanks to Thomas Metzinger and Jesse Prinz for comments.

NOTES

1. Even if he does not call it an argument.
2. In fact its status is more complicated: It is a necessary but not a sufficient condition (sect. 3.2).

Neural substrates of visual percepts, imagery, and hallucinations

Stephen Grossberg

Department of Cognitive and Neural Systems and Center for Adaptive Systems, Boston University, Boston, MA 02215. steve@bu.edu
<http://www.cns.bu.edu/Profiles/Grossberg>

Abstract: Recent neural models clarify many properties of mental imagery as part of the process whereby bottom-up visual information is influenced by top-down expectations, and how these expectations control visual attention. Volitional signals can transform modulatory top-down signals into supra-threshold imagery. Visual hallucinations can occur when the normal control of these volitional signals is lost.

Pylshyn quite rightly opposes any view of visual imagery that takes the form of a naive "picture theory" (see target article, sects. 1.1 and 5.1). He proposes instead to reduce experiences of imagery to a kind of "thinking" (sect. 1.1). His question: "Is the 'mind's eye' just like a real eye?" even leads him to derisively ask if the mind's eye has properties like a blind spot (sect. 5.1). People who would, in fact, view imagery as a "picture" in the mind come dangerously close to falling into the trap of naive realism. But by recoiling too far from this unsupportable extreme view about imagery, Pylshyn seems to embrace too much the "thinking" end of the dialog between "seeing" and "thinking."

To fully discuss how imagery relates to visual perception, one needs to consider all the facts that are known about vision and how they resemble or differ from those of imagery. Pylshyn provides a nice sample of such comparisons. My comments will summarize some conclusions drawn from neural models of visual perception.

These models gain their predictive force from their ability to quantitatively simulate perceptual data. The most recent models go so far as to quantitatively simulate the responses of identified cortical cells in known anatomical circuits and perceptual properties that they control; for example, Grossberg and Raizada (2000), Raizada and Grossberg (2001). These models shed light on many of the facts and issues raised by Pylyshyn, and suggest that deciding between thinking and seeing (or imagining) is not an either-or decision. Rather, there are bottom-up and top-down interactions between seeing and thinking, and the top-down interactions, in the absence of bottom-up data, can give rise to an experience of imagery when they are modulated by volition. Perhaps more important than these particulars, the models provide a theoretical rationale for why imagery exists, and constitute a rigorous theoretical framework in which it can be analyzed. Thus, I would contend that there is an emerging theory of imagery, but it is not a thing unto itself. Rather, it is part of a larger neural theory of visual seeing and thinking.

First, let me state some of the general conclusions from this theoretical work. The first model, called Adaptive Resonance Theory, or ART (Grossberg 1999b), suggests how brain mechanisms of learning, attention, and volition may give rise to mental imagery during normal behaviors, and to hallucinations during schizophrenia and other mental disorders. It is proposed that normal visual (and other) learning and memory are stabilized through the use of learned top-down expectations. These expectations learn prototypes that are capable of focusing attention upon the combinations of “critical features” that comprise conscious perceptual experiences. When top-down expectations are active in an attentional priming situation, they can modulate or sensitize their target cells to respond more effectively to matched bottom-up information. They cannot, however, fully activate these target cells. These predicted matching properties have been supported by neurophysiological experiments; for example, Bullier et al. (1996), Lamme et al. (1998), Reynolds et al. (1999), Sillito et al. (1994).

A recent embodiment of ART mechanisms within the laminar circuits of visual cortex, called the LAMINART model, suggests how the learned prototype is realized by the on-center of a top-down on-center off-surround network (Grossberg 1999a). The modulatory property of such a top-down expectation is achieved through a balance between top-down excitation and inhibition within the on-center. Volitional signals can shift the balance between excitation and inhibition to favor net excitatory activation. Such a volitionally-mediated shift enables top-down expectations, in the absence of supportive bottom-up inputs, to cause conscious experiences of imagery and inner speech, and thereby to enable fantasy and planning activities to occur. If these volitional signals become tonically hyperactive during a mental disorder, the top-down expectations can give rise to conscious experiences in the absence of bottom-up inputs and volition. These suprathreshold events help to explain key data properties about hallucinations (Grossberg 2000). The level of abstractness of learned prototypes may covary with the abstractness of imagery and hallucinatory content.

Given this theoretical context, the following remarks briefly respond to some of Pylyshyn's concerns about imagery: (1) Both bottom-up activation of visual percepts and top-down cognitively-activated and volitionally-modulated imagery are possible within the visual system. Seeing, imagery, and thinking are not mutually exclusive concepts. There is no contradiction in claiming that visually-based imagery exists and that it can be manipulated by cognitive constraints. (2) Visual representations are not like images on the retina. Rather, depthful boundary groupings and surface representations are formed through hierarchical and interstream interactions in areas V2 to V4 to represent occluding and occluded surfaces, both modally and amodally (Lamme et al. 1999; Schiller 1994). FACADE theory predicts that the final modal figure-ground separated visual representation is formed in V4 (Grossberg 1994, 1997; Kelly & Grossberg 2000). (3) Top-down expectations and attention operate at all levels of this hierarchy and can

reorganize cell properties using higher-level constraints; for example, Bullier et al. (1996), Lamme et al. (1998), Reynolds et al. (1999), Sillito et al. (1994). (3) When higher-level visual and cognitive representations and their top-down expectations act, they do not always have effects that are equivalent to bottom-up activation by visual scenes. For example, one would not expect an imagery percept of a Necker cube to be bi-stable if the top-down expectation is already biased to one interpretation.

In summary, the “imagery debate” is often carried out as a thing unto itself, without engaging the greater theoretical and modeling literature in vision. Recent neural models clarify why imagery exists, as well as some of its mechanistic substrates as part of a larger theory of vision and cognition.

ACKNOWLEDGMENTS

Author was supported in part by the Air Force Office of Scientific Research (F49620-01-1-0397), the Defense Advanced Research Projects Agency and the Office of Naval Research (ONR N00014-95-1-0409), and the Office of Naval Research (ONR N00014-95-1-0624).

Problems with a “cortical screen” for visual imagery

David Ingle

39 Pratt St., Framingham, MA 01702. lake@gis.net

Abstract: I support Pylyshyn's skepticism that visual imagery reflects a reactivation of the spatial layout of active neurons embedded within a topographical cortical map of visual space. The pickup of visual information via successive eye movements presents one problem and the two visual systems model poses another difficulty.

One of Pylyshyn's most telling critiques of Kosslyn's “screen projection” theory of visual imagery deserves more amplification than he himself offers. He notes that, as we pick up information from a visual scene during a series of eye saccades, disparate spatial locations are sequentially projected on the same foveal representation within the striate cortex. If mental imagery (MI) of a remembered object involved reactivation of the striate neurons, which were active during the scan sequence, we would witness a jumble of MIs. My own experiences with unusual “visual persistence” (VP) effects makes this point clear. I can fixate on an object for as little as one second, close my eyes and continue to see a pencil, a face, or a word for at least ten seconds as a vivid positive after-image. (My VPs may result from use of a morphine-agonist for chronic pain, but they are confirmed by a second visual scientist whose images are most intense during hypoglycemia.) Unlike retinal after-images, our persisting images do not move as we turn eyes or head to the side – rather, they remain fixed in body-centered coordinates. After I have someone set out three small objects (unknown to me) in a triangular array on a white background and then briefly inspect each one in turn through a cardboard tube, I see for about ten seconds after eye closure persisting images of the three objects in their real world locations. That is, I see a spatial configuration that was never present on the retina during the three brief visual fixations. (A detailed account of several VP phenomena is in preparation.)

If my VPs reflect continued activity of shape-sensitive and color-sensitive neurons within a “cortical screen,” its location must be more central than areas 17 and 18, since these neurons have receptive fields (RFs) that move with saccades (e.g., in monkeys). Although objects and colors are sharply discriminated by many neurons within the inferotemporal cortex (Desimone et al. 1984), here there does not appear to be a two-dimensional map, but only cells with very large RFs which register the identity but not the location of their optimal stimuli. The parietal cortex is one candidate for the spatial screen on which my VPs are localized, since the receptive fields of some parietal cells (in the monkey) also maintain

a constant spatial location during eye and head movements (Anderson et al. 1990).

However, the map recently found within the parietal cortex (Sereno et al. 2001) is not based on the usual correspondence of visual RF locations with neuronal positions within the two-dimensional cortical sheet. Here, it is the optimal eye movement directions associated with cell discharge that are mapped in respect to cortical coordinates. If the parietal cortex provides a map of intended eye movements, then it must be actively searching for a target for its intrinsic map to come alive. The problem is that prominent visual objects, VPs and mental images (MIs) are well localized when the gaze is fixed or when a scene is presented tachistoscopically.

Another problem with the screen hypothesis is that my VPs remain localized in body-centered space, after I rotate my body more than 90 degrees. Then the VP is located outside my visual field. I can also obtain vivid MIs of my hands (opening, closing, or rotating) when each hand is held in the far periphery. These MIs are not altered by turning my head to one side, so that one MI is outside my visual field. While there is no current physiological evidence for a brain map which represents the rear (unseen) field, my unpublished studies have shown that one can localize an MI of a previously seen target, after rotating away from it, so that it is located behind the back. A group of 30 college students localized this rear field MI as accurately as they localized remembered targets in the frontal field (replication earlier findings of Attneave et al. 1977). Yet it is possible that some part of the parietal cortex participates in representation of body-centered spatial coordinates (including the space behind the head) since I reported (Ingle 1990) on a single individual whose small parietal damage is associated with total loss of short-term memory for contralateral target directions after self-rotation. Relevant spatial coding in the parietal cortex might be found if recording experiments were done in passively rotated monkeys who were trained to reach in the dark for recently seen targets.

A deep problem with the screen hypothesis is the very nature of the “two visual systems” model, as updated by Ungerleider and Mishkin (1982) and by Milner and Goodale (1995). The coding of object identity (via shape and color-selective neurons) occurs within a separate visual stream (ventral) than that encoding visual location (dorsal stream). If temporal and parietal reverberations are the direct sources of mental imagery (or of my VPs), the objects and their spatial frameworks must merge within some further brain system to which both dorsal and ventral stream targets project. One candidate for such a perceptual synthesis is the prefrontal cortex, where some neurons are found to encode both object identity and spatial location (Rao et al. 1999). So far there is no evidence for a spatially organized “screen” in this region where the two systems appear to merge. It remains possible that the registration of a (nonidentifiable) object moving through space does depend on shifting activity within a parietal map, but it appears that the confluence of space and identity occurs in a system no longer operating in maplike coordinates. I don’t wish to exaggerate my difference with Kosslyn, since I favor his hypothesis that visual imagery uses much of the same neural machinery as activated in direct perception of an object or scene, but I favor the idea of a brain system which no longer operates in a “map-like” fashion.

Ultimately, a central map of visual space must be read out in terms of motor coordinates: that is, perception of a given translation or rotation of an object in space has a precise equivalent in terms of body, head or hand movement required to produce the identical change. One result of my own VP experiments is that one or two persisting images can be moved about (after eye closure) as my hands move while holding them. A key experiment involves fixating on a small object (bottle or toothbrush), closing the eyes, and then moving it to the periphery in one hand, so that the VP is now vividly seen far to one side. If the second hand then moves from my lap to the symmetrical peripheral field, the single VP suddenly appears within the second hand, as if it had instantly jumped across 150 degrees of empty space. Somehow the brain activity

underlying the VP (in both the temporal lobes?) is co-opted by activating the motor (or premotor) system of the second hemisphere.

These facts suggest that my VPs are unusually vivid correlates of a process preceding the memory formation which underlies mental imagery. The strong linkages of VP location to hand movements suggest that MIs can be useful in representing peripheral objects to be grasped or manipulated while the subject is focused centrally. MIs may be used in guiding routine manual tasks, as well as used for creative solutions to new challenges.

How do we define “sameness” of the processing of mental images and general reasoning processes?

Margaret Jean Intons-Peterson

Department of Psychology, Indiana University, Bloomington, IN 47405.
intons@indiana.edu

Abstract: This commentary raises questions about the central concepts in the null hypothesis presented by the author of the target article and urges expansion of the treatment of mental imagery to forms of sensory imagery beyond the visual.

Professor Pylyshyn provides a stimulating and provocative review of research on visual imagery, which will serve as an incentive to expand exploration of this tantalizing and challenging aspect of human cognition. In recognition of his contribution I raise three issues that I hope will further such an expansion. The issues involve what Pylyshyn calls his “null hypothesis;” namely,

that reasoning with mental images involves the same form of representation and the same processes as that of reasoning in general, except that the content or subject matter of thoughts experienced as images includes information about how things would look. (Abstract, last sentence)

My first question deals with the definition of the word “same” and the second asks how the similarity of processing visual images and general reasoning is to be assessed when the last phrase allows for the modification of visual image processing by information about “how things look.” The third issue addresses the use of the term “mental” images when the article focuses exclusively on visual images.

How, exactly, is the word “same” to be defined? What are its philosophical concomitants? Excluding the final phrase of the null hypothesis for the moment, does Pylyshyn mean that the form of representation and the processes invoked by mental images must be identical to those invoked during reasoning in general, to satisfy his null hypothesis? Or does he assume a more fluid interpretation of “sameness,” such that any similarity between processes stimulated by imagery and general reasoning is sufficient to accept his null hypothesis? As any one who has followed studies of concept identification knows, the word “same” may refer to a range of possibilities, extending from exact, unequivocal, duplicative identity to similarity of various degrees. Hence, the definition of “sameness” must be specified, for it is not simply a pedantic concern. The concept of identity appears to exclude any deviations from precise replication, save, perhaps, those resulting from imperfect measurements and uncontrolled factors, including information about how things look, if we include the last phrase of the null hypothesis. If Pylyshyn equates identity with sameness, he is adopting a form of the null hypothesis so strong as to be almost meaningless, given our inability to control all relevant factors. If he means anything other than a cloning-type identity, he is adopting a form of the null hypothesis potentially so weak as to eliminate the possibility of refutation, particularly if information about visual appearance is allowed. It would be useful to know how Pylyshyn deals with these conundrums.

My third concern is that the title, the statement of his null hypothesis, and the general thesis of the article will encourage his readers to apply his reasoning to all mental images, even though his examples deal exclusively with visual ones. In fact, I regret that he narrowed his coverage to visual images because the inclusion of other forms of images, such as those often labeled as auditory, olfactory, gustatory, kinesthetic, and so forth, have the potential to provide important tests of his general discussion and even the expression of the null hypothesis by substituting appropriate final words such as “hear,” “smell,” “taste,” and “feel” for “look.” Far less research has been devoted to these types of imagery, to be sure, but there could be striking support for his central hypothesis if substantially similar (identical?) types of processing occurred in each of these sensory-proprioceptive areas, when the simulated sensory inputs were taken into account. Perhaps the next stage of imagery research is to explore the similarities and differences among various types of imagery.

Imagery in multi-modal object learning

Martin Jüttner^a and Ingo Rentschler^b

^aNeuroscience Research Institute, Aston University, Birmingham, B13 9DH United Kingdom; ^bInstitut für Medizinische Psychologie, Universität München, 80336 Munich, Germany. m.juttner@aston.ac.uk
ingo@imp.med.uni-muenchen.de

Abstract: Spatial objects may not only be perceived visually but also by touch. We report recent experiments investigating to what extent prior object knowledge acquired in either the haptic or visual sensory modality transfers to a subsequent visual learning task. Results indicate that even mental object representations learnt in one sensory modality may attain a multi-modal quality. These findings seem incompatible with picture-based reasoning schemas but leave open the possibility of modality-specific reasoning mechanisms.

In his target article, Pylyshyn focuses a lot on the depictive nature of mental images, and on the assumption that examining such images involves the same mechanisms as those used in visual perception. Linking imagery and vision in such a way implies that picture-like representations are explanatory for vision in its own right. Notwithstanding the functional role of early visual mechanisms, one should keep in mind that vision ultimately is not about perceiving flat, two-dimensional (2D) pictures. Rather, it is about sensing a three-dimensional (3D) space and 3D objects embedded within that space. In this respect, the discussion about pictorial versus non-pictorial formats prevailing in the imagery debate is reminiscent of the discussion about the nature of mental object representations.

The latter has been dominated by two opposing views. On the one hand, it has been postulated that objects are mentally represented in terms of “symbolic” 3D, object centred, part-based descriptions (e.g., Biederman 1987; 2000; Marr & Nishihara 1978). On the other hand, there have been studies providing evidence for a more “pictorial” 2D representation of 3D objects, in terms of multiple, viewer centred views, among which the visual system interpolates if necessary (e.g., Edelman 1995; Poggio & Edelman 1990; Tarr et al. 1998).

The two hypotheses have been mainly put to test in mental rotation paradigms in which a test object is presented from different perspectives and the changes of error rate or response latency in identification tasks are being measured as a function of viewing angle. However, it has been shown that within this paradigm the two alternative explanations may not lead to readily distinguishable predictions. First, the dependency on viewpoint is itself dependent on object familiarity (Tarr & Pinker 1989), demonstrating the necessity to take into account learning processes. Second, a closer inspection of the apparent complementary approaches shows that the assumed viewpoint-invariance of the 3D symbolic

descriptions only holds under certain conditions (Biederman & Gerhardstein 1995). Conversely, representations in terms of multiple 2D views may become quasi-independent from viewpoint, if the number of views is sufficiently large, or if the interpolation mechanism between views becomes more efficient due to training (Edelman & Bühlhoff 1992). Not surprisingly, the interpretation of such mental rotation experiments has remained controversial (see, e.g., Rentschler et al. 2000).

Yet spatial objects may not only be perceived visually but also by touch. Experiments on haptic object recognition demonstrate that the identity of familiar objects may be established very quickly and seems to be mediated mainly by 3D structural information (Klatzky & Lederman 1995; 1999). Indeed, there is an intrinsic similarity between visual and tactile object recognition, in that both are based on the extraction of basic features, such as contours and their spatial arrangement, which together define an object. This raises the possibility that object recognition may benefit from a multi-modal integration of sensory information.

In a recent study (Rentschler et al., submitted) of *trans-modal* object learning we investigated whether prior object knowledge acquired in either the *haptic* or *visual* sensory modality transfers to a subsequent *visual* learning task. Three molecule-like models, each being composed of four spheres, served as learning objects. The objects were generated both as virtual models (to be displayed and manipulated via the computer mouse in a virtual-reality environment) and as physical models. The experiment consisted of two phases, an exploratory phase and a learning phase. During the exploratory phase the subjects explored the objects either haptically (being blindfolded and using the physical object models) or visually (using the virtual object models). In the subsequent phase of visual learning the subjects were trained in a supervised-learning paradigm to recognize a set of 2D views of the learning objects. Both the duration of the exploration phase and the amount of training given in the learning phase were the same for all subjects.

Three groups of children and adolescents in the age ranges 8–9 years, 10–11 years, 13–14 years, plus a fourth group of adult subjects (> 20 years) participated. Each age group was subdivided into three subgroups which were assigned to one of the conditions: *haptic* (haptic exploration + visual learning), *visual* (visual exploration + visual learning) and *control* (visual learning only). Figure 1 summarizes visual recognition performance upon completion of the learning phase. The data show that the sensory modality employed during the exploratory phase had a distinct impact on the




Figure 1 (Jüttner & Rentschler). Visual recognition performance as a function of subject age upon completion of a multi-modal object learning task. The task involved combinations of either a visual or a haptic exploration phase and a visual learning phase. The control condition only involved visual learning.

subsequent visual learning. Moreover, this effect revealed a significant interaction with age. Specifically, for subjects aged 13 years and above learning performance was significantly higher if preceded by a haptic rather than a visual exploration. For children aged 10–11 years visual and haptic exploration proved equally efficient, whereas for younger children (8–9 years) only a visual exploration gave them a significant advantage relative to the control condition.

The implications of these results are threefold. First, the transfer from haptic to visual learning suggests that mental object representations may attain an intrinsically multi-modal quality even if the training involves only one specific sensory modality at a time. Second, the fact that the haptic-to-visual transfer proved to be more efficient than the visual-to-visual transfer shows that such representations are not compatible with any notion of a pictorial format. Otherwise one should have expected the reverse finding. Third, the age-dependent dissociation between the *haptic* and the *visual* condition argues against the notion that the former simply fosters the involvement of imagery mechanisms (such as mental rotation) more strongly than the latter, thus leading to a better learning performance. Rather, the significant interaction between the factors of *age* and *condition* indicates that haptic and visual exploration contribute independently to the ontogenesis of mental object representations.

In summary, our results suggest that any attempt to characterize mental object representations in terms of spatial images must be futile, because of the intrinsically multi-modal nature of such representations. However, they also provide some evidence that Pylyshyn's "null hypothesis," according to which reasoning is a unitary (i.e., modality-unspecific) mechanism, may require revision.

Mental imagery doesn't work like that

Stephen M. Kosslyn, William L. Thompson,
and Giorgio Ganis

Department of Psychology, Harvard University, Cambridge, MA 02138.
smk@wjh.harvard.edu www.wjh.harvard.edu/~kwn/

Abstract: This commentary focuses on four major points: (1) "Tacit knowledge" is not a complete explanation for imagery phenomena, if it is an explanation at all. (2) Similarities and dissimilarities between imagery and perception are entirely consistent with the depictive view. (3) Knowledge about the brain is crucial for settling the debate. (4) It is not clear what sort of theory Pylyshyn advocates.

Pylyshyn has done a service by assembling in one place his arguments against depictive theories of imagery. Although his basic points are essentially the same as those he presented in Pylyshyn (1981), our responses have been augmented by recent developments in cognitive neuroscience. In this brief reply, we focus only on what we take to be the most central points of Pylyshyn's arguments. But first, let us try to be clear on what we take to be the central issue: Does visual mental imagery rely (in part) on a distinct type of representation, namely, one that depicts rather than describes? By "depict" we mean that each portion of the representation is a representation of a portion of the object such that the distances among portions of the representation correspond to the distances among the corresponding portions of the object (as seen from a specific point of view; see Kosslyn 1994, for a more complete characterization). The issue is not whether images perfectly preserve perceptual phenomena; they obviously do not. Nor is the issue whether imagery shares processing mechanisms with perception; it does, but in theory those shared mechanisms need not rely on depictive representations. Nor is the issue whether knowledge (tacit or explicit) can affect imagery; it clearly can (indeed, many of the uses of imagery rely on this characteristic). The issue is whether imagery relies in part on a qualitatively distinct type of internal representation, which is not used in language.

Tacit knowledge. We have five comments regarding the appeal to tacit knowledge as an explanation for empirical findings about imagery.

1. The fact that tacit knowledge *might* explain a result does not imply that it *does* explain a result. Unless Pylyshyn can devise ways to discover whether subjects do in fact have specific tacit knowledge (and can use it to affect their behavior appropriately), then the claim is vacuous. Indeed, actual empirical studies have shown that subjects do not necessarily know what "actually perceiving an object" in specific circumstances would be like (Denis & Carfantan 1985). Given that some of the phenomena that occur in imagery, such as the "oblique effect" (i.e., the fact that oblique line gratings are less resolved than vertical or horizontal ones), can be demonstrated only in the laboratory, the burden of proof is on Pylyshyn to show that people somehow know about such properties of their visual systems and can use this knowledge to produce behavior like that which would have occurred in the corresponding perceptual situation.

2. Pylyshyn does not offer a theory of mechanism. Tacit knowledge could in fact be represented by depictive representations. The tacit knowledge theory is a theory of content, not of format – but the issue at hand is about format, whether mental images rely on a distinct type of internal code. We applaud Pylyshyn for proposing alternative interpretations for specific results (such as, attentional crowding as an account for mental image scanning results, e.g., Finke & Pinker 1982); these accounts can be tested. But we need a specific theory of imagery, which can then be subjected to empirical test. Pylyshyn's criticisms of the depictive theory are ad hoc patches where various processes are invoked here and there to explain empirical results. No cohesive whole is presented. The depictive theory, in contrast, presents a coherent, internally consistent view of how mental images may be processed. Such a view is more parsimonious than the patchwork that Pylyshyn assembles to account for the various findings he discusses.

3. The depictive view has predicted many empirical results that are now "explained" post hoc by the tacit knowledge view. Indeed, given the fact that we apparently cannot know in advance exactly what tacit knowledge people in fact do have (e.g., according to Pylyshyn's logic, they have tacit knowledge of the oblique effect but not color mixing), the tacit knowledge view provides no firm grounds for making empirical predictions.

4. The issue of cognitive penetration is a red herring: Of course, knowledge can influence imagery – but this does not imply that all properties of imagery are a result of one's knowledge. The claim that people cannot visualize what they have never seen is inconsistent with a large literature on the role of imagery in creativity, where people clearly imagine novel shapes (e.g., Finke et al. 1992).

5. One reason the field has embraced cognitive neuroscience approaches is that tacit knowledge per se cannot selectively modulate particular neural mechanisms. For example, the fact that most studies of imagery find activation of topographically organized areas – which truly depict information – cannot be easily dismissed (more on this below).

Imagery and perception. Virtually all theories of imagery claim that imagery and perception share common mechanisms. The issue focuses on the nature of those mechanisms. We have the following observations.

1. Pylyshyn claims that "it ought to be an embarrassment to picture-theories" that they postulate so many similarities between the mind's eye and our own eyes. For example, citing Kosslyn (1978) Pylyshyn writes: "it seems that the mind's eye has a visual angle like that of a real eye . . . , and that it has a field of resolution which is roughly the same as our eyes . . ." and "[i]t even appears that the 'mind's eye' exhibits the 'oblique effect' . . ." (sect. 7.3, para. 2). Pylyshyn seems to confuse theoretical claims with empirical findings; these characteristics have been empirically demonstrated – which hardly seems a reason for embarrassment. Like it or not, that's the way the studies came out. The depictive theories made such predictions, which were successful; we infer that

Pylyshyn would not have made such predictions, based on his view.

2. Pylyshyn has greatly oversimplified the results of the many studies that have now addressed ways in which objects in images can be reinterpreted. The bottom line is that people can in fact reinterpret images if they are given ways to cope with the limited working memory capacities. Moreover, people can “break up” the elementary shapes organized by early vision to “see” new patterns in images (e.g., see Rouw et al. 1997; 1998).

3. Pylyshyn is correct: Images do not preserve the earliest visual representations. However, to our knowledge, no depictive theorist has ever claimed that they do. Rather than being like “primal sketches” (in Marr’s terminology), they are like 2.5-D sketches; they incorporate organized units.

4. The relation between the conscious experience of imagery and the underlying representation is not necessarily simple or straightforward. Pylyshyn asks how a patient could have cortical blindness and still experience imagery. If the experience of imagery arises from structures that receive inputs from neural structures that implement depictive representations, then the answer is straightforward: the experience occurs when visual memory representations are activated (in the inferior temporal lobe), which can occur via top-down inputs (e.g., from frontal cortex) even when the early visual structures are disrupted. However, if the depictive representations cannot be formed in early visual cortex, then the patient should not be able to reorganize the pattern in the image. To our knowledge, this critical study has not been conducted – and thus Pylyshyn should be careful when asserting categorically that cortically blind people have perfectly normal imagery abilities.

5. As Pylyshyn states, many low-level visual phenomena are not present in imagery. However, all of these phenomena are stimulus-driven, by input from the eyes. In contrast, visual mental images are not driven by input from the eyes, and thus cannot be expected to mimic effects of such bottom-up processing. Color mixing begins to occur at the retina, for example, and thus imagery cannot be used to anticipate the results of novel mixtures. Motor tracking is another example of bottom-up processing, which we would not expect imagery to mimic.

6. Pylyshyn is right in saying that we do not have a panoramic display in imagery or vision, but this is irrelevant to the issue at hand. Our images depict only a small slice of the visual world, but that does not imply that they are not depicting this portion.

7. Pylyshyn claims that many of the imagery results can be explained in terms of attention per se. Our responses are as follows: (a) Why not assume that imagery is what allows some of the corresponding attentional phenomena to occur? It is not clear, for example, whether subjects are using imagery as a way to attend. What, exactly, is meant by “attentional”? (b) Pylyshyn claims that most imagery studies require subjects to “imagine something while looking at a scene;” thus “. . . superimposing or projecting an image onto the perceived world” (sect. 5.3, para. 1); this claim is simply false. Most imagery studies have no such requirements. Moreover, comparable results are obtained when imagery experiments are performed with eyes closed and eyes open (e.g., in mental scanning), which makes the “spatial indexing” idea untenable as a general explanation of imagery phenomena. (c) It is not clear how attention per se can explain the representation of shape in imagery, as is required for many imagery tasks (such as those requiring verifying subtle properties of objects when one’s eyes are closed).

The importance of the brain. Pylyshyn dismisses neuroscientific data on numerous grounds. We, in contrast, feel that such facts are crucial for the current debate, for the following reasons.

1. Anderson (1978) proved that any theory that posits a representation (e.g., a depictive image) with processes that operate on it can be mimicked by another theory that posits a different representation (e.g., a list) with a different set of processes. In the alternate theory, the changes in the representation are compensated for by corresponding changes in the process. Thus, Pylyshyn is

correct in noting that the mental scanning results do not necessarily implicate a depictive representation; they can of course be explained in other ways (e.g., see Kosslyn & Pomerantz 1977). Pylyshyn claims that adding constraints to a system to ensure that patterns in an array are interpreted as depictive representations is ad hoc, as ad hoc as making up a theory post hoc to explain the effects of distance on mental image scanning. However, Anderson also pointed out that neurophysiological data could constrain theorizing, preventing the theorist from playing fast-and-loose with the characterization of the representation and process. One major advantage of shifting the theory from a computer metaphor to one rooted in the brain (compare Kosslyn 1980, with Kosslyn 1994) is that we cannot simply make up properties of representations and processes as we see fit. In the brain, the projections to higher areas from early topographically organized areas (which truly depict information) preserve the topography, with the receptive fields becoming increasingly large as one goes deeper into the system. The actual physical wiring is designed to “read” the depictive aspects of the representation in early visual cortex. In so doing, the interpretive function is not arbitrary; it is tailor-made for the representation, which is depictive. Form not only follows function, but in some cases function follows form! Neuroanatomy is in fact relevant for cognitive theories.

2. Pylyshyn implies that modern depictive theories posit that a picture is projected onto visual cortex when we entertain mental images (sect. 7.2), and that “the visual system is involved in imagery and that it examines a pictorial display” (sect. 7.3, para. 2). Contemporary depictive theories assume that the medium that supports depictive representations (early visual cortex, for neurologically oriented theories) sends signals to areas that store visual memories; this input is interpreted by matching to the stored memories (and is also sent to areas that interpret spatial relations). In no case is there a “mind’s eye” that is “looking” at something. There is no need of a homunculus for interpreting patterns in images, any more than there is a need for a homunculus in visual perception. This parody does not serve to further the debate.

3. Pylyshyn apparently misunderstands how patterns of activation in topographically mapped cortex represent information. The cortical magnification factor indicates the number of mm of cortex that are devoted to processing 1 degree on the retina at a given eccentricity. The cortical magnification factor decreases as stimuli move toward the periphery; the rate of decrease is maximal at the fovea, where, in humans, each degree of visual angle is allocated over 2 cm of striate cortex (Area 17). However, for the present issue it is irrelevant whether the cortical representation of a “large” peripheral pattern looks larger or smaller than that of a “small” foveal pattern to an outside observer: The crucial point is that in the context of the processing system (the connections to other brain areas, in this case) 1 cm of cortex in rostral Area 17 represents a larger spatial extent than 1 cm of caudal Area 17. From the “point of view” of the higher-level visual areas that “read” these representations – which is the only point of view that matters – the more anterior parts of Area 17 represent increasingly larger swaths of space. This is no different from the long-resolved issue as to why we don’t see the world upside down, given that the retinal image is inverted!

4. Pylyshyn is simply incorrect when he states that most studies of imagery have found activity only in visual association areas, not topographically organized regions of cortex. In fact, of the 21 fMRI studies of imagery we are aware of, 15 reported activation in Areas 17 or 18 (both of which are topographically organized, and hence implement depictive representations). Thompson and Kosslyn (2000) provide a review, which is updated in Kosslyn and Thompson (2002, under review).

5. In discussing topographically organized areas, Pylyshyn seems unaware of neuroimaging findings that have demonstrated clear retinotopic organization of the human visual cortex (similar to that found in other primates; see DeYoe et al. 1994; Hasnain et al. 1998; Sereno et al. 1995; Van Essen et al. 2001). Furthermore, he claims that no similar evidence to the Tootell et al. (1982) data

have been produced for imagery. We cannot fault Pylyshyn for not being aware of the very recent results of Klein et al. (submitted); nevertheless, these findings are merely the most recent in a steady stream of neuroimaging findings that point in the same direction. In this study, subjects visualized checkered bow-tie like patterns that were either oriented vertically or horizontally. Event-related fMRI was used to monitor activation in Area 17 while subjects performed the task. When the subjects visualized the shape vertically, the pattern of activation neatly fell along the vertical meridian of Area 17; when they visualized it horizontally, the pattern neatly fell along the horizontal meridian! In another condition, subjects saw the figure in the two orientations – and the results were virtually identical to what was found in imagery. This is very strong evidence that the topographic properties of Area 17 are in fact activated during imagery.

6. Pylyshyn wishes to dismiss the findings of Kosslyn et al. (1999), who used transcranial magnetic stimulation (TMS) to impair processing in medial occipital cortex – and showed that this effect in turn disrupted both imagery and perception (to the same extent). But this is merely one of several studies, which nicely converge in implicating early visual cortex in visual imagery. For example, Sparing et al. (2002) showed that visual mental imagery increases the excitability of medial occipital cortex, as indicated by the fact that TMS to this region evoked more phosphenes following imagery; in contrast, a control auditory task did not have this effect. Moreover, Kosslyn et al. (1996) statistically removed the effects of variations in blood flow in all other brain areas and still found that the degree of activation in Area 17 per se predicted response times in an imagery task. Thus, the activation in Area 17 cannot be written off as an incidental by-product of activation elsewhere in the brain. Kosslyn et al. (2001) provide a recent review of the relevant literature.

7. Without question, topographically organized cortical areas support depictive representations that are used in visual perception. These areas are not simply physically topographically organized, they function to depict information. For example, scotomas – blind spots – arise following damage to topographically organized visual cortex; damage to nearby regions of cortex results in blind spots that are nearby in the visual field. Moreover, transcranial magnetic stimulation of nearby occipital cortical sites produces phosphenes or scotomas localized at nearby locations in the visual field. These facts testify that topographically organized areas do play a key role in vision, and that they functionally depict information.

8. Pylyshyn asks how color and texture could be integrated into a depictive image unless the image is literally depictive. However, even if an image of a green object were in fact literally green on the cortex, this would accomplish nothing. The issue is how physical states of the brain are “read” by other parts of the system. Tye (1991) suggests that depictive representations may be annotated; properties such as color and texture may be represented elsewhere, with pointers to specific parts of the depictive representation. This sort of hybrid representation seems reasonable (especially given the fact that brain damage can result in dissociations among these properties), and preserves the distinctive geometric aspects of depictive representations.

9. Perhaps an argument *ad absurdum* can help us understand some of the implications of the claim that early visual cortex uses propositions to represent visual events. Assume for the moment that early visual cortex relies on propositions to represent incoming visual images, including any information it receives from the lateral geniculate nucleus (LGN). Why stop here? Isn't the LGN also using propositional representations? And why not go all the way, and conclude that the retina uses propositions to represent the images that fall onto it? We see no reason why the arguments offered in favor of propositional representation in early visual cortex would not also apply to earlier visual stages. Now turn the argument on its ear: If the retina employs depictive representation, why wouldn't the LGN, which is topographically organized? And if the LGN does, why would early visual cortex – also topo-

graphically organized – discard this information and use propositions?

Form of a theory. Pylyshyn suggests that our approach has been fruitless, and that a good theory would begin by laying out general constraints and boundary conditions. Such a theory would have the same form as theories of language. We have the following reactions.

1. Pylyshyn claims that compositionality guarantees that any image representation system must be topographically organized – and thus the notion that a representation can be “depictive” is meaningless. Not so. A genuine depictive representation gives rise to emergent properties that must be actively computed in other systems. For example, the depictive representation of the letter “A” makes explicit and accessible the triangular enclosed shape; given only a description of the segments and how they meet, considerable processing is required to derive the fact that such a shape exists. Certainly, the existence of the enclosed shape could be included in the description, but now consider spatial relations among portions of the shape: There are a near-infinite number of spatial relations that are evident in a depictive representation. For example, in a depictive image, it is just as obvious that the top of the A is above the center bar as it is that it's above the left bottom terminus, and so on. These relations are immediately accessible in the depictive representation. In contrast, every one of them would need either to be explicitly mentioned or computed as needed in a nondepictive representation.

2. Depictive representations are not composed of discrete symbols. A depictive representation can be divided up in any arbitrary way and the parts still represent a portion of the object or scene (e.g., think of cutting up a photograph randomly – the fragments will still be representations of portions of the object or scene); this is not true of a propositional representation.

3. Pylyshyn discusses imagery simply as another form of reasoning, ignoring the many other uses of imagery, such as in memory encoding and retrieval, mental practice, pain control, and emotion regulation. In contrast to the views Pylyshyn advocates, depictive theories have led to a large number of empirical discoveries about all these functions of imagery. It is not clear how tacit knowledge can account for the role of imagery in these other activities.

4. Finally, we are intrigued by Pylyshyn's ideas about how to formulate better theories, and urge him to get on with it. As soon as he has formulated an alternative theory, we will be delighted to conduct studies to evaluate the fruitfulness of its predictions against those of our theory.

Single cells in the visual system and images past

Glenn E. Meyer

Department of Psychology, Trinity University, San Antonio, TX 78212.
gmeyer@trinity.edu <http://www.trinity.edu/gmeyer>

Abstract: Various techniques have attempted to localize imagery. However, early findings using single cell recordings of human receptive fields during imagery tasks have had little impact. Reports by Marg and his coworkers (1968) found no evidence for imagery in human Area 17, 18, and 19. Single cells from humans suggest later imagery-related activity in hippocampus, amygdala, entorhinal cortex, and parahippocampal gyrus.

Myriad tests have tried to localize imagery. Most physiological attempts have not been directly at the cellular level (PET, EEG, fMRIs, etc.). However, there have been previous and more direct tests of early localization that are not well known.

One of the best tests of imagery loci might be recording from human visual cells during imagery. In fact, this has been done (Marg 1970;1973; Marg et al. 1968). Single cells, from Areas 17 (1 cell), 18 and 19 (4 cells), were recorded from humans. Receptive

fields were similar to those in monkeys. Then there was a test of imagery.

Marg et al. (1968, p. 350) state:

Attempts to have the patient mentally control unit rhythms which he could hear over a loudspeaker were fruitless. Similarly, attempts at mental imagery, even of the effective target, seen moments previously, did not noticeably influence the unit response. These cells seem to have no role in mental visual imagery.

And Marg (1970, p. 154) reports:

None of the units or their plotted receptive fields could be influenced by a patient's efforts to change them. For example, we increased the audio gain until the patient could hear the pulses of the unit firing in his cortex and then asked: "Can you do anything to influence it? Can you increase it or decrease it, or affect it in any way?" No matter how much the patient tried to influence the response, we could detect no changes. We also brought the target into the receptive field and asked: "Did you hear that sound when the target was brought here? Now, the target is withdrawn. Imagine it is there and try to make the same sound come from the loudspeaker." No one succeeded in doing that.

Wilson, et al. (1983) report no other receptive field mappings in this part of human visual cortex until their own. However, sample size was small. Ehrlichman and Barrett (1983), in reviewing EEG imagery studies, feel that you never know if subjects are really "imaging" when told to. Marg depended on subjective reports. If their data are believed, then the proposition of early visual system localization has failed, but this direct test is little known.

There are other possible sites outside of the classic visual system. Object related images were reported with stimulation of posterior hippocampus (Adams & Rutkin 1970; Halgren et al. 1978a; Horowitz 1970). Adams and Rutkin (1970) and Halgren et al. (1978b) report visual sensations such as flashing lights and colored balls with hippocampal stimulation. Cartoon-like or television-like reports can be produced by medial temporal lobe stimulation (Halgren et al. 1978a). These data suggest involvement of the hippocampus in imagery that is not surprising given its visual inputs, some receptive field organization, and roles in memory (Halgren et al. 1978a; Horowitz 1970; Wilson et al. 1983).

Does subject initiated imagery activate the hippocampal formation? Halgren et al. (1978b) tested this. Human subjects performed various visuo-spatial tests while hippocampal cells were recorded. The tasks were congruent with Farah's (1984, p. 250) criteria for imagery. Five cells fired strongly during recall of recent events. One fired only when the patient was asked to remember spatial aspects of his room. However, activation was not found for memory tasks involving color photographs with unfamiliar rooms. This suggests that the hippocampus plays some role in imagery. Hippocampal involvement in coding memory for places in the environment has been suggested in animals (Wilson et al. 1983). Interestingly, Parsons et al. (1987) report that the "noted amnesic HM" demonstrated no improvement in a mental rotation task after several days of training. They conclude that limbic structures are at least a necessary component for the improvement of a skill for mental rotation (p. 5). Data such as Halgren et al.'s or Parsons et al.'s are not supportive of early linkage propositions.

Extrastriate human visual physiology, imagery and its interface with the rest of cortical and subcortical processing is mainly explored by external scanning methods. However, there is one suggestive cellular level result. Kreiman et al. (2001) recorded from neurons in hippocampus, amygdala, entorhinal cortex, and parahippocampal gyrus while subjects were imagining previously viewed images. Stimuli were faces showing emotions, household objects, spatial layouts, cars, animals, drawings, famous faces, food, and complex patterns. For the neurons that fired selectively for both vision and imagery, the majority had identical selectivity. Firing rates during vision and imagery were highly correlated. These data argue for imagery loci that are far from early visual areas.

Testing human single cells isn't going to be frequent. How about other primates? Could a mental rotation curve be obtained while

a monkey is in the stereotaxic and recordings are made from an appropriate neuron involved in the imagery process? Perhaps a "linguistic" chimp could be "told" to image in part of the visual field when a neuron covering that spot is captured by the electrode. This would assume that rotation represents image processing, which is controversial as many propositional models suggest similar functions. Do monkeys image like we do? Pigeon "mental rotation" isn't similar to ours (Hollard & Delius 1982) with flat functions unlike sloped human ones. Bees (Collett & Kelber 1988) demonstrate spatial abilities and researchers refer to them as having images. Pigeon or bee imagery with their concomitant neurophysiology is not yet well understood.

Single cells may not resolve if images are pictorial or propositional. It does seem that later visual structures are very active in imagery. One must wonder though: if such cells are activated when we see things, why would they be doing something different when we image? Are we seeing storage of propositional structures during vision and retrieval during imagery? What brain chunk actually looks at an image? The homunculus rears its head. In any case, propositional models must deal with the use of imagery by the common folk. To quote the Everly brothers:

Whenever I want you, all I have to do is Drea-ea-ea-ea-eam, dream, dream, dream.

Personally, anytime, night or day, I have never tried to taste the lips of a propositional data structure.

Imagery and blindness

Susanna Millar

Department of Experimental Psychology, University of Oxford, Oxford, OX1 3UD, United Kingdom. susanna.millar@psy.ox.ac.uk
<http://www.psych.ox.ac.uk>

Abstract: My concerns are about the phrase "the nature of imagery," and the interpretation of findings with blind people. This discussion considers reports of imagery by congenitally totally blind people, and what should not be inferred from comparing efficiency levels of blind and sighted people in spatial tasks.

What can be meant by asking about the "nature" of mental imagery? The target article does not make it clear what the question is. I assume that it is not about the ontological status of mental imagery. The clarification of hitherto intractable philosophical body-mind (e.g., single, dual, or epiphenomenal) problems cannot be a precondition for empirical enquiry.

Images are clearly not explanations. But does any one disagree with Pylyshyn on that? The discrepancy in reports of "imageless thought" by the Würzburg school and of visual imagery by Titchener (1909) has been known for over a hundred years.

However, the same doubt applies to what Pylyshyn calls "any principles to which we have conscious intellectual access" (sect. 2, para. 1). It is not obvious that conscious access to principles is a good criterion for how people solve logical problems, or how they come to understand or to apply the principles of Euclidean geometry. Pylyshyn's "reason to be skeptical about what one's subjective experience reveals about the form of a mental image" (sect. 2, para. 4) must also apply to other subjective evaluations of how we solve logical problems.

The "nature" of reported imagery is not necessarily "visual." The fact that congenitally, totally blind people report imagery in other modalities is relevant to the empirical question of what, if anything, may be modality-specific about subjective experience.

Indirect empirical tests of subjective experiences have not been as futile as Pylyshyn suggests. It has been shown, for instance, that the accuracy of Eidetiker who report very vivid imagery, can be tested and suggests memory effects (e.g., Haber & Haber 1964). In principle, therefore, methods of testing modality-specific effects of inputs on memory are relevant.

Reported imagery and empirical tests for modalities other than vision. Pylshyn argues cogently for careful scrutiny of methods and conditions in interpreting empirical findings. But the same applies to studies with blind participants.

The general reader is unlikely to be aware that only a minute proportion of blind people is congenitally totally blind. Legally blind people differ widely in degrees of residual vision and in early vision. Not all questions about blindness require participants who totally lack visual experience from birth. However, findings based on averaging imagery responses from participants whose visual status can only be conjectured (e.g. Johnson 1980) are not evidence that vision is irrelevant, as suggested in the target article. Decisive relations between visual experience and responses to imagery questions for different modalities have been found (e.g., Schlaegel 1953).

Congenitally totally blind people report that they do experience imagery. The response of a colleague, asked to imagine his room at home, may be taken as typical. He mentioned feeling cold air on his face, and hearing some form of echoing sound on entering his room late at night, touching his desk after walking a certain distance to the right, and hearing a difference in the sound of his footfall and in sensations underfoot on reaching the rug before the fireplace.

A more extensive (as yet unpublished) study showed that younger blind participants also talked about echoing sounds, touching obstacles, the feel of the ground underfoot, and moving in certain directions when asked how they walked from one building to another. The absence of visual terms was striking, because the same blind people habitually use visual terms in other contexts (e.g., “yes, I see what you mean”).

Introspective reports may, of course, be descriptions of remembered or projected activities rather than of current images of cold air, sounds, or touch. There is no “litmus test.” But converging findings that the imagery we experience subjectively is not confined to a single modality contributes to the probability that modality-specific aspects of inputs affect recall.

An alternative method is to test modality-specific effects in memory by comparing effects of actual and imagined movements on recall (Finke 1979). We found that recall of a target movement by young congenitally totally blind children was biased only slightly less by merely imagining shorter or longer movements during a delay period, than by actually executing the biasing movements during delays (Millar & Ittyerah 1991). It is possible that the intention to execute a biasing movement mobilised kinaesthetic sensations, and retrieval of “tacit” knowledge. But such alternative interpretations of the findings also require precise empirical tests of what “tacit knowledge” was involved, and how it was retrieved while doing nothing overtly.

What can be inferred from studies with blind participants? There is good evidence that visual experience is not necessary for solving spatial problems, including mental spatial rotation, and good evidence also that such problems constitute a major difficulty for young congenitally totally blind children. The apparent paradox is explained by memory overload in difficult conditions, procedural knowledge, but, importantly, also by the availability and congruence of reference information from different sensory sources (Millar 1988; 1994).

The important point is that the level of efficiency at which a participant arrives, or whether it is higher, lower, or equal to the efficiency of another person or group, is not evidence about the question what strategies either participant used, let alone that they necessarily used the same heuristics. Relative effects, for instance, of using external and/or body-centred reference cues, can be tested empirically (e.g., Millar 1979; 1981; 1985; Millar & Al-attar 2000; 2001; 2002). But neither equal efficiency by blind and sighted participants, nor differences in efficiency tell us anything about how they solved a given spatial problem, let alone that they used the same strategy.

Pylshyn’s null hypothesis is a formal description of problem solving tasks. It makes no predictions about how people actually

go about solving different types of problems or how we might test the relation between the heuristics people use and what they say they experience. Our questions about imagery may well be misleading. But the contention that we are “deeply deceived by our subjective experience of mental imagery” (target article, sect. 1.1, para. 3) seems odd in the context of evolutionary biology.

Visual imagery and geometric enthymeme: The example of Euclid I.1

Keith K. Niall

Defence and Civil Institute of Environmental Medicine, Defence Research & Development Canada, Toronto, M3M 3B9, Canada.

Keith.Niall@drdc-rddc.gc.ca

Abstract: Students of geometry do not prove Euclid’s first theorem by examining an accompanying diagram, or by visualizing the construction of a figure. The original proof of Euclid’s first theorem is incomplete, and this gap in logic is undetected by visual imagination. While cognition involves truth values, vision does not: the notions of inference and proof are foreign to vision.

Seeing a thing is different than thinking about it. As links or relations between vision and visual imagination are brought to light, this distinction between vision and cognition can be obscured as a consequence. When one imagines seeing something, is that more like seeing or more like thinking? It may seem like a combination of the two. (It is crucial to ask why we want to say this.) Yet there is a specific way in which seeing something is different from thinking about it. That distinction does not admit half measures, just as Pylshyn’s test of cognitive impenetrability does not admit half measures. The distinction is simple. Thinking involves reputed entities which are truth-valuable, while vision does not. Depending on one’s account, such things that bear truth values have been called true propositions or false sentences or just thoughts.

There has been much talk about the metric properties of visual imagery, and its Euclidean or otherwise-geometric description. But it is helpful to remember that no one proves Euclid’s first theorem by vision. Where truth values do not apply, the notions of demonstration and inference have no foothold. We may use vision to avoid stout obstacles or to dodge leaping pussycats, but vision provides no demonstration of that very first theorem of Euclid. (Still, vision is intentional: that is a different matter.) Far from being Euclidean, vision doesn’t get on at all with Euclid. It isn’t that Euclid’s first proposition is too difficult: the first proposition comes before the *pons asinorum* of the fifth proposition, the one which separates able learners from dunces. Vision does not demonstrate any theorem of Euclid, rightly or wrongly. Nor is it that the first proposition involves microscopic quantities or horrendously complicated relations. Euclid’s first proposition does require a demonstration that two circles intersect. The centre of each circle lies on the circumference of the other. (If we can’t visualize that two such circles intersect, what can we visualize?) The problem is made clear by an old lesson in geometry.

Euclid’s first proposition is “on a given finite straight line to construct an equilateral triangle.” The line segment is AB . One circle has centre A and radius AB ; the other circle has centre B and radius BA . From the point C at which the circles intersect, the line segments CA and CB are joined to the line segment AB . We have already gone beyond Euclid’s definitions, postulates, and common notions. To speak strictly, there is already a gap in the proof, a gap unnoticed by the eye. There is nothing in Euclid’s definitions, postulates, and common notions that enables us to show the two circles intersect. Zeno of Sidon noticed this in the early first century B.C.E. (Heath 1956; it seems safe to assume Zeno did not notice it as a visual illusion, or anything of the sort). David Hilbert (1899/1999) brought it to attention once more, by his celebrated axiomatization of Euclidean geometry. His achievement has been inter-

preted as making geometric proof a thoroughly axiomatic business, rather than one of visualization or geometric intuition.

But for centuries this gap in the proof of Euclid's first proposition went unnoticed and unheeded. Perhaps a complete proof existed in the collective intuition of European civilization, but that seems a wild conjecture. If one draws chalk marks on a blackboard, of course the chalk circles will intersect. And when those chalk circles intersect, they are seen to intersect. It is as plain as day that two generously overlapping circles will intersect. Similarly, one can imagine the construction of the present Figure 1 (see below), and the circles will be imagined to intersect. That act of imagination is no proof the circles intersect, and in that sense it does not correspond to Euclid. It might be thought to be an appropriate extension to Euclid's geometry, or else it might be thought to represent what Euclid really meant. (That is, one can be tempted to revise history to suit one's own ideas.) Yet the gap in proof and its history tell us that imagination is not proof, just as perception is not cognition. Vision and visual imagination do not have the power of logical demonstration: centuries can pass, and a gap in logic will not be noticed by vision alone. Cognition does involve truth-valuable items, and involves notions like proof, inference, and demonstration.

Why should the gap in Euclid's proof, and the inapplicability of truth values to visual images, be important to the psychology of visual imagery? It is important because of a claim that is smuggled into discussion of visual imagery. That is, the claim that items of visual imagery constitute knowledge. Whether analogical or propositional, they are supposed to constitute knowledge. Pylyshyn (1984, p. 135) claims that "perception involves semantic-level principles – especially those of inference." In other words, the properties of mental imagery are represented to the mind without further effort or explanation. Such a story about visual imagery is a caricature of the development of knowledge, which can be slow, deliberate, and collective. Such a story "overlooks the need to give any account at all of the way the inner understander works, any account of the mechanics of inner representation, any account of what kind of reacting is comprehending." (Millikan 2000, p. 112).

The notion that properties of mental imagery constitute knowledge might be defended as part and parcel of some sort of Empiricism, perhaps British Empiricism. After all, a central tenet of Empiricism is that all knowledge proceeds from the senses. Knowledge may share the form of sensory information, whatever

such an assertion could mean. So it may be that this debate over mental imagery is a debate between one sort of Empiricist and another sort. Suppose that visual imagery is propositional, and that its propositions constitute some of our knowledge of geometry. Then either it is not Euclidean in the sense of being capable of proof, or else its propositions have been insufficient to demonstrate Euclid I.1.

Motion, space, and mental imagery

Romi Nijhawan and Beena Khurana

*School of Cognitive and Computing Sciences, University of Sussex, Brighton, BN1 9QH, England. romin@cogs.susx.ac.uk/beenak@cogs.susx.ac.uk
http://www.cogs.susx.ac.uk/users/romin/index.html
http://www.cogs.susx.ac.uk/users/beenak/index.html*

Abstract: In the imagery debate, a key question concerns the inherent spatial nature of mental images. What do we mean by spatial representation? We explore a new idea that suggests that motion is instrumental in the coding of visual space. How is the imagery debate informed by the representation of space being determined by visual motion?

The representation of space is critical to Pylyshyn's arguments against the notion that mental imagery is "inherently spatial" in nature. He supports his position by drawing a distinction between "intrinsic" properties of a mental representation and those that are transitory and attributable to people's beliefs. However, what does it mean for a representation to be "inherently spatial" in nature? We would like to take a new tack on the problem by considering a previously unexplored notion of how visual space comes about, assuming it is not present at birth in the sense we associate with it. We shall argue (see below) that spatial maps in the brain are trained during development by neural activity due primarily to movement in the world. This activity sets up spatial maps that are then referenced for the visual perception of both moving and stationary stimuli in the adult. Thus, the question to be posed becomes one of whether mental images have access to spatial maps (for example topographic representations). Perhaps, given Pylyshyn's litmus test of cognitive impenetrability, the question of consequence is whether images have obligatory access to spatial maps.

From the point of view of ontogeny and phylogeny a visual system that primarily processes movement is more primitive. The extreme periphery of the retina, which may be considered an older system, responds only to movement. Indeed, according to Richard Gregory (1979): "it seems that it is only the eyes of the highest animals which can signal anything to the brain in the absence of movement." However, developmentally, a primate's visual system during infancy may behave like a primitive visual system capable of responding only to motion (ontogeny recapitulates phylogeny). There are at least two important reasons why this may be so. First, movement (or change) is a much more potent stimulus for the visual system than stationarity (or continuity), particularly for the immature visual system which does not respond strongly even to high contrast stimuli unless motion is introduced (Hubel & Wiesel 1963). Second, given resource limitation and the association of movement with danger or food, detection of motion is of primary importance. Although an adult observer may hold a belief that an everyday visual scene (such as a parking lot) contains stationary objects relative to which some other objects are moving, we claim that motion is what sows the seeds for the coding of visual space. (By visual space we refer to the thing that extends between visual objects, and not to Newtonian space analogous to Ether). At a later stage in development, as the system matures and stationary objects become more effective, such objects are spatially localized according to the same principles that were first established by moving objects. Thus, for example, the rule that point *p* of a retinotopic map is associated with position *p'* of visual space, first established on the basis of neural activity due to motion, is then generalized




Figure 1 (Niall).

for stationary objects such that stimulation of point p by a stationary object also yields a position p' in visual space.

These considerations suggest that, while the representation of space called upon by moving objects is without a doubt “depictive,” stationary objects and mental images of stationary objects may not warrant the same conclusion. The inherently spatial character of motion is also supported by entities such as a Reichardt detector, with two spatially displaced input signals. Perhaps a spatial representation accessed through static images is weak, and only seems strong; just as we seem to have a detailed representation of a scene despite this representation being very limited, as change blindness experiments have convincingly shown. It may also be the case that the access to spatial representations by mental images is not obligatory; however, when achieved, the resulting performance of the observer is akin to that measured for visual objects. Thus, it may not be an all-or-nothing answer in terms of the spatial character of imagery, but rather that there is a hierarchy of accessibility to spatial maps, with moving items being the quintessential stimuli, followed by stationary stimuli, with static images holding the least accessibility rights. In this sense mental images are least like representations triggered by motion.

In sum, the crux of deciding whether a percept or a mental image is spatial in nature is whether it can and does access spatial maps. Given the above conjecture it may be that moving objects are privileged in that they have obligatory access to the spatial maps because they were the stimuli that gave rise to them. Stationary objects over the time course of development gain access to the same spatial representations set up through object in motion. Mental images of stationary objects may be particularly weak in terms of their ability to access spatial representations. Intriguingly, though mental images of moving objects may have obligatory access to spatial maps, it is this very form of mental imagery that is not available to observers, in that observers are unable to mimic or imagine smooth motion. Were that achievable, through a clever but at this juncture unspecified experiment, would Pylshyn accord spatial attributes, albeit limited, to mental images?

Is mental imagery prominently visual?

Marta Olivetti Belardinelli and Rosalia Di Matteo

*ECONA (Interuniversity Center for Research on Cognitive Processing in Natural and Artificial Systems) and Department of Psychology, University of Rome “La Sapienza”, I-00185 Rome, Italy. marta.olivetti@uniroma1.it
rosalia.dimatteo@uniroma1.it*

Abstract: Neuroimaging and psychophysiological techniques have proved to be useful in comprehending the extent to which the visual modality is pervasive in mental imagery, and in comprehending the specificity of images generated through other sensory modalities. Although further research is needed to understand the nature of mental images, data attained by means of these techniques suggest that mental imagery requires at least two distinct processing components.

The pictorial account of mental imagery rested on the demonstration that mental image processing follows the same rules that perceptual processing follows. However, this demonstration was performed quite exclusively with experimental investigation focusing on visual imagery. In particular, strong support for the analogical theory derived from Kosslyn’s demonstration that visual images have a spatial extent (Kosslyn 1994). As the debate grew, the evidence attained for visual imagery was extended to all imagery activity, and visual images have been taken as a paradigmatic “example” of a more general ability to generate and process internal objects regardless of the sensory modality of the single image.

The only way to legitimize this extension is by investigating: (1) the specificity of mental images from different sensory channels; (2) intermodal connections that could support the asserted pervasiveness of the visual modality in imagery.

In our opinion, comparing visuo-spatial mental images to im-

ages generated by means of other modalities may shed light on the issue of the depictive nature of mental images and on the role of mental imagery in reasoning.

As regards point (1), research on images that are not based on a visual representation is very rare. It is concerned mainly with the self-evaluation of the ability to form images in terms of vividness (Betts 1909; Sheehan 1967), although it does not try to define the dimensions along which each specific representation develops, nor does it examine the effects of the prominence of the visual modality in imagery, thus completely bypassing point (2).

As a consequence of the differences in the techniques employed, the data derived from neuroimaging and psychophysiological research are far from being conclusive. It is, however, possible that a more systematic investigation with these techniques could make an essential contribution to clarifying the nature of mental imagery. On this note, the literature is quite contradictory with regard to the modalities investigated. D’Esposito et al. (1997) concluded by means of fMRI that visual *association* areas, and not primary visual areas, were engaged during a visual image generation task.

Few studies examined the neural correlates of modality-specific processing different from the visual one. Fallgatter et al. (1997) examined ERPs for auditory, tactile, and visual imagery and found a distinct localization of the electrical brain activity. The evoked potentials were located mainly in the left hemisphere with tactile imagery, in the right hemisphere with visual imagery, and along the midline with auditory imagery, suggesting that different regions contribute to the generation of images in different modalities.

Zatorre et al. (1996), by using a PET study, showed that auditory imagery activates the same neural substrate involved in auditory perceptual processes, for example, the secondary auditory cortex and prefrontal associative areas, but does not activate the primary auditory cortex. Moreover, in the imagery conditions they found a significant increase of the cerebral blood flow (CBF) in two inferior fronto-polar regions that “may reflect some aspects of retrieval and/or generation of auditory information from long-term memory.”

Visual-tactile integration processes have been examined by Banati et al. (2000) by means of a PET study. Their cross-modal recognition task activated mainly associative areas (inferior parietal lobules, left dorso-lateral prefrontal cortex) and they suggested that these areas are responsible for the binding of information into hetero-modal representations.

Data on motor imagery reveal that the Internal Simulation of Movements (ISM) generally involves the processing of the supplementary motor area and the cingulate areas while usually the primary motor areas are not involved (Höllinger et al. 1999; Jeannerod 1995; Johnson et al. 2001).

The olfactory and gustatory modalities have been virtually ignored by the imagery research, although the olfactory system has been investigated by Zald and Pardo (2000) by reviewing PET and fMRI data on odor processing. They suggested that a restricted region of the orbito-frontal cortex plays an important role in recognition of odors, and that it responds differently depending upon the type of odor and /or the specific task demands.

From this research we can derive a general concordance about: (1) the specificity of imagery processing according to the modality, and (2) the activation of supramodal-associative areas in all the investigated modalities.

According to our knowledge, the first study concerning all possible imagery modalities is Olivetti Belardinelli (2001; see also Del Gratta et al. 2001). In this study, seven different imagery modalities (visual, auditory, tactile, organic, kinesthetic, olfactory, and gustatory), investigated by means of fMRI, exhibit a specific pattern of activation. In general, primary areas were never activated, while a compound pattern of activation was found in secondary areas and in amodal integrative areas. Visual, auditory, tactile, olfactory, and organic imagery activated the middle-inferior temporal regions and the inferior parietal lobules bilaterally, although the superposition of the activated areas among modalities was fairly

rough. Distributed activation was also observed in prefrontal areas, mainly in the middle orbital region with almost all modalities, except the kinesthetic one, which showed a bilateral activation in the cingulate cortex. Finally, visual and olfactory modalities exhibited the activation of the left hippocampal/fusiform gyrus, while visual and gustatory modalities showed the activation of the right cingulate cortex.

Taken together, all the cited studies show that mental imagery involves mainly the activation of associative cortical areas while there is little consensus about the involvement of primary cortical areas. Moreover, it has been found that the modality-specific areas were distinct, depending on the modality used as imagery cue. Finally, it has been shown that supra-modal associative areas in the parietal and in the prefrontal cortex were also activated. Whether these areas are involved in generating images, and whether they reflect either the generation process or the maintenance of mental images, are still open questions.

Although further research is needed to understand the relationship between the neural substrate of mental imagery and the nature of mental images, these data suggest that mental imagery requires at least two distinct processing components: a modality-independent component, presumably reflecting long-term memory retrieval processes (abstract/propositional recovery of object information), and a modality-specific component, reflecting short-term memory maintenance (concrete/analogical representation of perceptual objects).

Mental imagery is simultaneously symbolic and analog

John R. Pani

Department of Psychological and Brain Sciences, University of Louisville, Louisville, KY 40292. jrpani@louisville.edu
<http://www.louisville.edu/~jrpani01>

Abstract: With admirable clarity, Pylyshyn shows that there is little evidence that mental imagery is strongly constrained to be analog. He urges that imagery must be considered part of a more general symbolic system. The ultimate solution to the challenges of image theory, however, rest on understanding the manner in which mental imagery is both a symbolic and an analog system.

Professor Pylyshyn's article is one of the more insightful and informed analyses of theories of mental imagery in the long history of this topic (see Pani 1996). But although it is successful in pointing out the weaknesses of certain standard ideas, it does not fully describe the alternative.

A strength of Pylyshyn's analysis of image theory is the clarity of his primary goal: to consider whether image phenomena are due to a cognitive architecture that is intrinsically depictive or to a more general computational system that happens to generate depiction, or apparent depiction, when it is applied to particular problems. Much of the reach of this paper comes from the insightful application of this question to the experimental literature on mental imagery. Briefly consider three examples that I think exemplify the analysis.

The demonstration by Podgorny and Shepard (1978), that a block letter "imaged" over a grid decreases response time to events in the covered cells of the grid, has become a standard demonstration of the seamless interaction between imagery and perception. Although the demonstration is compelling, it likely has little to do with mental images. The block letter is generated by perceptual organization and selective attention within the perceived grid, much as one sees patterns in tile floors, and the associated behavioral effects concern high-level perception rather than imagery. Pylyshyn is to be thanked for challenging inaccurate characterizations of this experiment.

The second example concerns mental scanning of images (e.g.,

Kosslyn et al. 1978). A clearly appropriate, and important, conclusion to be drawn from this work is that people are capable of behaving as though they were consulting pictures in their heads, and that this capability is furnished neither by verbal encoding nor by an internalization of overt behavior. On the other hand, the common suggestion that scanning across images is something that people are generally constrained to do when imaging, is unwarranted. Pylyshyn makes this point effectively, and the clarification is welcome.

The third example concerns the importance of recognizing, as Pylyshyn does, that there are experimental demonstrations in which it seems that people really are constrained to generate analog imagery in order to reason correctly. Pylyshyn chooses mental paper folding (Shepard & Feng 1972) as a prime example of this, and he is right to do so. (The de facto standard test of spatial ability, the DAT Space Relations subtest [Bennett et al. 1989], is much like the paper folding task.) In discussing this task, Pylyshyn points out that the behavior that indicates analog representation is that people solve the problems in steps corresponding to folds. But this means only that the knowledge people bring to bear on the problem is organized in terms of what happens one fold at a time. This is a constraint that will force any representational system to be selectively analog as it computes the answer to the problem. That is, much of the analog nature of thinking comes from its computational context; reasoning that works from a set of premises to a desired conclusion may create a pattern of representations that corresponds directly to concrete objects and transformations. (See Pani 1996, for a similar analysis of imagery phenomena.)

Pylyshyn's ultimate theoretical goal is to find a formal characterization of the computational nature of imagery. I think it useful to point out, however, that the type of system he appears to favor fits well with one fairly broad perspective on the nature of imagery. To put it bluntly, mental imagery did not evolve to function in the manner of cameras and photocopiers. The presence of an image in mind, its structure, and its function are determined by the requirements of a computational system which generates only the information it must generate to achieve its ends. The computational role of imagery will typically determine that images are very little like pictures and that the individual doing the thinking is completely unconcerned with whether they are or not. This view of imagery was common in what used to be called American Functionalism and which now might be called adaptationism (Pani 1996). As Ladd (1894) put it, "If one arrives at the other side of the stream in safety, one does not notice or remember how each floating block of ice felt, as it was touched lightly with the toes – one's eyes and interests being set *on that other side*" (p. 284, emphasis in the original).

Although Pylyshyn's critique of image theory is admirable in many respects, the end result begs an important question. Why do so many of our mental symbols seem like mental pictures? It is important to remember in this regard that human vision and human thought are special purpose biological systems, and there is a great deal of hardware that is specialized for performing challenging jobs efficiently. Nothing Professor Pylyshyn says in this article eliminates the possibility that there is a dense mapping between concrete structures in the world and physiological structures in the brain, nor that generating thoughts about things in the world sometimes involves activating parts of those neural structures involved in the mapping. It remains possible, even probable, that when imagery occurs in reasoning – due to computational constraints that would force any system to be at least selectively analog – and using representations that are symbolic in every important sense of the term – the representations depend on tokens borrowed from the neural mapping of the visual world. Such symbols may be highly selective and abstract compared to photographs, but the visuospatial properties they do contain will lead people to report them as experiences in a mental world with an analog character. Put another way, what is missing from Pylyshyn's account is that a symbolic system may draw its symbols from a set of analog tokens, and it definitely will do that in cases where that

is the only way to get them. The representations are no less symbolic for being analog, and they are no less analog for being symbolic (Pani 1996).

ACKNOWLEDGMENT

I am grateful to Julia Chariker for comments on an earlier draft of this commentary.

Mental imagery in memory psychophysics

William M. Petrusic^a and Joseph V. Baranski^b

^aDepartment of Psychology, Carleton University, Ottawa, Ontario, K1S 5B6, Canada; ^bDefence Research and Development Canada, 1133 Shepard Avenue West, Toronto, Ontario, M3M 3B9, Canada.

bill_petrusic@carleton.ca joe.baranski@dcim.dnd.ca

Abstract: Imagery has played an important, albeit controversial, role in the study of memory psychophysics. In this commentary we critically examine the available data bearing on whether pictorial based depictions of remembered perceptual events are activated and scanned in each of a number of different psychophysical tasks.

Imagery in symbolic comparisons? In a landmark paper initiating contemporary studies of human memory processes using psychophysically based methods, Moyer (1973) asked his participants to make comparative judgements of the size of animals; for example, “which is larger, whale or moose?” Moyer also required his participants to provide estimates of the size of each of the animals named. He showed that response times were linearly related to the logarithm of the differences in standardized estimated sizes. Noting the striking parallel with Johnson’s classic (1939) plot of perceptual comparison times as a function of the logarithm of physical difference in length of the visual extents compared, Moyer argued that perceptual and symbolic comparisons might be based on common processes and/or representations and he referred to such studies with remembered magnitudes as *internal psychophysics*.

Although Moyer (1973, p. 183) was explicit in postulating an analogue basis for the representations of remembered animal size, he was careful not to prejudge the basis for this representation; “they may be positions along an imagined spatial dimension, temporal patterns in neural images, rich images, or an as yet unimagined possibility.” On the other hand, attesting to the charm and allure of the phenomenology of “rich images,” the imagery position has gathered considerable force. For example, Ashcraft (2002, p. 460), in the most recent edition of his introductory cognition text, states in discussing Moyer’s findings, “What is fascinating is that the judgments are being made on the basis of the visual image of the object. That is, the evidence suggests that when people make larger/smaller judgments about real-world objects, they retrieve mental images of the objects, then mentally scan the images to determine which is larger or smaller.” In fact, though, where does the empirical evidence stand on imagery in symbolic comparisons, now some nearly three decades after Moyer’s seminal work?

The range effect. Moyer and Bayer (1976) first trained participants to associate circles varying in size with nonsense syllables (CVCs) using a paired associate learning procedure. One group of participants worked with a relatively widely spaced set of stimuli and another group worked with a narrow range. In the second phase of their experiment, participants compared the sizes of pairs of circles directly (Circle–Circle comparisons) or pairs of circles from memory (CVC–CVC comparisons). Comparisons with perceived and remembered stimuli were uniformly faster when the wide range was used. Accordingly, Moyer and Bayer concluded that this *range effect* was sufficient evidence for an analogue based interval scale representation of size information in memory. However, on balance, the empirical evidence for the range effect, critical to establishing a necessary condition for analogue representations is mixed. The range effect is not always obtained (e.g., Banks

1977; Banks et al. 1982), although it can be if the necessary conditions are satisfied (e.g., Petrusic et al. 1998a).

Propositionally based-semantic coding theories of symbolic comparisons. In a series of forceful and striking papers, Banks and his students have provided a powerful antidote to the imagery scanning perspective. First, Banks and Flora (1976, Experiment 1) showed picture versus word effects were strictly additive with symbolic distance, establishing that the faster picture processing is due to encoding and not to specialized imaginal processing as posited by Paivio (1975). Second, Banks and Flora (1976, Experiment 2) demonstrated a symbolic distance effect with an abstract continuum, the intelligence of animals, portrayed either pictorially or verbally. Moreover, picture processing for this abstract continuum was faster than with words, although Paivio’s (1971; 1975) dual coding theory would predict faster processing with the words. Finally, Banks (1977) and Banks et al. (1982, Experiment 3) found that CVC–CVC comparisons are faster than Percept–CVC comparisons (i.e., the set-size reversal effect). The analogue-mental imagery view is contradicted by this finding because two mental image operations are required in the case of CVC–CVC comparisons whereas only one such operation is required in the case of Percept–CVC comparisons.

Failures to replicate. However, not all the evidence provided by Banks and colleagues has stood the test of replicability. The strictly additive effects of picture versus word effects on distance and semantic congruity effects are simply not evident in recent work by Shaki and Algom (2002). Moreover, Petrusic et al. (1998a) failed to obtain the set-size reversal effect so critical to rejection of the hypothesis that generation and scanning of two mental images must necessarily take longer than just one. Thus, taken together, the available evidence leaves a muddled and incomplete picture. It remains to be firmly established whether analogue representations, a necessary but not a sufficient condition for the pictorial view of mental imagery, even exist with symbolic comparative judgements.

Analogue representations and psychophysical methods. However, the evidence in support of analogue-based codes in psychophysically based tasks where the comparison cannot be resolved on the basis of ordinal, propositionally based codes alone is decisive. Baranski and Petrusic (1992) and more recently Petrusic (2001), first trained their participants to associate CVCs with visual extents. Subsequently, these CVC-labels served as standards in the method of constant stimuli with variable perceptual stimuli. They found: participants could perform the task with a high level of accuracy, larger Weber fractions were obtained for the remembered standards than for the perceptual standards, the Weber fractions exhibited end-point effects, and they systematically varied with set size, range, and acquisition conditions. Petrusic et al. (1998b) also trained their participants to associate labels with line lengths. Subsequently, participants indicated which pair of two pairs of labels corresponded to the more similar pair of perceptual referents. They then showed that the similarity comparisons were based on the computation of differences of differences of analogue based interval scale representations. Importantly, although these studies clearly implicate an analogue-based interval scale representation, there is nothing in these data that inexorably force pictorial based imaginal representations.

Magnitude estimation: Memory for elementary sensory magnitudes as re-perception. Moyer et al. (1978) and Kerst and Howard (1978) demonstrated that numerical magnitude estimates of the perceived and remembered sizes of objects were well described by power functions of their physical sizes. Moreover, they also asserted that the input to memory is a power function of the perceptual magnitude. Formally, $S_M(x) = aS_P(x)^B$, where B denotes the exponent for the transformation from perception to memory, and $S_M(x)$ denotes the subjective magnitude in memory and the psychophysical function for perception is given by, $S_P(x) = a_p x^P$, where a_p and P , respectively, denote the unit of measurement and the exponent on the subjective perceptual scale, and $S_P(x)$ denotes the perceptual magnitude of a stimulus with physical magnitude, x . According to the strict form of the *re-perception*

hypothesis, $B = P$. Consequently, the memory psychophysical function is given by $S_M(x) = a_M x^M$ with $a_M = aa_p$ and $M = P^2$. Thus, the re-perception hypothesis is the strongest form of the image as picture hypothesis in memory psychophysics.

However, its empirical status is also dubious. To date, a number of studies have provided support for the re-perception hypothesis with $B = P$, (e.g., Algom & Lubel 1994; Bradley & Vido 1984; Kerst & Howard 1978; Moyer et al. 1978). On the other hand, Moyer et al. (1982) showed that when $P > 1$ (e.g., heaviness and sweetness), the memory exponent is less than the perceptual exponent, although it should be larger than it. In addition, and also contrary to the re-perception hypothesis, Petrusic et al. (1998b) showed that the memory exponent (0.697) was considerably larger than the predicted square of the perception exponent ($0.564^2 = 0.318$) (see also Algom 1992).

Summary. The consensus of the evidence in memory psychophysics provides little, if any, support for the picture theory of mental imagery. Indeed, as Pylyshyn states, “nothing is gained by attributing a special format or special mechanisms to mental imagery” (sect. 1.2, para. 2).

ACKNOWLEDGMENT

This research was supported by a grant from the Natural Sciences and Engineering Research Council of Canada to W. M. Petrusic.

Neural representation of sensory data

Jonathan Polimeni^a and Eric Schwartz^b

^aDepartment of Electrical and Computer Engineering, Boston University, Boston, MA 02215; ^bDepartment of Cognitive and Neural Systems, Department of Electrical and Computer Engineering, Department of Anatomy and Neurobiology, Boston University, Boston, MA 02215. jonnyreb@bu.edu eric@bu.edu <http://cns-web.bu.edu/Profiles/Schwartz.html>

Abstract: In the target article Pylyshyn revives the spectre of the “little green man,” arguing for a largely symbolic representation of visual imagery. To clarify this problem, we provide precise definitions of the key term “picture,” present some examples of our definition, and outline an information-theoretic analysis suggesting that the problem of addressing

data in the brain requires a partially analogue and partially symbolic solution. This is made concrete in the ventral stream of object recognition, from V1 to IT cortex.

1. What is a “picture”? The core problem with the “picture theory” is the lack of a definition of the key term “picture.” A more correct term has been suggested earlier (Schwartz 1980) – *computational anatomy*, the properties of *locally regular feature maps*. The “little green man” problem is clarified by noting that:

(1) No known feature maps are isometric. They do not preserve metric information. They are “*distorted*,” but

(2) The lack of metric structure is irrelevant to their potential semantic content. There is no “little green man” to be confused by the (distorted) non-isometric maps of the brain. There are only neuroscientists looking inwards – and, if they wish not to be confused by what they see, they have only to learn about existing mathematical accounts of the feature maps of the brain. Figure 1 and Figure 2 are two examples. But the neurons of the brain, “looking” at feature maps from within, have no such problem!

2. Neural computation: A definition.

Definition 1 (Neural computation). Neural computation, like all computation, is based on a correct (i.e., expedient) choice of *data structure* and *algorithm*. Neural feature maps can be viewed as a form of data structure. There is little to say about neural algorithms at present – no one has ever observed a non-trivial neural network in vivo – but there are abundant experimental observations of neural data structures.

Definition 2 (Computational anatomy). Patterns of topographic mapping and columnar architecture are two prominent forms of (spatial) data structure in the brain. The key requirement is that nearby neurons in a laminar sheet must have trigger features that are nearby in some feature space.

Example 1 (Receptotopic maps of V1, V2, V3, V4, MT, MST, LGN, S. Colliculus, S1, A1, etc.). For receptotopy, the feature space is \mathbb{R}^2 , for example, the retinal surface, the body surface, or the cochlear surface. See Figure 1 for an example of human V1, V2, and V3.




Figure 1 (Polimeni & Schwartz). (a) “Retinal” view of US Naval Academy, high resolution. (b) Model of V1-V2-V3 complex, produced as a single dipole map function (Balasubramanian et al. 2002). The dipole map is a direct generalization of the familiar log-polar model of V1 topography (Schwartz 1994). V1 is the central “ovoid” region, V2 is the first surrounding “ring” and V3 the outer “ring” of cortex. (c) The USNA image is mapped via the complex dipole map to create an image model of the V1-V2-V3 complex. A face in a window of the USNA (a), which is visible in the original high resolution image, is clearly seen, repeated three times in the foveal representations of V1, V2, and V3. The entire campus of the USNA is compressed, via the highly non-linear cortical magnification factor, into the para-foveal and peripheral regions of the image. This figure represents *only* the topographic aspects, not the ocular dominance, orientation map, or other spatially represented data. The result may look confusing to a neuroscientist observer, but we believe that the brain has little problem interpreting this complex spatial data structure.




Figure 2 (Polimeni & Schwartz). Mapping a stereo-pair (a), (b) into a model of V1 layer IV ocular dominance columns. The disparity is represented as a visual “echo” or offset of repeating image elements. A non-linear cepstral filter extracts the stereo disparity in the form of a subsequent spatially mapped representation (c) (Yeshurun & Schwartz 1989) (see the proto-column model of (Landau & Schwartz 1994)).

Example 2 (Orientation columns in V1, direction columns in MT). The feature space is \mathbb{P}^1 (V1) and \mathbb{S}^1 (MT), orientation and direction, respectively. The target space is V1(or MT), the “pinwheel” pattern of orientation (or direction) tuning. The pinwheel structures result from the singularities associated with the different topological structure of the feature space and the cortical target (Schwartz & Rojer 1991; Wood & Schwartz 1999).

Example 3 (Ocular dominance columns in V1). The feature space is a (double-sheeted copy) of a visual hemi-field (\mathbb{R}^2). The target mapping is interlaced via a *proto-column* construction (see Landau & Schwartz 1994) to a locally regular map of the two half-fields, as shown in Figure 2.

Pylyshyn briefly mentions the existence of topographic structure in V1, but omits mention of the (approximately) 30–40 other visual topographic areas, as well as the other sensory modalities. Furthermore, he omits columnar structure entirely from this discussion. That is, the fact that the neo-cortex is largely organized in terms of feature maps and that these features maps are potentially semantic.

3. Information cost of addressing symbols. Pylyshyn’s key unstated assumptions are clear from his “null hypothesis”: “reasoning with mental images involves the same form of representation and the same processes as that of reasoning general” (target article, Abstract).

We agree that reasoning about pictures may well use the same processes as reasoning in general. The problem here is the *unstated* assumption that “reasoning in general” is symbolic. But, do we really know that reasoning is itself not mediated by spatio-temporal representations, that is, a “picture theory” of reasoning? There appears to be an implicit assumption in parts of the cognitive science community that computation is symbolically mediated. We will now present an argument in support of the idea that anatomy as data structure is an unavoidable consequence of the high cost of “addressing” symbolic data in the brain.

Recently, Rieke et al. (1998) demonstrated that the spike sequence of the H1 neuron of the fly is a temporal replica of the sensory stimulus. There are only two H1 neurons in the fly brain, one for each side. They call this idea “flynculus.” This does not seem like a good candidate for “symbolic” coding. The fly uses *time* to code time, not symbols to code time. Time is free, and the fly is short on neural space. The semantic meaning of an H1 spike is, in part, the time that it occurs.

The parallel is clear: attaching a spatial label to a spike in V1 is potentially expensive. There are about 10^5 resolvable spatial loca-

tions in the human visual field (see (Rojer & Schwartz 1990) for derivation) – 17 bits. The semantic content of a spike in V1 is probably no more than 2–5 bits. The obvious solution is to use physical space to code visual space. Progressing from V1 to V2 . . . and on to IT, the spatial precision becomes lesser and the semantic content of a spike becomes greater. It is expedient to pay the price for a symbolic code (i.e., axons, labeled lines, “grandmother cells”). It seems that one feasible solution for spatial coding of visual stimuli is a gradual transition from a largely (but not completely) spatio-temporal code near the periphery (i.e., V1, V2, . . .) to a largely (but not completely) symbolic code centrally (. . . , V4, IT).

4. Summary. In our analysis, we have not addressed the issue of “imagery.” It seems obvious that the real issue is visual representation, and the first area that needs clarification is the representation of visual stimuli, not mental re-creations of them.

Behavioral-level experiments are impotent, in principle, to address questions of neural representation. Purely symbolic and purely analog “machines” can easily mimic each other at the behavioral level.

If the brain is really the symbolic processor that Pylyshyn seems to envision, then it certainly has an inordinate fondness for “pictures.”

Time matters! Implications from mentally imaged motor actions

Markus Raab^a and Marc Boschker^b

^aCenter for Adaptive Behavior and Cognition, Max Planck Institute for Human Development, Berlin, 10195 Berlin, Germany; ^bDepartment of Movement Behavior, Vrije Universiteit, Amsterdam, 1081 BT Amsterdam, The Netherlands. raab@mpib-berlin.mpg.de mboschker@fbw.vu.nl www.mpib-berlin.mpg.de www.fbw.vu.nl/~mboschker/

Abstract: Pylyshyn provides sound arguments against the dominant picture theory of mental imagery. However, we claim that mental imagery is intrinsically dynamic and that the very nature of mental imagery will not be uncovered by studying static pictures. Understanding mental imagery of motor actions reveals that any theory of mental imagery should start off with the temporal nature of real-life experiences.

Pylyshyn’s criticism of the picture theory of mental imagery is inspiring and a welcome counterpart to those theorists who thought that they had resolved the debate (see Kosslyn 1994). Although we appreciate his contribution, we believe that a fundamental aspect

of all mental imagery is missing and that the debate, regardless of how it is defined by the participants, focuses on a rather artificial topic, which might be one of the factors that makes the debate seem so persistent and insoluble. In this commentary, we will focus on what the study of human motor and sports behavior can add to the ongoing imagery dispute.

The imagery debate, as a discussion on the nature of mental imagery, has neglected the temporal aspect and limited itself to visual imagery of pictures, disregarding all other kinds of imagery. We think the consideration of mentally imaged motor actions would contribute much in the search for a theory of mental imagery, a theory that should proceed from the inherent dynamic nature of all mental images and should not include solely geometric properties. For instance, Pylyshyn incorrectly argued, that “intrinsic properties of images are geometrical rather than dynamic” (sect. 2, para. 3; see also, Pylyshyn 1999). To sustain this, he insisted that baseball fielders predict the point where the ball will land. However, this is not what they seem to do (Gigerenzer & Selten 2001). Rather, they use continuous optical information (e.g., tracking optical acceleration of an approaching ball) to guide locomotion in catching fly balls (Oudejans et al. 1999).

Basic to the debate is a discussion on how visual perception operates. Mental images occur in the absence of the imagined scene; therefore, imagery is based on the recall of any kind of experience. The imagery debate is to a certain extent reducible to a discussion on what the nature, or in Pylyshyn's term, the cognitive architecture, of visual perception is. Kosslyn's notion of visual mental imagery, the picture theory, is certainly most closely related to the widely rendered idea that perception is about appointing meaning to detected stimuli and the processing of percepts. Pylyshyn, on the other hand, suggests that there is something special about mental imagery that is not found in the picture-like view of imagery, and although he does not elaborate in depth on this special characteristic of imagery, he also seemed to support the idea that visual perception starts with the detection of geometric properties.

It seems that both parties in the debate have adopted an inappropriate view of perception. Perception is not the passive detection of meaningless geometric properties, but the active pick-up of the significant features of an ever-changing environment. The imagery debate focuses on visual perception of motionless pictures and ignores the fact that behavior is about change and adaptation to changes in one's environment. The human nervous system, and the nervous system of every organism, is constructed to detect changes. When sensory organs (on the retina, in our muscles, or in our skin) are exposed to a constant stimulation, the produced signal quickly decreases. For example, it is very difficult to tell if the temperature is 8 or 14 degrees Celsius, but a change of half a degree is easily detectable. Visual perception is equipped to pick up the most relevant or meaningful features of the environment, and these features appear to be predominantly changes within the optic array. When mental imagery is based on previous experiences, it will thus reflect the dynamic nature of perception (Freyd 1987). We need a fundamentally different notion of perception and imagery that is based on the dynamics of the perceptual system.

Moreover, pictures are in fact rather artificial features. Although one could argue that distilled pictures are basic to perception and that motion is established by a specific sequence of adjacent pictures related to each other in time, this is more or less turning the world upside down. The human natural environment or “ecological niche” does not consist of computed displays or printed pictures. The human perceptual system evolved long before such things existed. Pictures are man-made, frozen moments in time that have little to do with normal behavior of organisms. It seems odd to imply that such an artificial activity could ever reveal the very nature of any behavior.

If dynamics are so essential in behavior, why not start with movement imagery, instead of the rather artificial activity of imagining pictures? But what do studies of mentally imagined motor

actions tell us about the nature of mental imagery? For instance, it is generally recognized that mental images are most vivid when all sensory modalities (visual, kinesthetic, haptic, auditory, olfactory, and taste senses) are involved and not only the visual one (see Janssen & Scheik 1994; Weinberg & Gould 1995). Second, a mental image of solely (visual) stimuli properties is not likely to affect behavior; for this, it is necessary to mentally imagine (motor) responses as well, that is, to simulate behavior (Lang 1979). Finally, Boschker et al. (2002) indicated that the effects of movement imagery are most pronounced when the imagined motor action is used to interfere with subsequent behavior, instead of enhancing it. These findings suggest that dynamic images of action scenes are not only possible to create, but highly effective in modifying behavior. This implies that mental imagery is all about responding: “our general view . . . is that the mind (i.e., the ensemble of cognitive events) is a system for organizing and directing responses” (Lang 1987, p. 408).

In conclusion, “in search of a theory” for mental imagery, the fundamental dynamics of every living creature should not be neglected. A theory of mental imagery of motor actions has to integrate temporal and kinesthetic properties of the image. These properties are neither pictorial nor spatial and are highly relevant to explain why temporal components of real and imaged actions are highly correlated (Frak et al. 2001) and why images are constrained by the biomechanical properties of the body (Munzert & Raab, in preparation).

ACKNOWLEDGMENT

We thank Peter M. Todd for helpful comments.

The imagery debate: *Déjà-vu* all over again?

Peter P. Slezak

Program in Cognitive Science, University of New South Wales, Sydney, NSW 2052, Australia. p.slezak@unsw.edu.au
http://www.arts.unsw.edu.au/sts/peter_slezak.html

Abstract: The imagery debate re-enacts controversies persisting since Descartes. The controversy remains important less for what we can learn about visual imagery than about cognitive science itself. In the tradition of Arnauld, Reid, Bartlett, Austin and Ryle, Pylyshyn's critique exposes notorious mistakes being unwittingly rehearsed not only regarding imagery but also in several independent domains of research in modern cognitive science.

Pylyshyn's return to the fray means that at least one thing may be said with certainty about the imagery debate: Despite Kosslyn's (1994) claim to have resolved the controversy, there has been no progress at all. Worse still, if Pylyshyn's null hypothesis is right, we don't have a viable theory of imagery of any kind. The “tacit knowledge” rival to pictorialism, is not itself an alternative theory but rather an indication of the direction in which an adequate theory might be sought – that is, as a theory of high-level belief or knowledge representation.

Pylyshyn's central criticism of pictorial theories echoes Descartes (1637/1985) who insisted that it is enough that the mind should adequately *represent* the properties of the world and does not have to *share* them. In the same vein, Edelman (1998) recently said nobody thinks that a mental representation of a cat is furry. Perhaps not, but it is telling that such views must be repeatedly refuted throughout the history of speculation about the mind. U. T. Place (1956) famously sought to counter a common objection to materialism by pointing out that, regardless of phenomenology, there is no green brain state when having a green after-image. Pylyshyn's charge has been that pictorialism commits a precisely parallel fallacy.

Visual images are no more likely to reveal underlying brain mechanisms than the myriad other things we are capable of thinking about. Pylyshyn's attribution of modern experimental results

to subjects' beliefs and expectations avoids the embarrassment of having to postulate vastly many and varied properties of the functional architecture. Thus, using the same experimental paradigm as the celebrated mental rotation and mental scanning experiments, we may obtain robust reaction-time evidence from a similar imagery task – “mental bouncing.” Subjects are asked to imagine holding a basket ball and are told to say “boing” (or press a key) each time it is imagined to bounce after being released. By parity of reasoning with mental rotation and scanning, the exponentially decreasing times between successive responses might lead one to conclude that the underlying processes have the property of inelastic deformations – that is, the brain is made of rubber. Or, perhaps we should say “quasi-rubber.”

For some reason, the case for spatial properties has seemed much more persuasive than the same point regarding rubberiness, greenness, or furriness. In view of their compellingness, such mistakes evoke Kant's distinction between mere errors and those deeper, inherent cognitive illusions. Thus, I disagree with Pylyshyn only regarding his optimism in hoping that, by repeating his powerful arguments loudly and slowly, he might succeed this time where he has failed before. Sufficient ground for my skepticism is the fact that the Imagery Debate is perhaps the most remarkable modern duplication of controversies concerning the nature of “ideas” which have persisted not just for thirty years but since the seventeenth century. In this recent re-enactment, Pylyshyn has played Arnauld (1683/1990) against Kosslyn's Malebranche (1712/1997). It is a striking and significant fact that the central error charged by Arnauld is that of ascribing corporeal properties to mental ones – exactly the one charged by Pylyshyn against pictorialists. In its modern guise, this is the charge of confusing properties of the *things* represented with properties of their *representations*.

Of course, Pylyshyn is not vindicated merely because he was anticipated by Descartes and Arnauld. The striking historical parallels suggest that the fundamental problems at stake do not arise in any essential way from the data of modern experiments and computational theories. Indeed, just as we would expect in this case, we see a recurrence of the same perplexities not only throughout history, but also in more or less independent domains of cognitive science today (see Slezak 2002a; 2002b; 2002c). Pylyshyn's case against pictorialism is strengthened when it is seen transposed and deployed in quite unrelated domains of cognition. Thus, for example, Carruthers (1996) defends the now unfashionable claim that we think in a natural language rather than Fodor's (1975) “language of thought.” Undeniably, the theory has a certain persuasiveness: Just as we seem to visualize in pictures, so we seem to think in language. However, paradoxically, this very intuitive plausibility provides the strongest case against such theories for, as Pylyshyn observes, we may be deeply deceived by our subjective experience. Ryle (1968) suggested that the very idea that we think “in” language is incoherent, and the introspective experience of talking to ourselves cannot support any claim about the vehicles of thought. Significantly, Ryle also mentioned the doctrine of mental pictures seen with the “mind's eye” as sharing the unintelligibility of thinking in language.

What these doctrines have in common is the mistake of assuming that we apprehend our mental states rather than just *have* them. It is clear why such an implicit conception leads to positing a representational format – sentences or pictures – which is paradigmatically the sort of thing requiring an external, intelligent observer – the notorious homunculus (see Slezak 2002a). Despite their evident irritation at this repeated accusation (Kosslyn et al. 1979, p. 574), computer simulation does not necessarily prove pictorialists' innocence. As Rorty (1979, p. 235) put it, there is no advance in replacing the little man in the head by a little machine in the head.

As Pylyshyn argues, resort to neuroscience is no help either. It is acutely ironic that pictorialism is thought to be vindicated by “developing” the cortex like a photographic plate to reveal a retino-

topic “picture” of the stimulus. Much earlier, Skinner (1963, p. 285) fantasized just such a result and remarked “In many quarters this would be regarded as a triumph in the physiology of vision. Yet nothing could be more disastrous.” Although we are unable to accept Skinner's solution, it is perhaps not surprising that he should see the *problem* very clearly in the dangers of homunculus pseudo-explanations. The topographical map on the cortex is a picture all right, but only for the *theorist* and not the monkey. As he remarked, “Seeing does not imply something seen.” (1963, p. 287).

Despite jaundiced views of philosophy as distinct from “strictly empirical science” (Finke 1989, p. 129; Finke et al. 1989, p. 54; Kosslyn 1994, p. 409), Pylyshyn's critique suggests that there remain grounds for Wittgenstein's (1953) gibe “in psychology there are experimental methods and *conceptual confusion*.”

Neuronal basis of imagery

Evgeni N. Sokolov

Department of Psychophysiology, Moscow State University, Moscow 103009, Russia. ensok@mail.ru

Abstract: The depiction of pictures as specified points in a functional space is achieved by vector encoding. Picture-selective neurons are added to the declarative memory in the process of learning. New neurons are recruited from stem cells through their proliferation and differentiation. Electrical stimulation of the temporo-parietal cortex produces subjective scenes of the past similar to imagery.

The basic concept discussed in the article refers to “depiction” – a specification of pictures as points in a functional space. The question arises how such a space is implemented within neuronal networks. To approach a solution of the problem a universal vector model of cognitive and executive processes was suggested (Fomin et al. 1979). Later, the model was tested in color vision (Izmailov & Sokolov 1991), in the emotional expression of faces (Sokolov & Boucsein 2000), and in stereovision (Vaitkevicius 2002). Experimental data confirmed the universal characteristic of the vector model.

According to the vector model, input stimuli are encoded by excitations of four modular neurons so that each stimulus is characterized by a specific excitation vector. Due to a normalization procedure occurring in the neuronal nets, all excitation vectors become of a constant length, constituting a hypersphere in the four-dimensional Euclidean space. The excitation vectors participate in the formation of picture-selective neurons of the declarative memory. Formation of highly selective neurons with respect to presented visual patterns was demonstrated in the temporal cortex of monkeys (Miashita et al. 1991).

The declarative memory gets extended in an adult organism by means of the recruitment of new to-be-learned neurons arising in the process of neurogenesis by proliferation and differentiation of stem cells. Newly generated neurons migrate to specific brain areas: hippocampus, temporal and prefrontal cortex, where they are incorporated into neuronal nets building up synaptic contacts with target cells (Gould et al. 2001). The neurogenesis is enhanced by the novelty of the neurons of the hippocampus.

The neurons of declarative memory are located in the temporo-parietal area indicated by Penfield as the “interpretive cortex” (Penfield 1958). Electrical stimulation of this area performed during an operation results in a bright scene visualized from the subject's past. Similar data were obtained by electrical stimulation of this area through implanted electrodes (Delgado 1969). Local damage of this area leads to visual agnosia (Luria 1966). It is assumed that electrically induced and voluntarily controlled imagery have common neuronal basis in the declarative memory.

The false dichotomy of imagery

Nigel J. T. Thomas

Natural and Social Sciences, California State University, Los Angeles, Los Angeles, CA, 90032-8202. njtt@earthlink.net
<http://www.members.leeds.ac.uk/n.j.thomas70>

Abstract: Pylyshyn's critique is powerful. Pictorial theories of imagery fail. On the other hand, the symbolic description theory he manifestly still favors also fails, lacking the semantic foundation necessary to ground imagery's intentionality and consciousness. However, contrary to popular belief, these two theory types do not exhaust available options. Recent work on embodied, active perception supports the alternative perceptual activity theory of imagery.

Pylyshyn's return to the fray of the imagery debate is very welcome. In typically trenchant fashion he sets forth the serious conceptual and empirical problems afflicting pictorial (including "quasi-pictorial") theories of imagery, showing how even vaunted neuro-imaging evidence fails to support it. Despite its surface appeal, pictorialism is almost certainly false.

Much as I value Pylyshyn's new contribution, however, I fear his re-entry into the debate may serve to further entrench a false dichotomy that seems firmly established in the minds of most cognitive scientists, and in the textbooks: the view that we are faced with a stark choice between some form of pictorialist theory of imagery,¹ or, alternatively, a "propositional" theory wherein imagery (quasi-perceptual experiences and associated empirical effects) is identified with *descriptions* couched in a computational language of thought (Fodor's [1975] *mentalese*). Pylyshyn is reticent about his positive theory of imagery (indeed, I think it has never been expounded in detail²), but clearly, under the guise of "the null hypothesis," he wants to sell us the same "propositional" descriptionist theory long associated with his name.

Because its details remain so underspecified, and because *mentalese* is supposed, *ex hypothesi*, to be able to represent anything that we can conceive, there are very few empirical constraints on descriptionism as it stands. Virtually any conceivable empirical observation could be accommodated without too much strain. Probably largely because of this lack of empirical content, descriptionism has remained unpopular, despite all the problems of pictorialism.

There is a worse problem, however. *Mentalese* is conceived by analogy to natural languages such as English and to computer programming languages and representation systems set up within actual working computer programs. Inasmuch as the symbols of such systems represent anything in the outside world, they do so by convention or by stipulation, which requires beings with minds to be around to do the stipulating or to settle upon the conventions. The whole point of *mentalese*, however, is to explain how minds are possible. Fodor postulated it largely to explain the constitutive *intentionality* of thought, the fact that thoughts, including mental images, are semantically meaningful, are thoughts or images *of* something or other. Thus, to explain *mentalese* semantics as stipulative or conventional would be viciously circular.

Admittedly, a lot of philosophical effort over the last quarter century has gone into trying to devise a naturalistic semantics for *mentalese*, one not dependent upon stipulation or convention. But this work has not even begun to converge upon any generally acceptable theory. The literature has become a casuistical morass, where every positive proposal (the underlying idea often, now, obscured beneath a mass of accumulated modifications) seems decisively refuted, even thoroughly incoherent, from the perspective of its rivals.³ If a naturally meaningful *mentalese* really could exist, it would explain an awful lot, but then again, so would a homunculus. It is past due time to admit that the quest for such a language is hopeless. Certainly we cannot take the conceptual legitimacy of *mentalese* for granted.

This need not threaten computational theories of specific cog-

nitive competencies and performances, which rarely need to invoke intentionality. Once we drop the requirement that computational cognitive representations should bear or ground intentionality, less problematic accounts of such representations become available (e.g., Cummins 1996; Horst 1996). However, the descriptionist is not just explaining competencies and performances, he is trying to explain imagery, a quintessentially intentional and conscious phenomenon (Sartre 1948; Thomas, in press). (Some philosophers hold that consciousness may be explicable in terms of representations, but these proposals rely upon the representations bearing intentionality; Lycan 2000.)

Picture theorists should not cheer, however. The only plausible account of the intentionality of mental pictures or quasi-pictures is that it derives from the intentionality of *mentalese* (Fodor 1975; Thomas, in press, sect. 3.2; Tye 1991). Without *mentalese* (or, worse, a homuncular mind's-eye), inner pictures will be neither intentional nor consciously experienced.

But the assumption that pictorialism and descriptionism exhaust our options for explaining imagery arises from mere historical accident. In the 1970s computational cognitive science was a new, exciting paradigm, but imagery – *prima facie* a thoroughly un-computational phenomenon, conscious and informal – was also a newly fashionable topic in psychology, with emerging experimental evidence demonstrating its objective reality and functional significance (Kessel 1972; Thomas, in press, sect. 2.1). The notorious "imagery debate" of that era was really about how and whether the evidence on imagery could be reconciled with symbolic computationalism, and Kosslyn's quasi-pictorialism (1980) soon emerged to rival Pylyshyn's descriptionist answer. Around the same time, several psychologists (sensitive, like Pylyshyn, to the defects of pictorialism) suggested alternative, non-computational mechanisms for imagery, versions of what I call *perceptual activity theory*⁴, but their voices were drowned by the clamor of the computationalists' urgent debate.

Circumstances today are very different. Symbolic computationalism has lost much of its luster, and is certainly no longer "the only game in town." With the emergence of embodied and situated approaches to cognition (not to mention connectionism and dynamical systems theory) we need no longer remain locked into a dichotomous choice of theories developed to appease symbolic computationalists. *Perceptual activity theory* comports well with these newer approaches to cognition and has distinct conceptual and empirical advantages over both quasi-pictorialism and descriptionism (Thomas 1999). It also suggests a promising approach to naturalizing intentionality and consciousness (Thomas 1999; 2001).

Pylyshyn's critique appeals to O'Regan's work, but O'Regan's conclusions (1992; O'Regan & Noë 2001) are incompatible with pictorialism and descriptionism alike. Visual experience, O'Regan holds, arises not from the presence of representations in the brain but from the active exercise of our "mastery of the relevant sensorimotor contingencies" (O'Regan & Noë 2001) as we explore our visual surroundings. *Perceptual activity theory* holds that imagery arises from vicarious exercise of such mastery: a sort of play-acting of perceptual exploration (Thomas 1999). Although the evidence does not support pictorialism, we should not thereby conclude that Pylyshyn's "null hypothesis" is true, or even null.

NOTES

1. Whether in the idiom of symbolic computation (Glasgow 1993; Kosslyn 1980), connectionism (Julstrom & Baron 1985; Mel 1986; Stucki & Pollack 1992), neuroscience (Kosslyn 1994), or whatever.

2. Except within *very* circumscribed task domains (Baylor 1972; Moran 1973).

3. Cummins (1997) persuasively refutes of a broad class of such proposals.

4. For example, Farley (1976); Hebb (1968); Neisser (1976); Sarbin and Juhasz (1970). See Thomas (1999, sect. 2.3) for further citations.

When is enough enough? The integration of competing scientific agendas

Jozsef A. Toth

Institute for Defense Analyses (IDA), Alexandria, VA 22311-1882.

jtoth@ida.org

Abstract: This commentary asks the reader to examine Pylyshyn's target article and the imagery debate at four levels of analysis – institutional, programmatic, empirical, and individual. It is proposed that the debate follows somewhat generic patterns of discourse at all four levels, but the discourse associated with one side of the debate may or may not be expressible and evaluated in terms of the other. The different sides of the debate might better serve cognitive science if they proceed as separate research programs in their respective sub-disciplines. A more inclusive program could result, however, if the opposing approaches could somehow unite.

The imagery debate might be viewed at four levels. Institutional, in the form of “hard” and “soft” economic support. Programmatic, within a broader framework, community, or “paradigm.” Empirical, as method, evidence, and argument. Individual, involving scientists, their desires, self-esteem, careers, status, students, peers. This commentary asks the reader to consider the following:

1. What methods are employed on each side of the debate, how sound are they, what results have they produced, and how have they been interpreted?

2. Rather than focusing on any differences, what do the two sides of the debate have in common?

3. Is this really a debate, or are both sides, at times, arguing for the same result?

4. What criteria or metrics should be established for resolving the debate?

5. If the debate is ever resolved, what are the implications at each of these four levels?

Research is usually terminated when it results in injury to the participants, landing the principal investigator and participants, or their survivors, in court, where penalties and damages are decided. Irrefutable evidence, as in the classic case of plate tectonics, can also resolve scientific disagreements, but is usually accompanied by much rancor before one side finally concedes. The imagery debate, nearing its fourth decade, poses minimal risk to participants, is still in the “rancor” phase, and could conceivably continue for years, though Kosslyn (1994; personal communication 2002) considers it over. The debate follows at least two paths, theories promoting symbolic representations and computational processes – that is, Pylyshyn's approach – and theories more grounded in physiology – that is, Kosslyn's approach.

Symbolic computational models emerged after the Second World War, challenging Gestalt psychology, cybernetics, and neo-behaviorism. With limited knowledge of the brain, technology and symbolic cognitive theory sometimes evolved rapidly in near lock step fashion, equating the mind with silicon and electromagnetism, rather than with tissue. The emergence of a technology and an analogous mental representation or process was more than a coincidence, at times advanced by the same individuals, such as the mapping of higher-level programming languages with corresponding symbolic models of thought. This resulted, for instance, in the physical symbol system hypothesis (PSSH) and the knowledge level (see Edwards 1996, for a critical historical account). Five decades later, large-scale analyses comparing the performance of several symbolic cognitive architectures with human subjects are finding significant deficits in the expected versus observed characteristics of these architectures (Gluck & Pew 2002). These results and other criticisms notwithstanding, Pylyshyn's use of the symbol metaphor follows a deeper philosophical tradition of which even mainstream cognitive architectures should only be considered an instance. That is, if the PSSH is not proven according to one research agenda, the possibility exists that it could be tested and proven in a future research agenda.

For the purposes of this commentary, the imagery debate is about representation and method, or the structure and process involved in manipulating visual or protovisual images in the mind (or

brain), and how cognitive scientists generate and gather evidence for their respective theories. Anderson's (1978) discussion on this matter provided key points for the argument even as it continues today, though at the time he did not have the benefit of recent brain imaging techniques. Pylyshyn's arguments are motivated, in part, by the PSSH, suggesting that a key functional feature of brain tissue is the ability to represent, perceive, store, retrieve, and manipulate symbols (personal communication 2002). Although a clear, disambiguated definition of “symbol” still eludes cognitive science, symbols, or something like them, nevertheless appear to be an emergent or epiphenomenal feature of the brain.¹ Support for Pylyshyn's views is derived from the PSSH (Pylyshyn 2002, personal communication), behavioral evidence, thought experiments, and longstanding tenets from linguistics and philosophy.

In contrast, Kosslyn's arguments are based on sub-symbolic properties of the same neural tissue. Thus, his pictorial or depictional hypothesis centers on the idea that a key representation involved in mental imagery is a multidimensional isomorphism or homomorphism between the world and regions of brain tissue. He has embraced a variety of methods – behavioral experiments, computational modeling and simulation, brain imagery and forensics (dissociative brain function as the result of trauma) – that have shaped a convincing theory of pictorialism. Kosslyn (1992) has usually, without fail, risen to each challenge with a programmatic and empirical response.

Surprisingly, both views realize different notions of the computing metaphor – symbols, as described above, and pictorial representations, which are “. . . like an array in a computer . . .” (sect. 5.1). The symbol is once or twice removed from the tissue according to the PSSH, and the multidimensional array is a simplified model of neural tissue. Just as symbols have their fundamental problems, one issue with the array pertains to the properties of each item, a neuron, in the array, and the fact that the neurons are not grouped like data structures in a procedural or object-oriented programming language. Thus, even sub-symbolic approaches adhere, in a sense, to a form of the PSSH, since they are modeled in the Von Neumann architecture.

In his article, Pylyshyn has again posed specific challenges to Kosslyn. One dispute addresses how the fundamental notions of pictorial or depictional representations fail to meet the requirements evident in the capabilities of symbolic manipulation mentioned above. The standard connectionist definition of input and output units in neural maps is too rigid, and permits only post hoc inferences from simulation results. More recent proposals, however, are beginning to address this matter, suggesting that a large number of modules each consisting of a relatively small number of neurons can function in a somewhat generic way. Prompted by Bartlett's dynamic notions of memory, and Edelman's work from the 1980s and 90s, neural simulations, brain imaging, and forensics may help provide answers to some of these problems (summarized in Clancey 1997). Even more intriguing is recent evidence suggesting that, after an otherwise healthy brain is denied visual stimuli, the tactile processing of Braille, can, in a matter of days, call upon visual neural tissue for this non-visual task (Hamilton & Pascual-Leone 1998). Likewise, as in the case of Michelle Mack (Grafman 2002, personal communication), the right hemisphere of the brain has been found to take on the role of language after the left hemisphere has been permanently rendered inactive as the result of infant trauma (see also Grafman 2000; Grafman and Litvan 1999; Romero et al. in press).

These kinds of modular, flexible, generic, and self-organizing capabilities of brain tissue, curiously, provide evidence for both sides of the debate, fulfilling requirements for symbolic and depictional/pictorial perspectives. On the one hand, these self-organizing capabilities support the hypothesis that the structural aspects of tissue needn't be considered as important as the functional – who cares where the neurons are located, so long as they get the job done? On the other hand, it is precisely these structural and functional capabilities that provide evidence for pictorialism; particularly when combined with evidence from brain imaging studies and forensics.

In their present form, both sides of the debate are at an obvious impasse, are most likely programmatically disjunct, and, from an epistemic perspective, any comparison between the two might even be logically incoherent (Bickhard & Terveen 1995; Bickhard personal communication 1994). But we hope not. This most recent exchange strongly suggests there are matters outstanding that haven't satisfied the cognitive science community.

What is the next step?

It may be the case that both sides of the imagery debate have something to offer towards a final, or interim, resolution. Different arguments in the debate appear to originate from particular sub-disciplines in cognitive science. At some point, metrics and criteria for evaluating the arguments and evidence according to these disciplines need to be established from an objective perspective. Self-imposed metrics and excessively stringent requirements for metrics designed for an opposing perspective may not be the best way to evaluate the debate, or research in general – that is, this calls for independent review (Emanuel et al. 2002). Model comparison and evaluation can be a tall order, and it is sometimes unclear the best way to proceed, that makes all parties happy (Gluck 2002; personal communication). Pylyshyn's requirement for compositionality will persist, prompting sub-symbolic approaches to address this prerequisite, which to an extent, some are. Kosslyn's requirements for the neuroscientific basis of imagery, however, will also persist, and physiology cannot be ignored, since it will only mature over time. The frame-of-reference problem, symbol grounding problem, and similar epistemic intractabilities, remain to be fully explained by symbolic theorists. They cannot be simply characterized as an annoyance, or recast as a mathematical puzzle as grand as Fermat's last theorem. Hybridization of what were once considered competing models is also a possibility, and can produce effective operational results, unfortunately, this may not be the answer for those trying to build sound theory.

But, as Clancey (1997; personal communication 1998) suggests, perhaps it is time to reconsider the either-or discourse and instead begin considering this problem in “both-and” terms. Rather than either pictorial or symbolic, it might better serve cognitive science to recast this debate into a problem in both pictorial and symbolic terms.

ACKNOWLEDGMENTS

The final preparation and research of this document was funded through IDA Task AK-2-1801, Cognitive Readiness via Advanced Distributed Learning, Task BE-2-1624, Advanced Distributed Learning Common Framework, and Task BE-2-1601, Defense Modeling and Simulation Office Mission Review, under contract DASW01-98-C-0067 from the Office of the Secretary of Defense.

NOTE

1. This view is complementary to the earlier notion that mental images were considered epiphenomena of symbolic activity.

Involvement of a visual blackboard architecture in imagery

Frank van der Velde and Marc de Kamps

Cognitive Psychology, Leiden University, Leiden, The Netherlands.
 vdvelde@fsw.leidenuniv.nl kamps@fsw.leidenuniv.nl

Abstract: We discuss a visual blackboard architecture that could be involved in imagery. In this architecture, networks that process identity information interact with networks that process location information, in a manner that produces structural (compositional) forms of representation. Architectures of this kind can be identified in the visual cortex, but perhaps also in prefrontal cortex areas related with working memory.

Pylyshyn argues that imagery results from a reasoning process based on structural representations. However, reasoning using mental imagery is “somehow different” from reasoning without it, and the former “in some sense” involves the visual system. This

raises the question of how the visual system could process structural forms of representation. To answer this question, it is useful to look at the relation between object-based attention and imagery. Object-based attention concerns the effect of a, for instance, memorized visual object (target) on the processing of current visual information. The involvement of memorized visual information shows a link with imagery.

Work on object-based attention suggests that the visual cortex possesses a “blackboard” architecture that provides compositional representation of visual attributes (e.g., color, form, motion, relative location). For example, consider the relative location of objects on a map, as in Figure 1 of the target article. In the visual cortex, objects are identified through a feedforward network of areas, going from the primary visual cortex to the higher areas in the temporal cortex. In this network, retinotopic representation is gradually transformed into a location-invariant identity representation. Similarly, location information is processed in networks that start at the primary visual cortex. These networks transform retinotopic information into location representations related to movements of different body parts.

To determine the relative location of objects on a map, a structural relation has to be established between the identity representation of the objects and the information about their current location on the map. This is a direct consequence of the compositional nature of the task. That is, objects on a map like Figure 1 of the target article can be identified irrespective of their location relative to one another. Monkey studies (e.g., Chelazzi et al. 1993) indicate how object information can be related with location information in a compositional manner. Once an object is selected as a target, a feedback process is initiated in the networks that produce object identification. This feedback process carries information about the identity of the object to the lower areas in the visual cortex (i.e., areas in-between the primary visual cortex and the higher-level identification areas). The feedback process interacts in these areas with the feedforward process described above. This interaction results in the enhanced activation of target representations in these areas, which results in the selection of target-related location information (Van der Velde & de Kamps 2001).

The crucial aspect of this process is the nature of the representations involved. Identity representation of an object is location invariant. However, in these lower “in-between” areas, the representations consist of conjunctions of (partial) identity information and location information. This is a direct consequence of the fact that these areas gradually transform the retinotopic information in the primary visual cortex into identity-based information. By means of the interaction process described above, information about the location of an identified object can be recovered, even though that information was lost at the level of object identification. In a converse manner, the selection of a location in the “in-between” areas can be used to select the information about the identity of an object, which will result in the identification of the object on the given location. This process is compositional because the same identity (object) representation can be related with different location representations, and the same location representation can be related with different identity (object) representations, through the interactions in these “in-between” areas described above. In computational terms, these areas provide a “blackboard” architecture because they link different “processors” to one another (van der Velde 1997). The “processors” in this case are networks for object identification as well as networks for location representation. However, the architecture can be extended with networks for color or motion processing (de Kamps & van der Velde 2001).

The compositional nature of this architecture provides information about how structural representations could be formed in imagery. For instance, the task of scanning a map as in Figure 1 of the target article by imaging a flying black dot moving from one object to another (Kosslyn 1980) could proceed as follows. The objects are used to identify locations in the manner described above, with one location as the starting point. Then, the imaging of a flying black dot consists of transforming the location representation of the first object into the location representation of the second

object in a continuous manner. However, one can also jump from one object to another in a discontinuous manner. In this case, one object is selected as a starting point, which results in the selection of its location as described above. Then, the other object is selected as a target, which results in the selection of the location of this second object as the new target location. The transformation between these two locations is then accomplished by a discontinuous transformation (jump).

The location of the blackboard architecture described above is in the visual cortex, in particular the higher and “in-between” areas of the visual cortex. Besides, in object-based and spatial attention, the architecture located in these areas will likely play a role in imaging tasks that are related to perception (e.g., imaging objects on a given visual display). However, as noted above, the crucial aspect of this architecture is the nature of the representations involved. As explained, these representations result from the nature of visual processing. However, this does not exclude the possibility that an architecture of this kind could (also) be located elsewhere in the brain. Functional neuroimaging studies, using working memory tasks, have identified areas in the prefrontal cortex with combined representations of objects and locations (D’Esposito 2001). These areas could be part of a blackboard architecture of visual working memory. Maps as illustrated in Figure 1 of the target article indeed have to be memorized, before they can be used in imaging tasks (e.g., Kosslyn 1980). Imaging processes as described above could occur in such a working memory blackboard architecture as well. Given the nature of the representations involved, an integration between the different views on imagery (“spatial” vs. “structural”) could evolve on the basis of the processes in these architectures.

A visual registration can be coloured without being a picture

Edmond Wright

Faculty of Philosophy, University of Cambridge, Cambridge CB3 9DA, United Kingdom. elw33@hermes.cam.ac.uk www.cus.cam.ac.uk/~elw33

Abstract: Zenon Pylyshyn here repeats the same error as in his original article (1973) in starting with the premiss that all cognition is a matter of perceiving entities already given in their singularity. He therefore fails to acknowledge the force of the *evolutionary* argument that perceiving is a motivated process working upon a non-epistemic sensory registration internal to the brain.

Zenon Pylyshyn is on firm ground in attacking the notion of there being a *picture* in the brain for visual mental imagery. He appears to have made a case against Stephen Kosslyn’s notion of an inner picture (Kosslyn 1994).

However, nothing in Pylyshyn’s argument undermines the possibility of there being a registration in the brain. This possibility does not occur to him because he does not take account of Roy W. Sellars’ proposal that all sensory experiences are only “structurally isomorphic” to input at the sensors, that is, they are “differentially correlated” to it, not necessarily in direct ratio (Sellars 1932, p. 86). This implies that sensory phenomena of any kind are utterly *unlike* what triggers them, so that there is no external “colour” to match neural colour (e.g., the actually experienced red), no external “smell” to match neural smell, and so on for all the modalities. *A fortiori*, since real external pictures are therefore actually uncoloured, there cannot be pictures in the brain. But, nevertheless, this theory can still claim without inconsistency that there is a neural-colour registration in the brain. This also renders useless Pylyshyn’s complaint that some inner “eye” would have to scan a field in the manner of the real eye, for the direct sensory experience can be of that very scanning without any supposed movement of a supposed eye (does a TV screen have to move to “scan” round a cup²). In any case, real eyes have evolved to pick up light rays, which are uncoloured, and there are no light rays in the brain.

The experience of stereoscopic space also must bear no pictorial resemblance to the real space with which it is correlated (stereoscopic space can be turned inside out; see Nakajima & Shimojo 1981). There is no point therefore in claiming, as Pylyshyn does (sect. 5.2), that there cannot be a registration because it could not be on “a physical surface” in the brain, for there is no requirement that the neural registration be on a physical surface similar to that of an external object; structural isomorphism rules out any such similarity. Kant was on the right track (Kant 1787/1964, p. 69): sensory stereoscopic space is thus not like real space (although in some form it is *in* real space in the brain).

It might be said that Pylyshyn half-grants an inner registration when he says “Yet it is quite possible that both vision and imagery lead to the same kind of experience because the same symbolic, rather than pictorial, form of representation, underwrites them both” (sect. 6.1), but he is wrong in using the word “symbolic.” Add to Sellars’ suggestion the proposal that such structurally isomorphic fields of whatever sense modality are fundamentally non-epistemic, that is, evidence as material, as brute, as the input (Collins 1967; Wright 1996, pp. 24–28). The evolutionary advantage of this is that the motivational system is then free to select portions of what is sensed as guides to action (Piaget 1970); it can adjust these portions as new contingencies arise (human beings have evolved the ability through communication to speed up the spread of such adjustments through the species). The key advantage of such separation of non-epistemic registration and motivated epistemized selection therefrom is the openness to a continual renewal of adaptation. So the sensory fields are material, involuntary registrations that present evidence that is not in itself “informative,” no more than the grain of a piece of wood is informative, though one may be able to work out that a particular line indicates, not “symbolizes,” say, a dry summer in 1987.

It is therefore an easily understandable prejudice to begin, as Pylyshyn does, with the notion that things and persons are given, which he is doing if he talks of the sensory as “symbolizing.” He is committing the “Entity Fallacy” (Wright 1992). All his examples are of an object, which is thought of as beyond adjustment in its singularity. Singularity is merely a feature of the selection process and does not guarantee the “singularity” of portions of the external flux; if it did, *the evolutionary advantage of adaptation would be lost*. This is what “reasoning” is most often about. The only place Pylyshyn does address the possibility of adaptation is in his discussion of re-interpreting the mental image (sect. 6.5), where he denies the possibility. However, if he had taken note of my specific response to his original article on mental imagery (Pylyshyn 1973; Wright 1983), he would have seen several examples of credible re-interpretations from the internal registration both for mental images and after-images. Here he has set up a straw man, for he only considers cases where the whole of one interpretation of an image turns into another interpretation of *the same whole*, but this is insufficiently general, for there are cases where the boundaries of the original percept are *not* preserved. Here is an experiment to prove it. Ask good audiles to hear in their minds the following constantly repeated without a pause: ‘Bell-I-Mud-Dum’ (Skinner 1957, p. 282). After a while a few of the subjects will laugh (even one is enough). For those subjects the perceived boundaries (of the original four “singular” words) will have shifted over the non-epistemic base.

In conclusion: I am one of those people who can have a short nap of ten minutes and awake refreshed. Of late, upon waking I have been having a short but vivid hallucination. In the top left-hand corner of my vision appears some *non-objectified* imagery, as of Lego bricks or printed circuits seen under moving water (though “they” are none of these, since the imagery is in constant phantasmagoric transformation). It is so vivid I can half-open my eye to produce a “split-screen” effect: I can see simultaneously the hallucination and a window in front of me, both the *non-epistemic* imagery and the *epistemized* window. The hallucination still shows faintly when I open my eyes wide. It seems counter-intuitive to maintain that they are not on the same inner “display.”

Generic assumptions shared by visual perception and imagery

Qasim Zaidi and A. Fuzz Griffiths

College of Optometry, State University of New York, New York, NY 10036.
 qz@sunyopt.edu fuzz@sunyopt.edu
 www.sunyopt.edu/research/zaidi.shtml

Abstract: What is difficult to imagine is also surprising to perceive. This indicates that active visual imagery is an integral part of active visual perception. Erroneous mental transformations provide clues to prior assumptions in visual imagery, just as visual illusions provide clues to perceptual assumptions. Visual imagery and perception share generic assumptions about invariants in images of rigid objects.

Look at the picture in Figure 1. This is the image one sees when looking up at a solid object that consists of a white board attached to a vertical black baseboard. Now close your eyes and imagine rotating the baseboard 180° around the vertical axis.

Does the rotated object look like Figure 2? Every observer who has tried this visual game with the actual solid object has answered in the negative. Why is this so? Observers reason as follows: “In the object in Figure 1, the free edge of the white board is higher than the attached edge. Since this is a rigid object, after rotation the free edge should stay above the attached edge, whereas in Figure 2 it is below the attached edge.” In fact, Figures 1 and 2 are two views of the same object, a white parallelogram (internal angles 45° & 135°) attached orthogonal (horizontal) to the black baseboard (Griffiths & Zaidi 2000). In Figure 1, the free edge is closer to the observer, whereas in Figure 2 the attached edge is closer. Having been informed of this fact, now try imagining a rotation from Figure 1 to Figure 2. It is unlikely that you will be able to imagine the rotation

while retaining the assumption that the parallelogram is attached rigidly to the baseboard. In visual imagery, it seems to be assumed that the fixed spatial relationships between features of a rigid object are preserved in all images of the object. Figures 1 and 2 show that this assumption can be wrong.

The reader may have noticed that the parallelograms do not look horizontal in either of these pictures, so obviously there is a visual illusion involved. In fact, despite being allowed to handle the solid object and to look at it from all angles, observers are constantly surprised when the object is placed in the view shown in Figure 1 and the parallelogram appears almost rectangular and tilted up (Griffiths & Zaidi 2000). The image of the front edge of the parallelogram is not horizontal because in perspective the farther corner of the edge is projected down towards eye-height. The illusion demonstrates that a visual percept can be immune to knowledge of the object, and that the perceptual system cannot discount the effects of the image formation process. Why is it surprising when the parallelogram appears tilted? A plausible explanation is that an observer imagines that it should remain orthogonal to the baseboard in all images, because relationships among features should be invariant in images of rigid objects. This points out that visual perception and visual imagery are intertwined processes and that any discordance between them is a source of surprise. It is possible that we register a percept as an illusion when an object appears different from the shape we imagined based on an earlier view.

The objects in Figures 1 and 2 were designed to represent the salient features of a building, two views of which are shown in Fig-




Figure 1 (Zaidi & Griffiths). A solid white parallelogram shape with internal angles of 45° and 135°, viewed from below with the front edge oriented 45° towards the camera.




Figure 2 (Zaidi & Griffiths). A solid white parallelogram shape with internal angles of 45° and 135°, viewed from below with the front edge oriented 45° away from the camera.

ure 3 (Griffiths & Zaidi 2000; Halper 1997). Notice that in the left panel the balconies seem implausibly tilted up, whereas in the right panel they implausibly appear tilted down. A frontal view of the building reveals the balconies to be horizontal parallelograms. This knowledge never weakens the illusion.

When the object shown in Figure 1 is slowly rotated to the view in Figure 2, observers report a percept of non-rigidity. Does this rotating rigid object appear non-rigid because it violates our expectations about the images of a rigid object? After being shown this rotation, observers can close their eyes and imagine a rotation that is similar to the percept, but only by giving up the assumed rigidity of the object. However, this rotation seems to involve the visual memory of a geometrical transformation, rather than the mental geometrical transformation of a visual memory. Since violations of image assumptions used in mental transformations of rigid objects create visual percepts that are surprising, this suggests that visual imagery continuously sets expectations in active visual perception.

A three-dimensional curl illusion described by Griffiths and Zaidi (1998) illustrates that stereo and motion-parallax effects are neither incorporated in visual imagery nor anticipated in perception. The reader should cut the C-shaped object (Fig. 4) out of a flat piece of rigid material, and try to imagine what this object will look like when raised and viewed monocularly. When the flat object is held at an elevation of 45° , it surprisingly acquires an upward curl when one eye is closed and stereo information removed. Now the reader should imagine what the object will look like if it is moved up or down. When the rigid object is moved, it appears to flex and change shape, and be more or less curled as the elevation above eye level is increased or decreased. Now imagine the shape of the object if the translation velocity is faster. Most observers imagine similar percepts whether the motion is slow or fast. In actuality, the curl disappears and the object appears rigidly flat if the object is moved quickly. Thus, the perceptual system uses information provided by motion-parallax and stereo, but mental imagery is unable to take the presence or absence of such



Figure 3 (Zaidi & Griffiths). (Left) A view of the apartment building “The Future,” located at 200 East 32nd Street in New York City. The balconies appear to tilt up, with tilt becoming more apparent the higher up the building one looks. (Right) The same building, but viewed from the opposite side. The balconies now appear to tilt downwards.




Figure 4 (Zaidi & Griffiths). The shape shown at left is constructed from a pair of identical semicircles, placed parallel to one another, one radius apart, and joined by a pair of parallel straight lines. In order to better view the effect at arm's length, we suggest using a radius of 4 inches, and cutting the shape out of stiff card or some other rigid material. The stimulus is viewed monocularly, placed parallel to the ground and elevated above the line of sight. Raising and lowering the stimulus as indicated will generate the effects described in the text.

information into account. This mismatch produces unanticipated percepts that are surprising to observers.

Pylyshyn has argued persuasively against a picture theory of visual imagery. As an alternative, erroneous mental visual transformations could be studied for clues to the generic assumptions used in imagery, just as visual illusions have provided clues to generic perceptual assumptions. Visual imagery and perception seem to share the generic assumption that spatial relationships among features of a rigid object are preserved in all images of the object. This implies that both perceived and imagined visual representations include spatial relationships. We have demonstrated that what is difficult to imagine is also surprising to perceive and vice versa (also illustrated by Fig. 6 of Pylyshyn). This suggests that active visual imagery is an integral part of active visual perception.

ACKNOWLEDGMENTS

This work was funded by NEI Grants EY07556 and EY13312, to Qasim Zaidi.

Author's Response

Stalking the elusive mental image screen

Zenon W. Pylyshyn

Rutgers Center for Cognitive Science, Rutgers University, Piscataway, NJ

08854-8020. zenon@ruccs.rutgers.edu

<http://ruccs.rutgers.edu/faculty/pylyshyn.html>

Abstract: After thirty years of the current “imagery debate,” it appears far from resolved, even though there seems to be a growing acceptance that a cortical display cannot be identified directly with the experienced mental image, nor can it account for the experimental findings on imagery, at least not without additional ad hoc assumptions. The commentaries on the target article range from the annoyed to the supportive, with a surprising number of the latter. In this response I attempt to correct some misreadings of the target article and discuss some of the ideas and evidence introduced by the commentators – much of which I found helpful,

even though they do not alter my basic thesis. I also further develop the idea that the spatial character of images may come from the way they are connected to our immediate or immediately-recalled environment (by attention or by visual indexes) and towards which we may orient while we are imaging, thus leaving the alleged spatial properties of images outside the head and freeing image-representations from having to be displayed on any surface.

R1. Introduction

Nearly thirty years after the first round of the current “imagery debate” (Pylyshyn 1973) it appears that we may have made a small amount of progress. Although the views expressed by the commentators are as varied as those of the research communities from which they come, the large number of at least partially supporting commentaries, from many disciplines, came as a welcome surprise since the picture theory is still the dominant view. I had assumed that the target article would be widely reviled since it asks that one put aside one’s intuitions about what is in the mind/brain. This appears to be an exercise that few theorists care to take up, for reasons that **Slezak** and **Gold** clearly lay out: It is often called the “intentional fallacy” and scientific arguments over it go back to the beginning of the Renaissance.

Because Kosslyn’s views about mental imagery are widely accepted, and because his theory is the most explicit statement of the view I am criticizing, I will devote disproportional space to the long commentary by **Kosslyn, Thompson & Ganis** (hereafter **Kosslyn et al.**). I appreciate the well-organized set of responses by Kosslyn et al. that clarify their stand on a number of the issues on which Kosslyn and I have disagreed over the years. I especially appreciate the clear statement in the latter part of their introductory paragraph (beginning “The issue is not . . .”) that seems to represent at least a softening of Kosslyn’s earlier position. The position they outline in their commentary also diverges in a number of other respects from the canonical view presented in Kosslyn 1994. For example, Kosslyn et al. no longer appear to claim that what we experience in mental imagery corresponds to the content of the depictive display. They now maintain (item 4 of their section “Imagery and perception”) that the experience of imagery does not originate from the depictive representation (which they claim is in V1), but from the activation of memory representations in the inferior temporal lobe. This is a puzzling revision. For, if what we experience is not in the depictive display but elsewhere in memory, why do they appeal to properties of the display to explain what happens when subjects do certain things to the image that they experience (e.g., examine or scan it)?

In spite of this and other apparent revisions of the Kosslyn position, I will argue that no amount of fine-tuning will save the depiction theory.

R2. Is imagery primarily a problem for neuroscience to solve?

I said in the target article that disagreements about the nature of mental imagery are more than a question of different interpretations of data. As **Slezak**’s historical perspec-

Response/Pylyshyn: Mental imagery: In search of a theory

tive serves to highlight (and which Dennett 1991 and Thomas, in press, have also documented), the disagreements rest on much deeper preconceptions and illusions – which is why they arouse such passionate reactions, and why the long “imagery debate” does not appear to have resolved the basic disagreements. Slezak points out that even though many people recognize that a representation of a cat does not have to be furry, nonetheless “it is telling that such views must be repeatedly refuted throughout the history of speculation about the mind.”

Kosslyn et al. (as well as **Polimeni & Schwartz**) feel that the reason the “debate” was not settled earlier is because behavioral data are incapable of resolving the question of the nature of representations underlying imagery, so we have had to turn to the findings of neuroscience. Kosslyn et al. cite Anderson’s (1978) indeterminism thesis to support this view – although it is well known that data, no matter how much of it there is nor what form it takes, *always* underdetermine theory.¹ However interesting and important the neuroscience findings are, it still remains the case that the problem most of us are trying to solve is not a neurophysiological one but a psychological one: Anyone who studies mental imagery wants to understand the nature of a particular phenomenon that arises in behavior and in experience – and that means understanding its formal and information-processing characteristics as well as its instantiation in the brain. The search for neural correlates or for neural mechanisms takes it for granted that we know what they are correlates of, or what functions the mechanisms are computing, and this depends on our understanding of the phenomena and in having at least some idea – preferably one that is not obviously wrong – of how it works. It also depends on certain assumptions about how function maps on to structure; an assumption that is very often questionable, as it is in the studies of vision and imagery (see the cautionary note in Young 2000). The idea that only neuroscience can provide the answers we seek is sheer prejudice, although a widely shared prejudice that finds expression in the commentaries by **Kosslyn et al.**, **Polimeni & Schwartz**, **de Haan & Aleman**, and **Toth**.

Several commentators seem to feel that even raising the issue of the nature of mental images at this time, given how much has been written about it, is somehow in poor taste. For example, **de Haan & Aleman** claim that we have “gone beyond” the debate and that “recent research allows us to formulate new theoretical ideas concerning how we are able to mentally imagine the outside world.” But they don’t say what these new theoretical ideas are, and I very much suspect that they are the very same old theoretical ideas that were criticized in the target article. Similarly, **Toth** urges us “to reconsider the *either-or* discourse.” This is all very nice but what matters in the present context is whether the arguments I presented are valid, because if they are, the form of the “discourse” becomes moot. Progress on conceptually difficult problems, like mental imagery, is unlikely to be furthered by homilies about how science should proceed, how we cannot ignore neuroscience, and how we should strive to accommodate both sides. When someone presents neuroscience evidence that bears on the issues it will *not* be ignored and when a non-vacuous version of picture theory is proposed we can then consider the option of some “hybrid” form of representation. But we are not there yet.

R2.1. The involvement of the visual brain in mental imagery

I tried to make it clear in the target article that the disagreements are about what, if anything, is special about mental images that distinguishes them from other forms of cognitive representation. The argument is *not* about whether certain parts of the brain are involved in both vision and mental imagery. The notion that some of the same neural circuits are activated in both cases is raised by a number of commentators, including **Burgess, Bartolomeo & Chokron, Chatterjee, de Haan & Aleman, Grossberg, Olivetti Belardinelli & Di Matteo, Pani, and Polimeni & Schwartz**. But the assumption that the vision-like experience of imagery derives from the deployment of some of the same neural structures as are activated in corresponding episodes of visual perception, is not in dispute. What is in dispute is the further assumption that both vision and imagery use a special form of representation, one that has become known as “depictive.” Even if one does not make this assumption, the overlap of vision and imagery is a platitude unless one has at least the beginnings of a model of what the two have in common other than the untenable assumption that they both examine some inner display.

For this reason I welcome attempts, such as **Grossberg’s**, to develop such a model, in which top-down and bottom-up information interact to produce visual experiences on the one hand, and hallucinations or visual images on the other, and which also account for some of the major differences between vision and imagery. On the other hand, it must be admitted that this is only a very small step towards understanding either the experience of imagery or its information-processing properties, and is moreover no help at all in explaining the vast array of behavioral imagery phenomena, such as documented in **Kosslyn (1980)**. (Although, as I suggest in sects. 3 and 4 of the target article and in sect. R6 of this response, a great many of these phenomena are unlikely to be explained in terms of any imagery-specific mechanisms, but rather in terms of general mechanisms such as inference from tacit knowledge.)

R3. Depiction and the “Tootell Display”

We see in the kind of neuroscience evidence that is sought and in the way it is interpreted, that investigators like **Kosslyn et al.** are seeking evidence for an old, intuitively appealing idea of what goes on when we “see” with our “mind’s eye,” an idea I called the “picture theory,” that goes back to Descartes and was revived in modern times by **Kosslyn, Shepard, Paivio, and others**. This is clear from the fact that **Kosslyn et al.** take as their central paradigm the finding of **Tootell et al. (1982)** and similar, more recent findings cited in their commentary: namely, that in visual perception and mental imagery there is a literal two-dimensional layout of neural activity in the visual cortex that “resembles” (**Kosslyn’s** term) what it represents. Moreover, even more critically, both vision and imagery *exploit* this spatial layout and rely on the fact that a particular structure-preserving mapping defines the content of the image or percept, in a way that would not be true of a symbolic form of representation (even though the latter would likely also be encoded spatially in the brain).

This is a strong thesis with considerable intuitive appeal. It does not do justice to **Kosslyn’s** highly focused research

program that he and his colleagues back off from this thesis each time a counterexample is raised, as they do when they admit that images are not exactly like any possible picture (e.g., they contain “predigested information” – again **Kosslyn’s** term), that they are laid out in a “functional space,” that the process of inspecting images is not really like the process of visual perception, or that, even though the depictive display explains imagery phenomena, its contents are not what we experience when we image. The way that images are supposed to be like pictures is quite clear, not only in the **Kosslyn (1994)** text I quoted in the target article, but also in the points made in the **Kosslyn et al.** commentary: Images and pictures are both laid out in space in a way that “preserves metrical properties” such as “large versus small,” “near versus far,” as well as other geometrical properties (e.g., “square” vs. “circular”). And it had better be the case that the space they are laid out in is literal space in the brain, and not some softened “functional” version; otherwise none of the predictions that involve size, distance, relative location, or shape would follow from intrinsic properties of the image (and if they do not follow from intrinsic properties, then the depictive theory does not differ from the “null hypothesis”).

Many empirical arguments have been cited in support of the view that images are special insofar as they are “depictive,” or picture-like. But as I argued in the target article (and documented with quotations in **Pylyshyn 1981**), the appeal to “depictive representations” in the literature has largely been a shell game in which picture theorists get all the predictive value out of the literal 2D assumption, and then immediately repudiate that assumption and retreat to a less specific version of the proposal (“it’s how the process accesses it that matters”) which, alas, no longer explains what it was supposed to explain, or at least has no advantage over the null hypothesis.² To avoid misunderstandings about how pictorial the entity hypothesized by **Kosslyn et al.** and others really is (and whether it is colored or 3D), let us call the sort of neural activity map found by **Tootell et al.**, a “Tootell Display” or “Tootell Screen” instead of a “picture.” (I have not attempted a mathematical definition of “picture” because the picture-theorists do not have anything that definite in mind.) **Polimeni & Schwartz’s** definition as a “*locally regular feature map*” will do, since it suggests that the display might be a wildly but continuously distorted version of the proximal stimulus. A more technical way of putting it is that a picture can be any continuous mapping of the proximal stimulus that preserves local topology – such as a homeomorphism or locally affine transformation. (This definition of “picture” allows for the “cortical magnification factor” that **Kosslyn et al.** raise in their commentary.)

While there is neuroscience evidence for a Tootell Display, such a display cannot explain the facts of imagery that are driving the imagery research program (e.g., mental scanning, the effects of different image “size,” mental rotation, and the other lines of research I discussed in the target article) for at least the following two reasons.

1. In a large number of cases the experimental phenomena are not attributable to properties of the mental image at all, but to people’s ability to simulate what they believe they would see if they were to witness certain events (e.g., if they were to scan their attention from place to place on a map). In other words, the facts in question do not constrain brain structures at all: any structures that could recall past

episodes, draw inferences, generate sequences of thoughts, and estimate time intervals such as time-to-collision, would do. The purpose of introducing the “null hypothesis” is not primarily to claim that images must be symbolic, but to show that many phenomena that imagery theorists are concerned with can be accounted for by *any* theory adequate to carry out reasoning, including a symbolic one that implements a “language of thought” (the reason I take the latter to be the “default” form is treated briefly in sect. R4.1 of this response). One reason why proposals, such as **Gottschling’s**, or **Pani’s**, that images are analog would not work in general, is because the phenomena are very often not a result of any properties of the image itself. This is not to say that there are no analog representations of particular magnitudes in the brain – see section R4.3.

2. The Tootell Display proposal is not compatible with what we know about either mental imagery or vision. This was the point of the eight reasons I gave in section 7.2 of the target article: they are just a sampling of reasons why, *even if a detailed 2D pattern of activity mirroring the experienced image were found in the cortex during imagery*, this could not be what is responsible for the empirical phenomena of mental imagery (by which I mean the kind of phenomena that have been extensively documented in Kosslyn 1980; 1994). If activity corresponding to the Tootell Display were found during episodes of mental imagery, this would indeed be a very interesting finding, but contrary to general belief we would not be the least bit closer to explaining the many phenomena of mental imagery than we were before: the Tootell Display is likely to have as small a role in theories of mental imagery as it has in fact had in theories of visual processing. In his commentary, **Dennett** has some characteristically picturesque examples for putting the same point. Dennett’s point is that, when viewed in a certain way, many interesting patterns might be observed in the brain *without those patterns being exploited by the brain* to compute certain particular behavioral phenomena. The question of whether the *patterns* are exploited is also independent of the question of whether certain brain regions are involved (causally or incidentally) in the activity of imaging, so the efficacy of rTMS (repetitive transcranial magnetic stimulation) is not directly relevant to this question (*pace* **Kosslyn et al.** and **de Haan & Aleman**). It is also independent of exactly how the spatial distribution of the display is implemented in detail in the brain: Whether this is done on a cell-by-cell basis or whether it uses vector encoding, as suggested by **Sokolov**, the conceptual issue remains the same.

Ingle underscores some of the inadequacies of the Tootell Display that I had mentioned and provides some additional suggestions. He reminds us that in vision the contents of the Tootell Display change several times each second so its contents do not actually correspond to how we see the world. The important point that is often ignored is that what we experience, in both vision and imagery, is a spatial configuration that was *never* present on the retina or on the Tootell Display, so such a display could not correspond to the contents of either vision or imagery. The obvious way to deal with this devastating criticism is to assume that mental imagery (and visual experience) is associated with some other stable panoramic display, as opposed to the Tootell Display, onto which the contents of the Tootell Display are transferred in registration with eye movements. But, as **Ingle** and others have pointed out, there is no evidence for

that sort of display in the brain (and there is plenty of evidence that such a display is never constructed in the course of saccadic integration; see O’Regan 1992; O’Regan & Noë 2001). In fact, as Ingle also points out, visual information that forms the basis of object recognition is carried by a distinct pathway (the ventral system) and converges with information about location – likely in the coordinates of motor control – that is carried by another pathway (the dorsal system). These two systems do not merge (as they would have to, if they corresponded to the contents of visual or imagery experiences) until perhaps the prefrontal cortex, where there is no evidence for any sort of topographical organization. In their commentary, **Kosslyn et al.** wisely disavow such a stable panoramic display (see point 6 in their section “Imagery and Perception”). Interestingly, Kosslyn (1994) clearly embraced it (see Chs. 4 to 7, especially pp. 85–94) and argued that since such a display is used by vision, it is natural to assume it is also used by imagery.

Van der Velde & de Kamps have also focused on the merging of identity and location information from the “two visual systems” and provide a sketch, using a “blackboard architecture,” of how these two types of information might be coordinated – a major problem in models of visual perception. While the term “blackboard” might suggest a spatial layout, none of the appeals to such an architecture in AI or vision have had that implication. The van der Velde & de Kamps proposal itself only claims that retinotopically organized information is transformed continuously into either the identity of each object or into “location representations related to movements of different body parts.” So far this is compatible with the null hypothesis. But they also suggest that the blackboard may be located in the visual cortex (which might be taken as support for the depictive view). But, as **Ingle** points out, the earliest place where ventral and dorsal visual systems converge is roughly in the prefrontal cortex, where there is no evidence of topographical organization.

R4. The “null hypothesis” and its detractors

R4.1. Why should we assume a symbolic system as the “null hypothesis”?

The purpose behind my “*null hypothesis*” proposal seems to have been widely misunderstood, leading to some arguments that are not germane. **Thomas** questions why I consider the symbolic alternative to be the “null” or default case, and perhaps I did not provide enough discussion of this point. **Millar** says that the null hypothesis is a “formal description of problem solving tasks” that “makes no predictions about how people actually go about solving different types of problems.” The null hypothesis is not a formal description of anything; it is a proposal, laid on the table largely as a foil, about what form our representations might take when they are experienced as visual images. The reason why I treat this particular option as the “default” is because (1) we know something about it (since it includes all the various formal languages and symbolic calculi for which we have a formal semantics to tell us how the meanings of complex structures are composed from the meanings of their parts); and (2) because it meets certain minimal requirements that must be met by any system adequate for the representation of knowledge and for reasoning. In par-

ticular, we know that a recursive system of symbols has properties such as productivity, compositionality, and systematicity, and that these are essential for reasoning and knowledge representation (for a detailed argument on this point, see Fodor & Pylyshyn 1988).

Even though the format used by a circumscribed part of the system need not meet all these requirements, it will still have to face the problem of providing a seamless interface between its form and the form used in reasoning, since both vision and imagery do play a role in reasoning (and since translating between a picture and beliefs is what the entire brain does, hence the ever-present danger of a homunculus regress that **Kosslyn et al.** like to pooh-pooh).³ I take considerations such as these to present a serious challenge to those who advocate some proposal other than the symbolic one. Saying that this proposal “seems odd in the context of evolutionary biology” (as **Millar** does) is a strange way to meet this challenge. It buys into the impression fostered by some connectionists (and also discussed in Fodor & Pylyshyn 1988) that if it does not *look like* a nervous system, it must be biologically implausible.

R4.2. How does a symbol system deal with various particular phenomena?

Many commentators presented interesting imagery findings that they claim would be difficult for a symbolic system to accommodate (e.g., **Arterberry, Craver-Lemley & Reeves, Bartolomeo & Chokron, Burgess, Chatterjee**, and to a lesser extent **Jüttner & Renstchler**). But why do they think that? Is it because some of the findings imply the processing/storage of metrical information? Is it because the operations on the representations are globally “spatial,” such as rotating something or viewing it from another viewpoint? These examples are often presented in the tone of voice, “What do you have to say about *that*?” Well, in most cases I have nothing to say about them except that they are not grounds for favoring the alternative: No picture theory could even begin to address them without a collection of ad hoc stipulations. And with the benefit of such stipulations any theory, including a symbolic calculus, could do equally well. It is worth reiterating that what makes certain operations (e.g., rotation) appear natural for a pictorial representation is that it so readily invites the intentional fallacy that **Gold** discusses: It is easy to forget that what is natural to rotate is the (solid, rigid, physical) thing that is imagined, not the mental representation or its brain encoding. It is definitely not natural to rotate a Tootell Display (or its contents), or to examine it from a different perspective, an operation that **Bartolomeo & Chokron** felt was unnatural in a symbolic system. Some analog operations do make sense (e.g., rotation of a dihedral vertex through a small angle in 3D, as postulated by Marr & Nishihara 1976), and I look forward to the development of a theory that incorporates them, but it is unlikely that it will contain a depictive display (Marr & Nishihara’s SPASAR mechanism computes an analog operation on a symbolic representation).

R4.3. Do we need analogs as well as symbol structures?

Despite the above reasons for putting forth a symbolic form as a default proposal, there may be good and sufficient reasons for rejecting this form of representation for images

and, instead, to adopt something more intuitive, like a depictive form. The point that I argued in detail in the target article (for the devil is in the details) is that the known experimental findings do not provide such reasons.

This is not to suggest that there are no important unresolved issues in the symbolic option. One that may be at the heart of many of the examples commentators have raised is this: Discrete symbol systems do not seem to have a generally satisfactory way of dealing with the representation and manipulation of real-valued magnitudes. They do, of course, have the venerable numeral systems – even systems with arbitrarily expandable precision, such as used in the Dewey decimal system – but these may not be completely satisfactory for such purposes as, say, controlling biomotor systems, or perhaps even for accounting for such phenomena as the “symbolic distance effect” cited by **Petrusic & Baranski**, or various abilities that implicate a metrical form of representation (as, for example, in mathematical reasoning, see Dehaene et al. 1998; Gallistel & Gelman 2000). **Polimeni & Schwartz** make the not unreasonable suggestion that a hybrid analog-symbolic system might be required to model vision and other cognitive functions, although whether they actually do or do not remains an open empirical question.

Notwithstanding the various proposed roles for analogs in mental representations, there are many concerns regarding the proposal that image representations involve an analog component, not the least of which is that there is a question about the interpretation of much of the data that lead investigators to that assumption. The considerations discussed in connection with the tacit knowledge proposal (sects. 3 and 4 of the target article, as well as sect. R6 of this response) is one sort of worry (viz., that these data may reflect something other than the form of the representation); and various problems with the empirical demonstrations, discussed by **Petrusic & Baranski**, is another.

Yet another reason given for favoring analogs comes from the way we experience images and percepts. For example, **Pani** worries about why our experience of both vision and imagery has the character of continuity, and suggests that the imagery system may use “tokens borrowed from the neural mapping of the visual world . . . [so that] the visuo-spatial properties they do contain will lead people to report them as experiences in a mental world with an analog character.” The assumption that some of the same neural structures are involved in both imagery and vision does not entail that we need what Pani calls a “dense mapping” from the world to structures in the brain to account for the experience of dense vistas, any more than that we need an analog system to talk about dense layouts (I just did so in this sentence!). To ascribe density to these structures, on the grounds that we experience dense visual regions, is to make the same mistake as to ascribe color, shape, or size to the neural structures because we experience color, shape, and size in our images; it is the insidious intentional fallacy.

In spite of these concerns, I agree that for many purposes an analog representation system (in which a system of physical magnitudes in the represented domain maps to a different system of physical magnitudes in the representing domain) might be required. **Gottschling** is right to say that despite the indefensible matrix-type functional space, I have not excluded some interpretation of “functional space” that depends on a real analog of space (but see Note

6 of the target article). But *caveat lector*; there is more to the notion of an analog representation than meets the eye. Take, for example, the **Polimeni & Schwartz** argument for the efficiency of representing time by time and space by space, as opposed to using “expensive” symbolic codes. To reach this conclusion, a particular cost-accounting scheme had to be assumed, which could turn out to be unjustified. For instance, Minsky and Papert (1971) discuss an example of what looks at first glance like an efficient analog computing mechanism, but in which the efficiency turns out to be illusory when precision and physical constraints are taken into account.⁴ Furthermore, one should not be fooled into thinking that it is even clear what it is for a system to be analog – it is easy to get into the same difficulties in thinking about analogs as one does in thinking about “functional space,” and for similar reasons (see Block & Fodor 1972; Lewis 1971; Pylyshyn 1984).

R4.4. Formats and representations in different modalities

Intons-Peterson challenges my null hypothesis on the grounds that it lacks criteria for when something is the “same form of representation” as that used for general reasoning. And so it does; this was not intended as an operational definition, but a challenge to theories. When a worked-out proposal is available, the criteria will become clear. For example, if we were to take the predicate calculus as the proposed form of representation for general reasoning, and the “vivid representation” proposed by Levesque and Brachman (1985) as the form for mental imagery, the criterion for being “the same type of representation” would be met, since the latter is just a subset of the former (the vocabulary of nonlogical terms would, of course, be domain specific). In general, sameness can only be assessed relative to a theory, so Intons-Peterson is correct in saying that this is an underspecified proposal. She is also quite right to point out that there are forms of imagery other than visual ones, and that they could provide further evidence for my thesis. The work I am familiar with in other modalities does point to the importance of amodal space, as opposed to visual space, as an organizer of cognitive representation (a point that was also made by **Jüttner & Rentschler** in their commentary). Clearly, this is a direction that merits further research.

Olivetti Belardinelli & Di Matteo make a similar point regarding the advisability of extending imagery research into other modalities. In their review of a rather sparse literature on imagery in different modalities, these commentators found evidence for both modality-specific and amodal components to imagery. Moreover, they found that during imagery in the seven sensory modalities they surveyed, brain areas activated were generally specific to the modality in question, yet in no case were they the primary cortical areas associated with perception in that modality. While these commentators speculate that the amodal component might be propositional, they conclude that the modality-specific component is likely “concrete/analogical.” But this is just the conclusion that I have been suggesting is unwarranted. So long as we do not know how brain activity maps onto modal experience, and do not have evidence that information in different modalities is different in format, as opposed to content, the modal-amodal distinction does not show that we represent the two types of

Response/Pylyshyn: Mental imagery: In search of a theory

information differently; as I suggested, again in the spirit of a null hypothesis foil, they *could* be the same except for their content (i.e., what they are about).

The issue of different modalities raises the question of whether what we call spatial imagery, in the fullest sense of the word, arises in the blind. Both **Chatterjee** and **de Haan & Aleman** argue that totally cortically blind people have imagery in this strong sense, though **Millar** disputes this, at least for the congenitally blind, citing informal examples of congenitally blind people who only report sketchy images in modalities that (judging by her examples) do not implicate spatial imagery. Of course, without the experience of vision, blind people might not use color terms in exactly the same way as sighted people do. But the preponderance of evidence seems to show that the blind (including the congenitally blind) have an extremely well developed sense of space and shape, which they claim to experience in the form of images; and that they also show the same effects of scanning, mental rotation, and other signature imagery phenomena as do sighted people (though **Millar** claims that they are not as efficient at it). Moreover, blind children learn to use spatial terms to refer to space in very nearly the same way as sighted children do (Landau & Gleitman 1985). This suggests that sight is not necessary for most of the spatial phenomena that have been studied by imagery researchers, and therefore neither is the vision-specific Tootell Display (unless we are to assume that audition, olfaction, and the tactile and gustatory senses are also represented in a two-dimensional depictive display). Certainly **Jüttner & Rentschler**’s work shows that a cross-modal representation of object shapes is the rule rather than the exception, and hence that images of objects are amodal (or at least not solely visual), thus casting further doubt on the inherently visual depictive theory of imagery.

Jüttner & Rentschler and others, feel that the null hypothesis is also committed to an amodal conception of perceptual representation. But one of the principal differences among modalities is that they are about different things and require different concepts. For example, vision concerns visual properties like color or luminance, audition concerns auditory properties such as pitch or loudness, the tactile sense concerns properties such as smooth or sharp, and so on. Each modality has its own set of modality-specific concepts. Beyond those is a large array of modality-free concepts such as those referring to spatial properties, which apply to most modalities. A symbol system can prima facie accommodate such facts just the way that language does, by simply using different terms when dealing with different modalities. The point is that a modality is, at least to a first approximation, about content and not about format, so there is no need to hypothesize a different form of representation for each modality. Whether this makes symbolic systems amodal or not, is therefore mostly a matter of terminology.

R4.5. Are there other options besides pictures and symbols?

Commentators have offered a number of ideas for a “third option,” although the options that have been proposed are either too sketchy to judge, or else are problematic. For example, an alternative that **Chatterjee** has proposed is the *schema*, favored by many others as well. For **Chatterjee** this

is an abstraction that combines some of the features of both pictures and language. It is, in fact, not too far from Kosslyn's occasional description of "depictive displays" as containing "predigested information" or being "annotated." One version of the *schema* proposal that has been popular in artificial intelligence is sometimes called a *frame* (see, e.g., Minsky 1975). It is a structure that contains slots that are eventually bound to particulars as information becomes available. It also contains default assignments that prevail if no information to the contrary is available, and it contains procedures for transforming the structure or for determining what values go in the slots. That sort of schema is clearly a symbolic form of representation. What Chatterjee has in mind, however, and what many people would like to see, is a form of representation that not only has the properties of frames, but also has intrinsic spatial properties. But such a system would still require a literal spatial display on which to locate the descriptions or annotations, so the locations of the slots or descriptors would have to follow Euclidean principles. The only way out of this dilemma may be the one I have taken in my own theoretical work, in which I provide a mechanism called a visual index (or FINST) which provides a limited means for a descriptive representation to point to, or reference, objects in the world, thereby inheriting spatial properties from the world without actually *having* such properties itself (see below).

Thomas is one of the commentators who agrees that picture theory has all the problems I attribute to it, yet he is not favorably disposed towards the null hypothesis either. Rather, he advocates another approach that he calls *Perceptual activity theory*, which holds that "imagery arises from vicarious exercise of . . . [mastery of relevant sensorimotor contingencies]: a sort of play-acting of perceptual exploration." The version of activity theory developed by O'Regan & Noë (2001), which Thomas cites with approval, has many clear advantages over picture theory, though it also has some problems of its own (many of which it inherits from J. J. Gibson's direct realism, which minimizes the role of representations of any kind; see my commentary, Pylyshyn 2001b in the same issue of *BBS*).

Having said that, I should point out that many of the ideas I have proposed, in this response and elsewhere (see Pylyshyn 2001c), are very much in the spirit of perceptual action theory. For example, in the next section I suggest that many of the apparent spatial and directional properties of images could derive from real space, providing we have a mechanism for associating features or objects in images with corresponding objects in real space. This view has been developed in connection with a theory of visual indexes, which provides a mechanism for preconceptual links to objects in the world (the theory, alluded to only briefly in this response, is laid out in detail in Pylyshyn 1998; 2000; 2001a; 2001b; Pylyshyn, forthcoming). If we assume that the spatial quality of perceived space derives from the way we interact with, or are potentially able to interact with, objects in real space, we can explain why images appear to be spatial in the same way, so long as the spatial locations of objects in the images are bound to locations of objects in the world (as postulated in visual index theory). So in fact my views do not diverge so radically from **Thomas'** insofar as his views are sufficiently specified to allow comparison. But I continue to view symbolic forms of representation as the default or "null" hypothesis for reasons sketched in section R4.1.

R5. Where does the spatiality of images come from, if not the display?

R5.1. Space, motion, and motor control in visual imagery

Nijhawan & Khurana argue that perhaps the sort of static images I criticized are not strongly spatial because spatial maps are constructed from motion information and so are best accessed through motion stimuli. The idea that motion may be important in revealing the spatial character of representations is an interesting one. But, in fact, many of the experimental phenomena used to demonstrate the spatial nature of imagery do involve motion (e.g., mental scanning, mental rotation, mental paper folding, and other dynamic operations performed on images). Nijhawan & Khurana ask whether I would agree that images are spatial if smooth motion in an image could be demonstrated. But this question already assumes a reification of mental space. What if imagined continuous motion did not literally involve movement through some space by the continuous change of position? I have argued that imagining a greater distance does not involve a greater amount of some neural magnitude, such as cortical distance. Imagining that something is moving, and imagining that something is further away, could be nothing more than entertaining the thought that some particular things are moving or are further away. There is nothing strange about the idea of distinguishing motion from change in position in some space. Patients with cerebral akinetopsia or motion blindness (Zeki 1991) are able to see *that* objects change location over time, without seeing them *as* moving. Conversely, imagining something as moving could also occur without anything changing in real time (I once questioned why we assume that imagined time maps onto real time, but the question just seemed to puzzle people; Pylyshyn 1979a). In addition, of course, it could involve imagining something as being at a sequence of locations (as when observers simulate the experience of seeing something move, by thinking of it as being at a series of locations). But where are the locations, if not on an image display? My contention, sketched out in sections R4.5, R5.2, and R5.3 of this response, is that the locations in this case could be places occupied by certain objects in the world, perceived visually or in some other modality, or places that are recalled in terms of their relation to some currently perceived objects. (For instance, in the target article I described an experiment that showed that we were better at imagining uniform motion if we had a series of visible places to think of the object as being at, while we imagined it to be moving.) The point is that in imagining motion, the space through which an object passes need not be in the head.

Raab & Boschker also emphasize the dynamic quality of the perceptual act, although from a different point of view; one that focuses on the active exploration of visual information (this approach is consistent with that advocated by **Thomas**, and is one that has gained favor in computer vision). Whatever the merits of the approach, it has led researchers (e.g., O'Regan & Noë 2001) to focus on the information available in the world, instead of on its projection on a cortical display, which is clearly a step in the right direction. From this perspective, Raab & Boschker recommend that we study imagery by first studying dynamic and motor imagery. I am glad to endorse this recommendation, and indeed have already endorsed the idea that some parts

of the proprioceptive/kinesthetic/motor system may be used to provide the spatial properties usually attributed to images themselves. But the same intentional fallacy that pervades visual imagery theory also threatens theories of motor imagery; it is the illusion that external properties of the world are not only represented, but are *duplicated* in the brain. For example, the idea that we plan the execution of motor actions by carrying them out imaginally and observing what happens, presupposes that when we imagine actions, their consequence is automatically made available by virtue of some property of the imagery mechanism, without the involvement of knowledge, inference, and thought. This is exactly like the case of visual imagery, where it is typically assumed that when we imagine some process (like mental scanning), all we have to do is wait and “see” what happens without having to draw inferences – because the process by which the image unfolds is determined by the inherent properties of the depictive display.

The role that motor mechanisms play in the transformation of visual images is even less clear, notwithstanding evidence of correlations between visual image transformations and activity in parts of the motor system (Cohen et al. 1996; Richter et al. 2000), or the influence of real motor actions on visual transformations (Wexler et al. 1998). Some of the cortical activity observed during both motor performance and the mental transformation of visual images, may reflect the fact that these areas (e.g., posterior parietal cortex) compute higher-level functions required for extrapolating trajectories, for tracking, for planning, and for visuo-motor coordination (Anderson et al. 1997). Since many of these functions also have to be computed in the course of anticipating movements visually, it is reasonable that the same areas might be active in both cases. Although studying the interaction of imagery and the motor system is clearly important, at the present time we are far from justified in concluding that dynamic visual imagery is carried out *by means of* the motor system (or that visual operations *exploit* motor control mechanisms). This way of speaking suggests that our motor system can grasp and manipulate our images (a view that is quite compatible with the mental reification of the world that we find in much mental imagery theorizing).

R5.2. Superimposing images on the visual world: Attention and inhibition

In the target article I suggested that certain things are special about the case where images are combined with visual perception (although the same thing probably applies when images are combined with other perceptual modalities that concern space). The example discussed by **Gosselin & Schyns** falls into this category, and the comments in section 5.3 of the target article apply to it. For over thirty years, there has been evidence that observers are able to vary both their response bias (β) and sensitivity (d') at attended regions of a display (Bonnell et al. 1987; Downing 1988; Farah 1989; Mueller & Findlay 1987; Segal & Fusella 1969; 1970). Focusing attention on a region with a certain simple shape (e.g., an S-shape) while looking at a display of uniform noise, could thus lead to the perception of a figure in the shape of the region being attended: the noise in the attended region would simply be enhanced (either amplified in amplitude or raised by a pedestal function, depending

on whether the effect was due to a change of criterion or of sensitivity), creating a real perceptual effect at that region.

Gosselin & Schyns claim that since there was no information in the signal, the relevant information must have come from “an internally generated signal – that is, a mental image.” I agree with the first part of that claim but see nothing to be gained by calling a distribution of attention an “image.” Contrary to the way that Gosselin & Schyns put it, there is no reason to assume that observers had to have “knowledge of *all* the pictorial characteristics of an ‘S’” and even less reason to claim that what they have is “functionally isomorphic to an actual image of an ‘S’.” All they need is a description of the ‘S’ together with the ability to think the following sort of demonstrative thoughts while viewing certain regions in the display: “*this* region is where R_i would fall,” where R_i is the part of the overall representation of the shape that encodes a certain discrete part i of that shape (it might, for example, be a code for the top upward-concave semicircle of the ‘S’). Such a mechanism would allow observers to allocate attention to the appropriate region, and in so doing would enhance the sensitivity in that region without having to image the figure. Farah (1989) showed that instructions to *attend to* a region were even more effective in sensitizing a region than instructions to imagine the relevant-shaped region. Grossberg’s LAMINART model even provides a neural mechanism for how this can occur (by shifting the balance between excitation and inhibition on relevant neural circuits).

This view of what is going on is radically different from the one proposed by **Gosselin & Schyns**. For one thing, according to this view there is no such thing as space *in the image* (as required by the depictive theory), there is only the space where the observer is looking. To put this another way, the Gosselin & Schyns finding (and their sophisticated power-spectrum analysis) is compatible with the representation of the ‘S’ shape being symbolic, and thus without the need to assume a spatial display in the head at all. For example, it might be encoded as symbols denoting different *geons* (Biederman 1987), together with some symbolic form of location code used to direct visual attention to contiguous sets of pixels in the real display, thus leaving all spatial properties where they belong – outside the head. It would be more interesting if it could be shown that observers “project” more than a selected region onto a surface. For example, if they could project other visual properties such as color, shading, texture, depth, and so on, which had observable visual consequences that could not plausibly be attributed to the effect of attention.

The **Arterberry et al.** commentary also involves phenomena that arise when an image is “superimposed” over a viewed display. The demonstration of dissimilarities between vision and imagery discussed by Arterberry et al. is very interesting. Their account, in terms of the suppression of visual signals that compete with projected imagery, is particularly intriguing and has far-reaching ramifications for understanding the involvement of visual cortex in imagery. **Gottesmann** also reports evidence that activity in primary visual cortex, as monitored by EEG, is suppressed during vivid dreaming. **Meyer** points out that single-cell recordings also do not provide evidence for the activation of visual cortex in imagery, and **Sokolov** hypothesizes that visual information may be located in the temporoparietal area. Such findings suggest that visual imagery experiences

may not be associated with activity in the visual cortex (as Crick & Koch [1995] had already recognized).

Indeed, **Gottesmann's** finding that there is actual suppression of activity in the visual cortex during certain visual experiences, raises the intriguing possibility that the increased blood flow observed by neuroimaging techniques may not reflect increased information-processing activity (a possibility previously suggested by Fidelman 1994) but might perhaps reflect the active *suppression* of visual information processing. This would make sense in those experiments that involve examining an image during fMRI or PET scanning when there is some visual input, or even visual persistence of the sort that **Ingle** describes, or that Ishai and Sagi (1995) report, which would have to be ignored in carrying out the primary task (and this could be true even if the eyes were closed). Intriguing as it is, such speculation will clearly need to be submitted to careful empirical scrutiny.

R5.3. Visual neglect and spatial orientation in imagery

One often hears that reports of left hemifield neglect in both vision and mental imagery supports the view that both involve an image on a cortical display, since, if one side of the Tootell Display is damaged, the same deficit might be expected in both vision and imagery. But the idea that what is damaged in visual neglect is one side of a display, seems too simplistic⁵; it does not account for the dissociation between visual and imaginal neglect (Coslett 1997), for the amodal nature of neglect (the deficit shows up in audition as well as vision; see Marshall 2001; Pavani et al. 2002), for the fact that “neglected” stimuli typically provide some implicit information (Driver & Vuilleumier 2001; McGlinchey-Berroth et al. 1996; Schweinberger & Stief 2001), for the characteristic response bias factors in neglect (Bisiach et al. 1998; Vuilleumier & Rafal 1999), and for the fact that higher-level strategic factors appear to play a central role in the neglect syndrome (Behrmann & Tipper 1999b; Bisiach et al. 1998; Landis 2000; Rode et al. 2001). The “damaged display” view also does not account for the large number of cases of object-centred neglect (Behrmann & Tipper 1999a; Tipper & Behrmann 1996). Moreover, as Bartolomeo and Chokron (2002) have documented (and reiterate in their commentary), the primary deficit in neglect is best viewed as the failure of stimuli on the neglect side to attract attention.

I agree with **Bartolomeo & Chokron**, as well as with **Burgess** and with **Chatterjee**, that it would be odd for a symbolic encoding system by itself to have directional preferences, such as found in neglect, and I also agree that most cases of imaginal neglect are unlikely to be due to tacit knowledge. Having granted that, one must then ask why we should expect the explanation for such directional properties to be found in the format of representations or in the medium of the Tootell Display. Deficits such as neglect, whether in vision or in imagery, represent a failure to orient to one side or the other, and the direction may have more to do with direction in the world, than direction in an image. As in the case of the **Gosselin & Schyns** example discussed earlier, orienting is a world-directed response. There is considerable merit in Bartolomeo & Chokron's suggestion that perhaps “visual imagery involves some of the attentional-exploratory mechanisms that are employed in visual behavior . . . [so] the ‘perceptual’ aspects of visual mental images might thus result not from the construction

of putative ‘quasi-perceptual’ representations, but from the engagement of attentional and intentional aspects of perception in imaginal activity.” In other words, when attending to the left side of an image, patients are actually *orienting towards* the left side of the perceived world (or perhaps of their body). Even with eyes closed we have accurate recall, at least for a short time, of the location of things in the world immediately around us (see the remarks about this by **Ingle**), and it may be in relation to these world-locations that attention orients. As I speculated in section 5.3 of the target article, it may be generally the case that it is the physical space outside the head that gives imagery its putative spatial character and that it does so by virtue of how mental contents are associated with (or bound to) places in the perceived world. This interpretation is given further support by reports, mentioned by Bartolomeo & Chokron, that imaginal neglect can be modulated by peripheral manipulations, such as turning the head.

Although the case is clearest when the spatial layout is visually perceived while imagining, since in that case aspects of what is imagined can be associated with places in the perceived layout through visual indexes, there is no reason why this should not also hold when real space is sensed through other modalities, such as proprioceptive or kinesthetic modalities. Indeed, motor-space analogs of visual indexes, called Anchors, were proposed when the visual index theory was first introduced in Pylyshyn (1989). It is known that people are very good at orienting to stimuli that are not visually present (Attneave & Farrar 1977). The ability to bind objects of thought to the location of perceived (or recalled) external objects allows us to orient to them, thereby enhancing the illusion that things are laid out inside the head the way that the corresponding things are laid out outside the head, thus reinforcing the intentional fallacy.

R6. Tacit knowledge and cognitive penetrability again

In the target article and elsewhere, I have been at pains to point out that tacit knowledge does not explain all imagery phenomena (it does not, for example, explain all aspects of mental rotation or of the crowding effect in mental scanning or the oblique effect). But if it does explain some things (e.g., the scanning effect, the image size effect), then the extra apparatus of a depictive display is redundant because it plays no role in the explanation, however much it might give comfort to one's preconceived ideas. It's not that the postulated structure is necessarily false, but it is simply irrelevant to the data at hand. In these cases, if there are certain activity patterns on the Tootell Display, they are not the reason that you get the scanning effect, the size effect, or even the phenomenology of mental imagery. That is why it does not help to say that tacit knowledge may be encoded in the form of a depictive representation (as **Kosslyn et al.** suggest). Kosslyn et al. are correct to note that I use tacit knowledge to talk about content rather than form, so that if the phenomena can be explained by appeal to tacit knowledge, then assuming that such knowledge is encoded in a depictive manner is gratuitous. It could be encoded in protein molecules so far as these data are concerned (indeed, there is evidence in favor of such an idea as regards spatial memory, see Blum et al. 1999), because these data do not in any way constrain the format of the representation.

So why assume a depictive representation as a way of encoding tacit knowledge? The answer is surely that **Kosslyn et al.** believe that it is the spatial format, and not the content, that explains the data; but in this they are simply wrong for those cases (such as scanning) where tacit knowledge provides a better explanation. The demonstration that they are wrong consists in showing such things as that, if observers believed that in viewing a map it would take them longer to switch from viewing point A to viewing point B when A and B were, say, on opposite sides of the river on the map, then they *will* take longer when they do it from their image. If you need convincing, just imagine a map and imagine that switching your point of regard takes longer when the two imagined places are across the river (you could do that by thinking of them as moving more slowly through water or as swimming across *or just taking longer to get there without moving at all*). Go ahead, do the experiment: *it is your image, so it can have any property you want it to have!* I'm not saying that people do believe that it will take longer if the fixations are on either side of the river. But if they did believe it, for whatever reason that may strike them, then it is obvious that this is the way it would happen in their image. You might wonder whether this is only true of dynamical processes like scanning as opposed to basic geometrical properties. Can your image, for example, be non-Euclidean? This is where tacit knowledge is obviously relevant. In order to imagine such a thing as a non-Euclidean (or four-dimensional) space, you would have to have certain relevant knowledge; you would have to know what moving through a non-Euclidean space would *look like*, in the sense that you would need to know, for example, how shapes would change as you moved through this space. (Contrary to Kosslyn et al.'s claim, I don't say that you need to have *seen* things for you to be able to imagine them, but you need to know what certain aspects of them would look like, in the same sense that the patients that **Goldenberg** writes about in his commentary do not know what certain things would look like and consequently cannot image them.)

The main point about cognitive penetrability is not just that you can influence your image (as **Kosslyn et al.** assume), but that your image has no properties other than those you take it to have – which is precisely **Dalla Barba, Rosenthal & Visetti's** point about there being no surprises in your image. Anyone who does not believe that their image will do more-or-less what they will it to do is allowing scientific ideology to override common sense. Your image will even *look like* what you make it look like, usually how you believe that something does look (but not necessarily, since you can make your image of something look different from the way you believe it actually looks!). Of course many factors determine what you *do* make it look like, whether or not you can express (in language or in drawings) what something looks like and the conditions under which you can recall what something looks like. Studies showing that people cannot predict the results of imagery experiments (as in the Denis & Carfantan [1985] study that **Kosslyn et al.** cite) are beside the point, all you need is some idea, however vague or implicit or ineffable, about what it would be like if you were to see the thing you are supposed to imagine. Adopting the image mode of reasoning (i.e., focusing on the appearance of things, whatever that means in terms of a theory of imagery) may well alter the likelihood that you will think of or recall something or the other, just as being in a

certain place affects what you recall. But no conclusion can be drawn about the nature of images from such properties of imagery, as these are properties of memory and thought in general.

Goldenberg appears to agree with most of what I claim about certain phenomena of imagery being a result of knowledge, but he insists that, in contrast to the knowledge that functions in recognition, the knowledge involved in imagery is not tacit but explicit. Here, much depends on what you mean by "tacit." In very many cases the relevant knowledge is indeed explicit knowledge of how things look, inasmuch as that knowledge is available for answering questions. But it also seems to me that the knowledge that determines such properties of imagery as "the visual angle of the mind's eye" (Kosslyn 1978), are not available if one simply asks the subject, which is why I called it tacit; yet it can be rationally altered by providing the right experience or information, and it can be revealed in a variety of ways, which is why I call it knowledge. For Goldenberg the relevant distinction is between knowledge that can have a general effect in cognition (which he calls explicit), and knowledge that is part of the modular visual system and is only used in recognition (which he calls tacit). While that is a distinction clearly worth preserving, it is not the one I had in mind in appealing to tacit knowledge, so the apparent disagreement may well be merely terminological.

R7. Second-order and "structural" isomorphism

I do not claim, as **Amiri & Marsolek** suppose, that representations must be first-order isomorphic in order to be explanatory. What I said is that picture-theorists claim that images are spatially isomorphic (or homeomorphic) to a picture of what they depict, but that in order for this sort of isomorphism to explain typical imagery phenomena, the representation would have to be literally spatial rather than "functionally spatial." But second-order isomorphism of the sort that Shepard studied, though requiring pictures-in-the-head, is not sufficiently constraining: It is simply functional isomorphism and is compatible with a descriptivist position (see Pylyshyn 1984, Ch. 9). A system of representations that is second-order isomorphic to some domain is just a system of representations that allows a similarity measure among represented objects to be computed (as in the original similarity judgment study of Shepard & Chipman 1970). There is no reason why such similarity judgments could not be based on inferences drawn from symbolic representations. Second-order isomorphism by itself places no constraints on the form of the representation – it could be depictive or symbolic or anything else. One would need to at least know *why* the second order isomorphism held (what mechanisms were responsible), in order to infer the form of the representation itself.

The Edelman (1998) paper that **Amiri & Marsolek** cite, presents a mathematical analysis of the requirements that should be met by an adequate system of representation, which *inter alia* include second-order isomorphism. The further step that Amiri & Marsolek take of suggesting that it is the cognitive architecture, rather than the content of representations (and inferences drawn from them), that must be responsible for the second-order isomorphism is an interesting and substantive proposal. Given the modularity of vision, there are significant belief-independent

constraints built in to the early vision system, especially in computing an object's appearance. Nonetheless, second-order similarity of the sort that was demonstrated by Shepard and Chipman (1970) is in general cognitively penetrable (and readily altered by attentional strategies, as Shepard 1964 showed) and therefore unlikely to be part of the architecture.

I'm not sure whether **Wright** has something like second-order isomorphism in mind when he says that visual experiences are only "structurally isomorphic" to sensory inputs. But he makes a very different point when he goes on to claim that, although there are no pictures in the brain, there is "non-epistemic" storage, which he calls "inner registration," of sensory events. I assume that he equates these to uninterpreted sensory images. There is a great deal of empirical data showing that sensory records are not kept, but even if they were, mental images are certainly not like records of such sensory events. I argued in the target article that while some image reinterpretation may occur, this sort of reinterpretation is arguably not "visual," nor is the record that is reinterpreted a record of "non-epistemic" visual sensations or of their "structurally isomorphic" internal responses. While it is known that fairly rich sensory storage may be available for short times (Ishai & Sagi 1995; Sperling 1960), those sorts of iconic stores are very different from images constructed from memory. Wright's "bell-I-mud-dum" example involves reparsing a phonetic string in short-term memory (where it was arguably already "epistemic" inasmuch as it was likely encoded in terms of phones), but I am not aware of any evidence of such reparsings occurring from an image constructed from long-term memory.

R8. Other topics: Visual expectations, thoughts, and phenomenology

R8.1. Visual expectations or visual images?

Zaidi & Griffiths provide some interesting demonstrations of visual illusions that appear when a mental rotation is carried out on a perceived figure. From these they conclude, quite reasonably, that it is the assumptions that viewers make that results in their expecting the rotated figure to look different from what it actually does look like when it is physically rotated. Zaidi & Griffiths conclude that "active visual imagery is an integral part of active visual perception." I would not have put it that way since visual expectations are hardly equivalent to the sorts of images that we experience or that picture-theorists postulate. Visual expectations need not take the form of a projected picture, as opposed to some general prediction as to which elements should be where. In discussing the **Gosselin & Schyns** commentary in section R5.2, I have argued that a visual expectation, even one that involves detailed shapes and locations, does not need to be more than a spatial distributions of attention over a real scene, which is very different from a picture or a pattern of activity on a Tootell Display.

R8.2. Distinguishing images from thoughts

Niall is right that if images are propositional, then their propositional content is insufficient to demonstrate a simple Euclidean theorem. But who is so naïve as to think that

the visual system "sees" the entailments of Euclid's axioms in a figure (perhaps a person who believes we can think in pictures?). As Niall's example shows, we are deceived if we think we have represented all the diagram's geometrical properties in our image. But the lesson to draw from this is *not* that mental images are "non-epistemic" (to use **Wright's** term), or that they do not constitute knowledge. Diagrams in the world and on the retina are non-epistemic, but mental *representations* of them are epistemic; they constitute beliefs about how things look, which is why we can think about them and also why we can be mistaken about their true shape. They are, moreover, too impoverished to permit proofs of Euclid's First Proposition without additional non-diagrammatic representations (i.e., thoughts – which, contra empiricism, do not derive from sensations).

R8.3. The role of phenomenology

Dalla Barba et al. raise an interesting point concerning the role of phenomenology in the enterprise of understanding mental imagery. They say that in studying imagery, phenomenology is of the essence and it does support a picture theory of imagery because that is how we experience imagery. They assert that I am unfairly maligning phenomenology "for what it never pretended to be" and that "phenomenology has never aimed at causal explanation." This may well be the case, although my target was not phenomenology, but precisely the attribution of causal power to the experience itself (which is done implicitly and nearly universally). I find myself agreeing with much of the Dalla Barba et al. commentary (e.g., concerning the content of images and its qualitative difference from vision), which suggests that perhaps the proper use of phenomenological evidence may be a useful tool, although psychology has been justifiably suspicious of introspection since the failure of the method to deliver scientifically useful results at the turn of the (last) century. The issue of a phenomenological homunculus, raised by Dalla Barba et al., reduces to the intractable mind-body (or experience-experencer) problem, which is beyond the scope of the present article, not to mention the present author (but see Dennett 1991). In any case, the *phenomenological* homunculus (the experience of being a viewer of one's image) is irrelevant to a causal theory, as the authors admit at the outset.

R9. Conclusion

What is so unappealing about the current direction in the study of mental imagery is that it cannot seem to avoid what Pessoa et al. (1998) call "analytical isomorphism" – the assumption that what one will find in the brain is what appears in one's conscious experience. I recommend the following heuristic: If you feel yourself drawn by some body of data to the view that what is in your head is a smaller and perhaps less detailed version of what is in the world, then you had better stop and reconsider your underlying assumptions. While many readers were not persuaded by what I called the null hypothesis, it does appear that there has been a move away from naïve picture theory in several areas of imagery research. Many people are now objecting to the purely symbolic view by considering other options, rather than by insisting that it is obvious that imagery must exploit some sort of spatial display. Others are concentrat-

ing on studying the parallel mechanisms of vision and imagery, while rejecting the implication that this means there must be a picture-like object for vision to exploit. This is a conceptually difficult problem and the arguments will no doubt continue (despite the belief held by many writers that the debate has already been resolved by evidence from neuroscience). One can always hope that the next time around we may approach the question with a better appreciation of the general conditions that have to be met by an adequate theory. On the other hand, as **Slezak** intimates, we may be condemned, like Sisyphus, to repeat the task of correcting the intentional fallacy without end, creating employment for future generations of cognitive scientists and philosophers.

ACKNOWLEDGMENT

Work on this paper was supported by the National Institutes of Health Research Grant 1R01-MH60924. Send reprint requests to the author at Rutgers Center for Cognitive Science, Center for Cognitive Science, Psychology Bldg. Addition, Busch Campus, Rutgers University, Piscataway, NJ 08854–8020. Author email: zenon@rucss.rutgers.edu

NOTES

1. I responded to Anderson's version of the indeterminism thesis in Pylyshyn 1979c, and have written extensively on the notion of "strong equivalence" in cognitive science (e.g., Pylyshyn 1984), showing that mere input-output equivalence is not what cognitive scientists aim for, even without the benefit of neuroscience data.

2. **Kosslyn et al.** claim, "The depictive theory . . . presents a coherent, internally consistent view of how mental images may be processed." But so long as the coherence and predictive power come not from intrinsic properties of the "depictive" form of representation itself, but from a variety of ancillary assumptions about how the representation must be used and what restrictions are placed on accessing information from it, the depictive theory is coherent only in that it fits one's preconceptions, and its predictive power derives entirely from the independent constraints which any theory could adopt. Think of the added epicycles of "annotations," or the "predigested information," or the requirement that to get from A to B one has to pass through places that are "in between," even though they are so only by stipulation, or the appeal to the intuition that smaller images must be harder to "see," and so on, all of which are assumed for no "internally consistent" reason except to fit the data at hand, however they turn out. Was it really the depictive form of images that predicted the oblique effect? Once laid bare, the depictive theory is no less a patchwork than any other theory for explaining the experimental phenomena of mental imagery – which are unlikely to have a single cause in any case.

3. One hears over and over that since the process has been implemented on a computer, it shows that a homunculus is not needed ("any more than there is a need for a homunculus in visual perception" according to **Kosslyn et al.**). But if vision proceeded by examining the panoramic display that we experience, we *would* need a homunculus (as **Dalla Barba et al.** correctly point out). Moreover, what has been implemented on a computer is but a trivial fragment of this process. This fragment, with encouragement from the names given to various components (e.g., "visual buffer," "attention window") and operations ("look for," "generate image," "determine whether resolution is sufficient," "zoom," and so on), invites the assumption that the basic idea can be extended to model all of imagery without the intervention of an intelligent agent. Therein lies the slip "twixt cup and lip" where the homunculus lurks, barely hidden. The problem is that the homunculus has not been "discharged" (as Dennett 1978 would put it) until the "intentional loan" incurred in using these names is paid up.

4. The example Minsky and Papert (1971) describe is a method for finding the cheapest path between two cities connected by many roads through other cities. Represent each city by a ring and the cost of going between each pair of cities by the length of a string tied to the two rings. Then simply grasp the rings representing the starting and ending cities and pull tight. The solution is the set of strings that stretch straight across and it appears to be found instantly. But because of physical constraints that grow nonlinearly with the number of cities, the appearance is illusory, as Minsky and Papert show.

5. Kosslyn (1994) does not explicitly claim that the "depictive display" is damaged in cases of neglect, preferring instead to speak of the parallels between the vision and imagery systems. But to be consistent he *should* claim that the display is damaged, since the point of the display is that it allows one to explain spatial properties of imagery by appealing to spatial properties of the display. Simply saying that it shows that vision and imagery use the same mechanisms does not confer any advantage to the depictive theory, since *any* theoretical imagery format can claim that (including the null hypothesis, which is why it is there: to provide a test for the irrelevance of assumptions about the image format).

References

[Note: The letters "a" and "r" before author's initials stand for target article and response references, respectively]

- Adams, J. E. & Rutkin, B. B. (1970) Visual responses to subcortical stimulation in the visual and limbic system. *Confinia Neurological* 32:156–64. [GEM]
- Ahumada, A. J. & Lovell, J. (1971) Stimulus features in signal detection. *Journal of the Acoustical Society of America* 49:1751–56. [FG]
- Aleman, A., Rutten, G.-J., Sitskoorn, M., Dautzenberg, G. & Ramsey, N. F. (2001a) Activation of striate cortex in the absence of visual stimulation: An fMRI study of synesthesia. *NeuroReport* 12:2827–30. [EdH]
- Aleman, A., Schutter, D. L. J. G., Ramsey, N. F., van Honk, J., Kessels, R. P. C., Hoogduin, J. H., Postma, A., Kahn, R. S. & De Haan, E. H. F. (2002) Functional neuroanatomy of top-down visuospatial processing in the human brain: Evidence from rTMS. *Cognitive Brain Research* 14:300–302. [EdH]
- Aleman, A., Van Lee, L., Mantione, M., Verkoijen, I. & De Haan, E. H. D. (2001b) Visual imagery without visual experience: Evidence from congenitally totally blind people. *NeuroReport* 12:2601–604. [EdH]
- Algom, D. (1992) Memory psychophysics: An examination of its perceptual and cognitive prospects. In: *Psychophysical approaches to cognition*, ed. D. Algom. Elsevier Science. [WMP]
- Algom, D. & Lubel, S. (1994) Psychophysics in the field: Perception and memory for labor pain. *Perception and Psychophysics* 55:133–41. [WMP]
- Andersen, R. A., Bracewell, R. M., Barash, S., Gnadt, J. W. & Forassi, L. (1990) Eye position effects on visual, memory and saccade-related activity in anterior LIP and 7A of macaque. *Journal of Neuroscience* 10:1176–96. [DI]
- Andersen, R. A., Essick, G. K. & Siegel, R. M. (1985) Encoding of spatial location by posterior parietal neurons. *Science* 230(4724):456–58. [NB]
- Andersen, R. A., Snyder, L. H., Bradley, D. C. & Xing, J. (1997) Multimodal representation of space in the posterior cortex and its use in planning movements. *Annual Review of Neuroscience* 29:303–30. [rZWP]
- Anderson, J. R. (1978) Arguments concerning representations for mental imagery. *Psychological Review* 85:249–77. [SMK, rZWP]
- Armstrong, J. (1996) *Looking at pictures – an introduction to the appreciation of art*. Gerald Duckworth. [GG]
- Arnauld, A. (1683/1990) *On true and false ideas*. Translated with an introductory essay by Stephen Gaukroger. Manchester University Press. [PPS]
- Ashcraft, M. H. (2002) *Cognition (3rd edition)*. Prentice Hall. [WMP]
- Attneave, F. & Farrar, P. (1977) The visual world behind the head. *American Journal of Psychology* 90(4):549–63. [DI, rZWP]
- Avant, L. L. (1965) Vision in the ganzfeld. *Psychological Bulletin* 64:246–58. [aZWP]
- Balasubramaniam, M., Polimeni, J. & Schwartz, E. L. (2002) The V1-V2-V3 complex: Quasi conformal dipole maps in primate striate and extra-striate cortex. *Neural Networks* 15(10). (in press). [JP]
- Banati, R. B., Goerres, G. W., Tjoa, C., Aggleton, J. P. & Grasby, P. (2000) The functional anatomy of visual-tactile integration in man: A study using positron emission tomography. *Neuropsychologia* 38:115–24. [MOB]
- Banks, W. P. (1977) Encoding and processing of symbolic information in

- comparative judgments. In: *The psychology of learning and motivation, vol. 11*, ed. G. H. Bower. Academic Press. [WMP]
- (1981) Assessing relations between imagery and perception. *Journal of Experimental Psychology: Human Perception and Performance* 7:844–47. [aZWP]
- Banks, W. P. & Flora, J. (1977) Semantic and perceptual processes in symbolic comparisons. *Journal of Experimental Psychology: Human Perception and Performance* 3:278–90. [WMP]
- Banks, W. P., Mermelstein, R. & Yu, H. K. (1982) Discriminations among perceptual and symbolic stimuli. *Memory and Cognition* 10:265–78. [WMP]
- Baranski, J. V. & Petrusic, W. M. (1992) The discriminability of remembered magnitudes. *Memory and Cognition* 22:254–70. [WMP]
- Barolo, E., Masini, R. & Antonietti, A. (1990) Mental rotation of solid objects and problem-solving in sighted and blind subjects. *Journal of Mental Imagery* 14(3–4):65–74. [aZWP]
- Barsalou, L. (1999) Perceptual symbol systems. *Behavioral and Brain Sciences* 22(4): 577–660. [aZWP]
- Bartolomeo, P. (2002) The relationship between visual perception and visual mental imagery: A reappraisal of the neuropsychological evidence. *Cortex* 38(3):357–78. [PB]
- Bartolomeo, P., Bachoud-Lévi, A. C., Chokron, S. & Degos, J. D. (2002) Visually- and motor-based knowledge of letters: Evidence from a purely alexic patient. *Neuropsychologia* 40(8):1363–71. [PB]
- Bartolomeo, P., Bachoud-Levi, A. C. & Denes, G. (1997) Preserved imagery for colours in a patient with cerebral achromatopsia. *Cortex* 33(2): 369–78. [aZWP]
- Bartolomeo, P. & Chokron, S. (2001) Levels of impairment in unilateral neglect. In: *Handbook of neuropsychology, vol. 4, 2nd edition*, ed. F. Boller & J. Grafman. Elsevier Science. [PB]
- (2002) Orienting of attention in left unilateral neglect. *Neuroscience and Biobehavioral Reviews* 26(2):217–34. [PB, rZWP]
- Bartolomeo, P., D’Erme, P. & Gainotti, G. (1994) The relationship between visuospatial and representational neglect. *Neurology* 44:1710–14. [PB]
- Basso, A., Bisiach, E. & Luzzatti, C. (1980) Loss of mental imagery: A case study. *Neuropsychologia* 18:435–42. [CG]
- Baylor, G. W. (1972) A treatise on the mind’s eye. Unpublished doctoral dissertation. Carnegie Mellon University, Pittsburgh. (University Microfilms No. 72–12, 699.) [NJTT]
- Behrmann, M. (2000) The mind’s eye mapped onto the brain’s matter. *Current Directions in Psychological Science* 9(2):50–54. [aZWP]
- Behrmann, M., Moscovitch, M. & Winocur, G. (1994) Intact visual imagery and impaired visual perception in a patient with visual agnosia. *Journal of Experimental Psychology: Human Perception and Performance* 20(5):1068–87. [aZWP]
- Behrmann, M. & Tipper, S. (1999a) Attention accesses multiple reference frames: Evidence from unilateral neglect. *Journal of Experimental Psychology: Human Perception and Performance* 25:83–101. [rZWP]
- (1999b) Attention accesses multiple reference frames: Evidence from visual neglect. *Journal of Experimental Psychology: Human Perception and Performance* 25(1): 83–101. [rZWP]
- Behrmann, M., Winocur, G. & Moscovitch, M. (1992) Dissociation between mental imagery and object recognition in a brain-damaged patient. *Nature* 359(6396):636–37. [aZWP]
- Bennett, G. K., Seashore, H. G. & Wesman, A. G. (1989) *DAT – Differential Aptitude Tests of personnel and career assessment: Space relations*. Harcourt Brace/The Psychological Corporation. [JRP]
- Bernbaum, K. & Chung, C. S. (1981) Müller-Lyer illusion induced by imagination. *Journal of Mental Imagery* 5(1):125–28. [aZWP]
- Beschin, N., Basso, A. & Della Sala, S. (2000) Perceiving left and imagining right: Dissociation in neglect. *Cortex* 36(3):401–14. [NB]
- Beschin, N., Cocchini, G., Della Sala, S. & Logie, R. H. (1997) What the eyes perceive, the brain ignores: A case of pure unilateral representational neglect. *Cortex* 33(1):3–26. [NB, aZWP]
- Betts, G. H. (1909) *The distribution and functions of mental imagery*. New York Teachers College (Contribution to Education Series, No. 26, 1–99). Columbia University Press. [MOB]
- Bickhard, M. & Terveen, L. (1995) *Foundational issues in artificial intelligence and cognitive science: Impasse and solution*. Elsevier Science. [JAT]
- Biederman, I. (1987) Recognition-by-components: A theory of human image understanding. *Psychological Review* 94:115–48. [MJ, rZWP]
- (2000) Recognizing depth-rotated objects: A review of recent research and theory. *Spatial Vision* 13:241–54. [MJ]
- Biederman, I. & Gerhardstein, P. C. (1995) Viewpoint-dependent mechanisms in visual object recognition. *Journal of Experimental Psychology: Human Perception and Performance* 21:1506–14. [MJ]
- Bisiach, E. & Berti, A. (1990) Waking images and neural activity. In: *The psychophysiology of mental imagery*, ed. G. Kunzendorf & A. A. Sheikh. Baywood. [PB]
- Bisiach, E., Capitani, E., Luzzatti, C. & Perani, D. (1981) Brain and conscious representation of outside reality. *Neuropsychologia* 19(4):543–51. [PB, NB]
- Bisiach, E. & Luzzatti, C. (1978) Unilateral neglect of representational space. *Cortex* 14:129–33. [PB]
- Bisiach, E., Luzzatti, C. & Perani, D. (1979) Unilateral neglect, representational schema and consciousness. *Brain* 102(3):609–18. [NB]
- Bisiach, E., Ricci, R., Lualdi, M. & Colombo, M. R. (1998) Perceptual and response bias in unilateral neglect: Two modified versions of the Milner Landmark task. *Brain and Cognition* 37(3):369–86. [rZWP]
- Blackmore, S. J., Brelstaff, G., Nelson, K. & Troscianko, T. (1995) Is the richness of the visual world an illusion? Transsaccadic memory for complex scenes. *Perception* 24(9):1075–81. [aZWP]
- Block, N. J. (1981) Introduction: What is the issue? In: *Imagery*, ed. N. Block. MIT Press. [aZWP]
- Block, N. J. & Fodor, J. A. (1972) Cognitivism and the analog/digital distinction. Unpublished manuscript. [rZWP]
- Blum, S., Moore, A. N., Adams, F. & Dash, P. K. (1999) A mitogen-activated protein kinase cascade in the CA1/CA2 subfield of the dorsal hippocampus is essential for long-term spatial memory. *Journal of Neuroscience* 19(9):3535–44. [rZWP]
- Bolles, R. C. (1969) The role of eye movements in the Müller-Lyer illusion. *Perception and Psychophysics* 6(3):175–76. [aZWP]
- Bonnel, A. M., Possami, C. A. & Schmitt, M. (1987) Early modulation of visual input: A study of attentional strategies. *The Quarterly Journal of Experimental Psychology* 39(4):757–76. [rZWP]
- Boschker, M. S. J., Bakker, F. C. & Michaels, C. F. (2002) Effect of mental imagery on realizing affordances. *Quarterly Journal of Experimental Psychology A: Human Experimental Psychology* 55A(3):775–92. [MR]
- Bosinelli, M. (1995) Mind and consciousness during sleep. *Behavioral Brain Research* 69:195–201. [CG]
- Bower, G. H. & Glass, A. L. (1976) Structural units and the reintegrative power of picture fragments. *Journal of Experimental Psychology: Human Learning and Memory* 2:456–66. [aZWP]
- Bradley, D. & Vido, D. (1984) Psychophysical functions for perceived and remembered distance. *Perception* 13:315–20. [WMP]
- Braine, L., Schauble, L., Kugelmass, S. & Winter, A. (1993) Representation of depth by children: Spatial strategies and lateral biases. *Developmental Psychology* 29:466–79. [AC]
- Brandt, S. A. & Stark, L. W. (1997) Spontaneous eye movements during visual imagery reflect the content of the visual scene. *Journal of Cognitive Neuroscience* 9(1):27–38. [aZWP]
- Braun, A. R., Balkin, H. J., Wesensten, N. J., Gwardy, F., Carson, R. E., Varga, M., Baldwin, P., Belenky, G. & Herscovitch, P. (1998) Dissociated pattern of activity in visual cortices and their projections during human rapid eye movement sleep. *Science* 279:91–95. [CG]
- Brigell, M., Uhlarik, J. & Goldhorn, P. (1977) Contextual influence on judgments of linear extent. *Journal of Experimental Psychology: Human Perception and Performance* 3(1):105–18. [aZWP]
- Broadhurst, F. M. (1964) Some aspects of the palaeoecology on non-marine faunas and rates of sedimentation in the Lancashire coal measures. *American Journal of Science* 262:865. Lithographic image in: Milton, R. (1997) *Shattering the myths of Darwinism*. Park Street Press. [JAT]
- Broerse, J. & Crassini, B. (1981) Misinterpretations of imagery-induced McCollough effects: A reply to Finke. *Perception and Psychophysics* 30:96–98. [aZWP]
- (1984) Investigations of perception and imagery using CAEs: The role of experimental design and psychophysical method. *Perception and Psychophysics* 35(2):155–64. [aZWP]
- Brooks, L. R. (1968) Spatial and verbal components of the act of recall. *Canadian Journal of Psychology* 22(5):349–68. [aZWP]
- Bullier, J., Hupé, J. M., James, A. & Girard, P. (1996) Functional interactions between areas V1 and V2 in the monkey. *Journey of Physiology (Paris)* 90:217–20. [SG]
- Burgess, N., Becker, S., King, J. A. & O’Keefe, J. (2001) Memory for events and their spatial context: Models and experiments. *Philosophical Transactions of the Royal Society of London B: Biological Sciences* 356:1493–503. [NB]
- Butter, C. M., Kosslyn, S., Mijovic-Prelec, D. & Riffle, A. (1997) Field-specific deficits in visual imagery following hemianopia due to unilateral occipital infarcts. *Brain* 120:217–28. [AC]
- Canon, L. K. (1970) Intermodality inconsistency of input and directed attention as determinants of the nature of adaptation. *Journal of Experimental Psychology* 84(1):141–47. [aZWP]
- (1971) Directed attention and maladaptive “adaptation” to displacement of the visual field. *Journal of Experimental Psychology* 88(3):403–408. [aZWP]
- Carlson-Radvansky, L. A. (1999) Memory for relational information across eye movements. *Perception and Psychophysics* 61(5):919–34. [aZWP]
- Carlson-Radvansky, L. A. & Irwin, D. E. (1995) Memory for structural information

- across eye movements. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 21(6):1441–58. [aZWP]
- Carpenter, P. A. & Eisenberg, P. (1978) Mental rotation and the frame of reference in blind and sighted individuals. *Perception and Psychophysics* 23(2):117–24. [aZWP]
- Carruthers, P. (1996) *Language, thought and consciousness*. Cambridge University Press. [PPS]
- Casey, E. (1976) *Imagining: A phenomenological study*. Indiana University Press. [aZWP]
- Cavallero, C., Cicogna, P., Natale, V., Occhionero, M. & Zito, A. (1992) Slow wave sleep dreaming. *Sleep* 15(6):562–6. [CG]
- Chambers, D. & Reisberg, D. (1985) Can mental images be ambiguous? *Journal of Experimental Psychology* 11:317–28. [aZWP]
- Charlot, V., Tzourio, N., Zilbovicius, M., Mazoyer, B. & Denis, M. (1992) Different mental imagery abilities result in different regional cerebral blood flow activation patterns during cognitive tasks. *Neuropsychologia* 30(6):565–80. [aZWP]
- Chatterjee, A. (2001) Language and space: Some interactions. *Trends in Cognitive Science* 5:55–61. [AC]
- (2002) Portrait profiles and the notion of agency. *Empirical Studies of the Arts* 20:33–41. [AC]
- Chatterjee, A. & Maher, L. (2000) Grammar and agrammatism. In: *Aphasia and language: Theory to practice*, ed. L. Gonzalez Rothi, B. Crosson & S. Nadeau. Guilford. [AC]
- Chatterjee, A., Maher, L. M., Gonzales-Rothi, L. J. & Heilman, K. M. (1995a) Asyntactic thematic role assignment: The use of a temporal-spatial strategy. *Brain and Language* 49:125–39. [AC]
- Chatterjee, A., Maher, L. M. & Heilman, K. M. (1995b) Spatial characteristics of thematic role representation. *Neuropsychologia* 33:643–48. [AC]
- Chatterjee, A. & Southwood, M. H. (1995) Cortical blindness and visual imagery. *Neurology* 45(12):2189–95. [AC, aZWP]
- Chatterjee, A., Southwood, M. H. & Basilico, D. (1999) Verbs, events and spatial representations. *Neuropsychologia* 37:395–402. [AC]
- Chelazzi, L., Miller, E. K., Duncan, J. & Desimone, R. (1993) A neural basis for visual search in inferior temporal cortex. *Nature* 363:345–47. [FvdV]
- Chokron, S. & Bartolomeo, P. (1999) Réduire expérimentalement la négligence spatiale unilatérale: Revue de la littérature et implications théoriques. *Revue de Neuropsychologie* 9(2–3):129–65. [PB]
- Chomsky, N. (1957a) Review of B. F. Skinner's *Verbal behavior*. In: *The structure of language*, J. A. Fodor & J. J. Katz. Prentice-Hall. [aZWP]
- (1957b) *Syntactic structures*. Mouton. [aZWP]
- Christman, S. & Pinger, K. (1997) Lateral biases in aesthetic preferences: Pictorial dimensions and neural mechanisms. *Laterality* 2:155–75. [AC]
- Christou, C. G. & Bulthoff, H. H. (1999) *The perception of spatial layout in a virtual world*. Technical Report 75, Max Planck Institute for Biological Cybernetics, Tübingen, Germany. [NB]
- Clancey, W. J. (1997) *Situated cognition: On human knowledge and computer representations*. Cambridge University Press. [JAT]
- Clancey, W. J. (1999) *Conceptual coordination: How the mind orders experience in time*. Erlbaum.
- Cocude, M., Mellet, E. & Denis, M. (1999) Visual and mental exploration of visuo-spatial configurations: Behavioral and neuroimaging approaches. *Psychological Research* 62(2–3):93–106. [aZWP]
- Cohen, M. S., Kosslyn, S. M., Breiter, H. C., DiGirolamo, G. J., Thompson, W. L., Anderson, A. K., Bookheimer, S. Y., Rosen, B. R. & Belliveau, J. W. (1996) Changes in cortical activity during mental rotation: A mapping study using functional MRI. *Brain* 119(Pt. 1):89–100. [rZWP]
- Colapinto, J. (2000) *As nature made him: The boy who was raised as a girl*. Harper Collins. [JAT]
- Collett, T. S. & Kelber, A. (1988) The retrieval of visuo-spatial memories by honeybees. *Journal of Comparative Physiology A* 163:145–50. [GEM]
- Collins, A. W. (1967) The epistemological status of the concept of perception. *Philosophical Review* 76:436–59. [EW]
- Coren, S. (1986) An efferent component in the visual perception of direction and extent. *Psychological Review* 93(4):391–410. [aZWP]
- Coren, S. & Porac, C. (1983) The creation and reversal of the Müller-Lyer illusion through attentional manipulation. *Perception* 12(1):49–54. [aZWP]
- Cornoldi, C., Bertucelli, B., Rocchi, P. & Sbrana, B. (1993) Processing capacity limitations in pictorial and spatial representations in the totally congenitally blind. *Cortex* 29(4):675–89. [aZWP]
- Cornoldi, C., Calore, D. & Pra-Baldi, A. (1979) Imagery ratings and recall in congenitally blind subjects. *Perceptual and Motor Skills* 48(2):627–39. [aZWP]
- Coslett, H. B. (1997) Neglect in vision and visual imagery: A double dissociation. *Brain* 120:1163–71. [rZWP]
- Craig, E. M. (1973) Role of mental imagery in free recall of deaf, blind, and normal subjects. *Journal of Experimental Psychology* 97(2):249–53. [aZWP]
- Craver-Lemley, C. & Arterberry, M. E. (2001) Visual imagery interference in a detection task. *Spatial Vision* 14:101–19. [MEA]
- Craver-Lemley, C., Arterberry, M. E. & Reeves, A. (1997) The effects of imagery on Vernier acuity under conditions of induced depth. *Journal of Experimental Psychology: Human Perception and Performance* 23:3–13. [MEA]
- (1999) Illusory illusory conjunctions: The conjoining of features of real and imagined stimuli. *Journal of Experimental Psychology: Human Perception and Performance* 25:1036–49. [MEA]
- Craver-Lemley, C. & Reeves, A. (1987) Visual imagery selectively reduces Vernier acuity. *Perception* 16:599–614. [MEA]
- (1992) How visual imagery interferes with vision. *Psychological Review* 99:633–49. [MEA]
- Crawford, H. J. (1996) Cerebral brain dynamics of mental imagery: Evidence and issues for hypnosis. In: *Hypnosis and imagination*, ed. R. G. Kunzendorf. Baywood. [aZWP]
- Crick, F. & Koch, C. (1995) Are we aware of neural activity in primary visual cortex? *Nature* 375(11):121–23. [CG, arZWP]
- Cummins, R. (1996) *Representations, targets, and attitudes*. MIT Press. [NJTT]
- (1997) The LOT of the causal theory of mental content. *Journal of Philosophy* 94:535–42. [NJTT]
- Currie, G. (1995) Visual imagery as the simulation of vision. *Mind and Language* 10(1–2):25–44. [aZWP]
- Dalla Barba, G. (2002) *Memory, consciousness and temporality*. Kluwer Academic. [GDB]
- Dalman, J. E., Verhagen, W. I. M. & Huygen, P. L. M. (1997) Cortical blindness. *Clinical Neurology and Neurosurgery* (Dec):282–86. [aZWP]
- Dauterman, W. L. (1973) A study of imagery in the sighted and the blind. In: *American Foundation for the Blind, Research Bulletin, January 1973*, pp. 95–167. American Foundation for the Blind. [aZWP]
- Davies, T. N. & Spencer, J. (1977) An explanation for the Mueller-Lyer illusion. *Perceptual and Motor Skills* 45(1):219–24. [aZWP]
- Dehaene, S., Dehaene-Lambertz, G. & Cohen, L. (1998) Abstract representations of numbers in the animal and human brain. *Trends in Neurosciences* 21(8):355–61. [rZWP]
- De Kamps, M. & van der Velde, F. (2001) Using a recurrent network to bind form, color and position into a unified percept. *Neurocomputing* 38–40:523–28. [FvdV]
- Delgado, J. M. R. (1969) *Physical control of mind*. Harper and Row. [ENS]
- Del Gratta, C., Di Matteo, R., De Nicola, A., Ferretti, A., Tartaro, A., Bonomo, L., Romani, G. L. & Olivetti Belardinelli, M. (2001) Sensory image generation: A comparison between different sensory modalities with fMRI. In: *Proceedings of the 7th Annual Meeting of the Organization for Human Brain Mapping*, ed. A. W. Toga, R. S. J. Frackowiak & J. C. Mazziotta. *Neuroimage* 13(6, part 2):S394. [MOB]
- DeLucia, P. R. & Liddell, G. W. (1998) Cognitive motion extrapolation and cognitive clocking in prediction motion tasks. *Journal of Experimental Psychology: Human Perception and Performance* 24(3):901–14. [aZWP]
- Dement, W. & Kleitman, N. (1957) The relation of eye movements during sleep to dream activity: An objective method for the study of dreaming. *Journal of Experimental Psychology* 53:339–46. [CG]
- Dement, W. & Wolpert, E. A. (1958) The relation of eye movements, body motility and external stimuli to dream content. *Journal of Experimental Psychology* 55:543–53. [CG]
- Denis, M. & Carfantan, M. (1985) People's knowledge about images. *Cognition* 20(1):49–60. [SMK, arZWP]
- Denis, M. & Kosslyn, S. M. (1999) Scanning visual mental images: A window on the mind. *Cahiers de Psychologie Cognitive (Current Psychology of Cognition)* 18(4):409–65. [aZWP]
- Dennett, D. C. (1978) *Brainstorms*. MIT Press/A Bradford Book. [rZWP]
- (1982) How to study consciousness empirically: Or, Nothing comes to mind. *Synthese* 53:159–80. [DCD]
- (1991) *Consciousness explained*. Little, Brown/Allen Lane. [DCD, arZWP]
- (1997) The Cartesian theater and “filling in”: The stream of consciousness. In: *The nature of consciousness. Philosophical debates*, ed. N. Block, O. Flanagan & G. Guezeldere. MIT Press. [VG]
- Descartes, R. (1637/1985) *Dioptrics*. In: *The philosophical writings of Descartes, vols. 1 and 2*, trans. J. Cottingham, R. Stoothoff & D. Murdoch. Cambridge University Press. [PPS]
- Desimone, R., Albright, T. D., Gross, C. G. & Bruce, C. (1984) Stimulus selective properties of inferior temporal neurons in the macaque. *Journal of Neuroscience* 4:2051–62. [DI]
- D'Esposito, M. (2001) Functional neuroimaging of working memory. In: *Handbook of functional neuroimaging of cognition*, ed. R. Cabeza & A. Kingstone. MIT Press. [FvdV]
- D'Esposito, M., Detre, J. A., Aguirre, G. K., Stallcup, M., Alsop, D. C., Tippet, L. J. & Farah, M. J. (1997) A functional MRI study of mental image generation. *Neuropsychologia* 35(5):725–30. [MOB, aZWP]
- De Vreese, L. P. (1991) Two systems for colour-naming defects: Verbal

- disconnection vs. colour imagery disorder. *Neuropsychologia* 29(1):1–18. [aZWP]
- De Yoe, E. A., Bandettini, P., Neitz, J., Miller, D. & Winans, P. (1994) Functional magnetic resonance imaging (fMRI) of the human brain. *Journal of Neuroscience Methods* 54:171–87. [SMK]
- Divadkar, V. A. & McNamara, T. P. (1997) Viewpoint dependence in scene recognition. *Psychological Science* 8(4):302–307. [NB]
- Dodds, A. G. (1983) Mental rotation and visual imagery. *Journal of Visual Impairment and Blindness* 77(1):16–18. [aZWP]
- Doricchi, F., Guariglia, C., Paolucci, S. & Pizzamiglio, L. (1993) Disturbances of the rapid eye movements (REMs) of REM sleep in patients with unilateral attentional neglect: Clue for the understanding of the functional meaning of REMs. *Electroencephalography and Clinical Neurophysiology* 87(3):105–16. [PB]
- Downing, C. J. (1988) Expectancy and visual-spatial attention: Effects on perceptual quality. *Journal of Experimental Psychology: Human Perception and Performance* 14:188–202. [rZWP]
- Driver, J. & Vuilleumier, P. (2001) Perceptual awareness and its loss in unilateral neglect and extinction. *Cognition* 79(1–2):39–88. [rZWP]
- Easton, R. D. & Bentzen, B. L. (1987) Memory for verbally presented routes: A comparison of strategies used by blind and sighted people. *Journal of Visual Impairment and Blindness* 81(3):100–105. [aZWP]
- Easton, R. D. & Scholl, M. J. (1995) Object-array structure, frames of reference, and retrieval of spatial knowledge. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 21(2):483–500. [NB]
- Edelman, S. (1995) Representation, similarity, and the chorus of prototypes. *Mind and Machines* 5:45–68. [MJ]
- (1998) Representation is representation of similarities. *Behavioral and Brain Sciences* 21:449–98. [HA, rZWP, PPS]
- Edelman, S. & Bühlhoff, H. H. (1992) Orientation dependence in the recognition of familiar and novel views of 3D objects. *Vision Research* 32:2385–4000. [MJ]
- Edwards, P. N. (1996) *The closed world: Computers and the politics of discourse in cold war America*. MIT Press. [JAT]
- Ehrlichman, H. & Barrett, J. (1983) Right hemisphere specialization for mental imagery: A review of the evidence. *Brain and Cognition* 2:55–76. [GEM]
- Emanuel, E. J., Wendler, D. & Grady, C. (2000) What makes clinical research ethical? *Journal of the American Medical Association* 283(20):2701–11. [JAT]
- Engel, S. A., Rumelhart, D. E., Wandell, B. A., Lee, A. T., Glover, G. H., Chichilnisky, E. J. & Shadlen, M. N. (1994) fMRI of human visual cortex. *Nature* 369:525. [SMK]
- Escher, M. C. (1960) *The graphic work of M. C. Escher*. Hawthorn Books. [aZWP]
- Fallick, A. J., Mueller, T. J. & Strik, W. K. (1997) Neurophysiological correlates of mental imagery in different sensory modalities. *International Journal of Psychophysiology* 25(2):145–53. [MOB]
- Farah, M. J. (1984) The neurological basis of mental imagery: A componential analysis. *Cognition* 18:245–72. [GG, GEM]
- (1988) Is visual imagery really visual? Overlooked evidence from neuropsychology. *Psychological Review* 95(3):307–17. [aZWP]
- (1989) Mechanisms of imagery-perception interaction. *Journal of Experimental Psychology: Human Perception and Performance* 15:203–11. [arZWP]
- (1994) Beyond “pet” methodologies to converging evidence. *Trends in Neurosciences* 17(12):514–15. [aZWP]
- (1995a) Current issues in the neuropsychology of image generation. *Neuropsychologia* 23:1455–72. [GG]
- (1995b) The neural bases of mental imagery. In: *The cognitive neurosciences*, ed. M. S. Gazzaniga. MIT Press. [aZWP]
- Farah, M. J., Levine, D. & Calvanio, R. (1988) A case study of mental imagery deficit. *Brain and Cognition* 8:147–64. [GG]
- Farah, M. J., Soso, M. J. & Dasheiff, R. M. (1992) Visual angle of the mind's eye before and after unilateral occipital lobectomy. *Journal of Experimental Psychology: Human Perception and Performance* 18(1):241–46. [aZWP]
- Farley, A. M. (1976) A computer implementation of constructive visual imagery and perception. In: *Eye movements and psychological processes*, ed. R. A. Monty and J. W. Senders. Erlbaum. [NJTT]
- Felleman, D. J. & Van Essen, D. C. (1991) Distributed hierarchical processing in primate visual cortex. *Cerebral Cortex* 1:1–47. [HA]
- Festinger, L. (1957) *A theory of cognitive dissonance*. Stanford University Press. [JAT]
- Festinger, L., White, C. W. & Allyn, M. R. (1968) Eye movements and decrement in the Müller-Lyer illusion. *Perception and Psychophysics* 3(5-B):376–82. [aZWP]
- Fidelman, U. (1994) A misleading implication of the metabolism scans of the brain. *International Journal of Neuroscience* 74(1–4):105–108. [rZWP]
- Finke, R. A. (1979) The functional equivalence of mental images and errors of movement. *Cognitive Psychology* 11:235–64. [SM, aZWP]
- (1989) *Principles of mental imagery*. MIT Press. [aZWP, PPS]
- Finke, R. A. & Freyd, J. J. (1989) Mental extrapolation and cognitive penetrability: Reply to Ranney and proposals for evaluative criteria. *Journal of Experimental Psychology: General* 118(4):403–408. [aZWP]
- Finke, R. A. & Kosslyn, S. M. (1980) Mental imagery acuity in the peripheral visual field. *Journal of Experimental Psychology: Human Perception and Performance* 6(1):126–39. [aZWP]
- Finke, R. A. & Kurtzman, H. S. (1981a) Mapping the visual field in mental imagery. *Journal of Experimental Psychology: General* 110(4):501–17. [aZWP]
- (1981b) Methodological considerations in experiments on imagery acuity. *Journal of Experimental Psychology: Human Perception and Performance* 7(4):848–55. [aZWP]
- Finke, R. A. & Pinker, S. (1982) Spontaneous imagery scanning in mental extrapolation. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 8(2):142–47. [SMK, aZWP]
- Finke, R. A., Pinker, S. & Farah, M. J. (1989) Reinterpreting visual patterns in mental imagery. *Cognitive Science* 13(1):51–78. [aZWP, PPS]
- Finke, R. A. & Schmidt, M. J. (1977) Orientation-specific color aftereffects following imagination. *Journal of Experimental Psychology: Human Perception and Performance* 3(4):599–606. [aZWP]
- Finke, R. A., Ward, T. B. & Smith, S. M. (1992) *Creative cognition: Theory, research, and applications*. MIT Press. [SMK]
- Fletcher, P. C., Shallice, T., Frith, C. D., Frackowiak, R. S. J. & Dolan, R. J. (1996) Brain activity during memory retrieval: The influence of imagery and semantic cueing. *Brain* 119(5):1587–96. [aZWP]
- Fodor, J. A. (1968) The appeal to tacit knowledge in psychological explanation. *Journal of Philosophy* 65:627–40. [aZWP]
- (1975) *The language of thought*. Crowell. [GDB, aZWP, PPS, NJTT]
- (1981) Imagistic representation. In: *Imagery*, ed. N. Block. MIT Press. [aZWP]
- Fodor, J. A. & Pylyshyn, Z. W. (1988) Connectionism and cognitive architecture: A critical analysis. *Cognition* 28:3–71. [arZWP]
- Fomin, S. V., Sokolov, E. N. & Vaitkevicius, G. G. (1979) *Iskusstvemye organy chuvstv (Artificial sensory organs)*. Nauka. [ENS]
- Fosse, R., Stickgold, R. & Hobson, J. A. (2001) Brain-mind states: Reciprocal variation in thoughts and hallucinations. *Psychological Science* 12:30–36. [CG]
- Foulkes, D. (1962) Dream report from different stages of sleep. *Journal of Abnormal and Social Psychology* 65:14–25. [CG]
- Fox, P. T., Mintun, M. A., Raichle, M. E., Miezin, F. M., Allman, J. M. & Van Essen, D. C. (1986) Mapping human visual cortex with positron emission tomography. *Nature* 323(6091):806–809. [aZWP]
- Frak, V., Paulignan, Y. & Jeannerod, M. (2001) Orientation of the opposition axis in mentally simulated grasping. *Experimental Brain Research* 136:120–27. [MR]
- Freimuth, M. & Wapner, S. (1979) The influence of lateral organization on the evaluation of paintings. *British Journal of Psychology* 70:211–18. [AC]
- Freud, S. (1875/1975) *Project for a scientific psychology. Standard edition, vol. I*. The Hogarth Press, 1975. [CG]
- (1900/1975) *The interpretation of dreams. Standard edition, vol. IV*. The Hogarth Press, 1975. [CG]
- (1911/1975) *Formulation of the two principles of mental functioning. Standard edition, vol. XII*. The Hogarth Press, 1975. [CG]
- Freyd, J. J. (1987) Dynamic mental representations. *Psychological Review* 94(4):427–38. [MR]
- Freyd, J. J. & Finke, R. A. (1984) Representational momentum. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 10(1):126–32. [aZWP]
- Funt, B. V. (1980) Problem-solving with diagrammatic representations. *Artificial Intelligence* 13(3):201–30. [aZWP]
- Gainotti, G. (1993) The role of spontaneous eye movements in orienting attention and in unilateral neglect. In: *Unilateral neglect: Clinical and experimental studies*, ed. I. H. Robertson & J. C. Marshall. Erlbaum. [PB]
- Gainotti, G., D'Erme, P. & Bartolomeo, P. (1991) Early orientation of attention toward the half space ipsilateral to the lesion in patients with unilateral brain damage. *Journal of Neurology, Neurosurgery and Psychiatry* 54:1082–89. [PB]
- Gallistel, C. R. & Gelman, R. (2000) Nonverbal numerical cognition: From reals to integers. *Trends in Cognitive Sciences* 4:59–65. [rZWP]
- Gibson, J. J. (1966) *The senses considered as perceptual systems*. Houghton Mifflin. [aZWP]
- Gigerenzer, G. & Selten, R. (2001) Rethinking rationality. In: *Bounded rationality*, ed. G. Gigerenzer & R. Selten. MIT Press. [MR]
- Gilden, D., Blake, R. & Hurst, G. (1995) Neural adaptation of imaginary visual motion. *Cognitive Psychology* 28(1):1–16. [aZWP]
- Glasgow, J. I. (1993) The imagery debate revisited: A computational perspective. *Computational Intelligence* 9:310–33. [NJTT]
- Glenberg, A. M. & Robertson, D. A. (2000) Symbol grounding and meaning: A comparison of high-dimensional and embodied theories of meaning. *Journal of Memory and Language* 43(3):379–401. [aZWP]

- Clisky, M. L., Tataryn, D. J. & Kihlstrom, J. F. (1995) Hypnotizability and mental imagery. *International Journal of Clinical and Experimental Hypnosis* 43(1):34–54. [aZWP]
- Gluck, K. & Pew, R. (2002) The agent-based modeling and behavior representation (AMBR) comparison project: Round III—Modeling category learning. *Proceedings of the 24th annual meeting of the Cognitive Science Society, Aug. 8–10 2002*.
- Goldenberg, G. (1992) Loss of visual imagery and loss of visual knowledge – a case study. *Neuropsychologia* 30(12):1081–99. [GG, aZWP]
- (1993) The neural basis of mental imagery. *Bailliere's Clinical Neurology* 2:265–86. [GG]
- (1998) Is there a common substrate for visual recognition and visual imagery? *Neurocase* 4:141–48. [GG]
- Goldenberg, G. & Armer, C. (1991) Visual imagery and knowledge about the visual appearance of objects in patients with posterior cerebral artery lesions. *Brain and Cognition* 15(2):160–86. [GG, aZWP]
- Goldenberg, G., Müllbacher, W. & Nowak, A. (1995) Imagery without perception – a case study of anosognosia for cortical blindness. *Neuropsychologia* 33(11):1373–82. [AC, aZWP]
- Goodale, M. A., Jacobson, J. S. & Keillor, J. M. (1994) Differences in the visual control of pantomimed and natural grasping movements. *Neuropsychologia* 32(10):1159–78. [aZWP]
- Goryo, K., Robinson, J. O. & Wilson, J. A. (1984) Selective looking and the Mueller-Lyer illusion: The effect of changes in the focus of attention on the Mueller-Lyer illusion. *Perception* 13(6):647–54. [aZWP]
- Gosselin, F. & Schyns, P. G. (in press) Superstitious perceptions reveal properties of internal representations. *Psychological Science*. [FG]
- Gottesmann, C. (1999) Neurophysiological support of consciousness during waking and sleep. *Progress in Neurobiology* 59:469–508. [CG]
- (2000) Each distinct type of mental state is supported by specific brain functions. *Behavioral and Brain Sciences* 23:941–43. [CG]
- Gould, E., Vail, N., Wagers, M. & Gross, C. G. (2001) Adult-generated hippocampal and neocortical neurons in macaques have a transient existence. *PNAS* 98(19):10910–17. [ENS]
- Grafman, J. (2000) Conceptualizing functional neuroplasticity. *Journal of Communication Disorders*, 33:345–356.
- Grafman, J. & Litvan, I. (1999) Evidence for four forms of neuroplasticity. In: *Neuroplasticity: Building a bridge from the laboratory to the clinic*, ed. J. Grafman and Y. Christen. Springer-Verlag.
- Gregory, R. L. (1965) Inappropriate constancy explanation of spatial distortions. *Nature* 207(4999):891–93. [aZWP]
- (1979) *Eye and brain* (3rd edition). Weidenfeld and Nicholson. [RN]
- Griffiths, A. F. & Zaidi, Q. (1998) Rigid objects that appear to bend. *Perception* 27:799–802. [QZ]
- (2000) Perceptual assumptions and projective distortions in a three-dimensional shape illusion. *Perception* 29:171–200. [QZ]
- Grossberg, S. (1994) 3-D vision and figure-ground separation by visual cortex. *Perception and Psychophysics* 55:48–120. [SG]
- (1997) Cortical dynamics of 3-D figure-ground perception of 2-D pictures. *Psychological Review* 104:1583–605. [SG]
- (1999a) How does the cerebral cortex work? Learning, attention, and grouping by the laminar circuits of visual cortex. *Spatial Vision* 12:163–87. [SG]
- (1999b) The link between brain learning, attention, and consciousness. *Consciousness and Cognition* 8:1–44. [SG]
- (2000) How hallucinations may arise from brain mechanisms of learning, attention, and volition. *Journal of the International Neuropsychological Society* 6:583–92. [SG]
- Grossberg, S. & Raizada, R. D. S. (2000) Contrast-sensitive perceptual grouping and object-based attention in the laminar circuits of primary visual cortex. *Vision Research* 40:1413–32. [SG]
- Grusser, O.-J., Selke, T. & Zynda, B. (1988) Cerebral lateralization and some implications for art, aesthetic perception and artistic activity. In: *Beauty and the brain. Biological aspects of aesthetics*, ed. I. Rentschler, B. Herzberger & D. Epstein. Birkhauser Verlag. [AC]
- Guariglia, C., Padovani, A., Pantano, P. & Pizzamiglio, L. (1993) Unilateral neglect restricted to visual imagery. *Nature* 364(6434):235–37. [PB, NB]
- Haber, R. N. & Haber, R. B. (1964) Eidetic imagery. *Perceptual and Motor Skills* 19:131–38. [SM]
- Halgren, E., Babb, T. L. & Crandall, P. H. (1978a) Activity of human hippocampal formation and amygdala during memory testing. *Electroencephalography and Clinical Neurophysiology* 45:585–601. [GEM]
- Halgren, E., Walter, R. D., Cherlow, D. G. & Crandall, P. H. (1978b) Mental phenomena evoked by electrical stimulation of human hippocampal formation and amygdala. *Brain* 101:83–117. [GEM]
- Halper, F. (1997) The illusion of the future. *Perception* 26:1321–22. [QZ]
- Halpern, A. & Kelly, M. (1993) Memory bias in left versus right implied motion. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 19:471–84. [AC]
- Hamilton, R. H. & Pascual-Leone, A. (1998) Cortical plasticity associated with braille learning. *Trends in Cognitive Sciences* 2:168–174.
- Hampson, P. J. & Duffy, C. (1984) Verbal and spatial interference effects in congenitally blind and sighted subjects. *Canadian Journal of Psychology* 38(3):411–20. [aZWP]
- Hans, M. A. (1974) Imagery and modality in paired-associate learning in the blind. *Bulletin of the Psychonomic Society* 4(1):22–24. [aZWP]
- Harman, G. (1990) The intrinsic quality of experience. In: *Philosophical perspectives 4: Action theory and philosophy of mind*, ed. J. E. Tomberlin. Ridgeview. [IG]
- Harris, J. P. (1982) The VVIQ imagery-induced McCollough effects: An alternative analysis. *Perception and Psychophysics* 32(3):290–92. [aZWP]
- Hasnain, M. K., Fox, P. T. & Woldorff, M. G. (1998) Intersubject variability of functional areas in the human visual cortex. *Human Brain Mapping* 6:301–15. [SMK]
- Hayes, J. R. (1973) On the function of visual imagery in elementary mathematics. In: *Visual information processing*, ed. W. G. Chase. Academic Press. [aZWP]
- Heath, T. L., ed. (1956) *The thirteen books of Euclid's elements, vol. 1*. Dover. [KKN]
- Hebb, D. O. (1968) Concerning imagery. *Psychological Review* 75:466–77. [NJTT]
- Heller, M. A. & Kennedy, J. M. (1990) Perspective taking, pictures, and the blind. *Perception and Psychophysics* 48(5):459–66. [aZWP]
- Henderson, J. M. & Hollingworth, A. (1999) The role of fixation position in detecting scene changes across saccades. *Psychological Science* 10(5):438–43. [aZWP]
- Herman, J. H., Erman, M., Boys, R., Peiser, L., Taylor, M. E. & Roffwarg, H. P. (1984) Evidence for directional correspondence between eye movements and dream imagery in REM sleep. *Sleep* 7:52–63. [CG]
- Hilbert, D. (1899/1999) *Grundlagen der Geometrie. 14c. Auflage*. B. G. Teubner. (Original work published in 1899). [KKN]
- Hinton, G. E. (1979) Some demonstrations of the effects of structural descriptions in mental imagery. *Cognitive Science* 3:231–50. [aZWP]
- Hobson, J. A. & McCarley, R. W. (1977) The brain as dream state generator: An activation-synthesis hypothesis of the dream process. *American Journal of Psychiatry* 134:1335–48. [CG]
- Hobson, J. A., Pace-Schott, E. F. & Stickgold, R. (2000) Dreaming and the brain: Toward a cognitive neuroscience of conscious states. *Behavioral and Brain Sciences* 23:793–842. [CG]
- Hobson, J. A., Stickgold, R. & Pace-Schott, E. F. (1998) The neuropsychology of REM sleep dreaming. *NeuroReport* 9:R1–R14. [CG]
- Hochberg, J. (1968) In the mind's eye. In: *Contemporary theory and research in visual perception*, ed. R. N. Haber. Holt, Rinehart and Winston. [aZWP]
- Hochberg, J. & Gellman, L. (1977) The effect of landmark features on mental rotation times. *Memory and Cognition* 5(1):23–26. [aZWP]
- Hoeng, P. (1972) The effects of eye movements, fixation and figure size on decrement in the Muller-Lyer illusion. *Dissertation Abstracts International* 33(6-B):2835. [aZWP]
- Hollard, V. D. & Delius, J. D. (1982) Rotational invariance in visual pattern recognition by pigeons and humans. *Science* 218:804–806. [GEM]
- Höllinger, P., Beisteiner, R., Lang, W., Lindinger, G. & Berthoz, A. (1999) Mental representations of movements. Brain potentials associated with imagination of eye movements. *Clinical Neurophysiology* 110:799–805. [MOB]
- Hong, C. C. H., Gillin, J. C., Dow, B. C., Wu, J. & Buschbaum, M. S. (1995) Localized and lateralized cerebral glucose metabolism associated with eye movements during REM sleep and wakefulness: A positron emission tomography (PET) study. *Sleep* 18:570–80. [CG]
- Horowitz, M. J. (1970) *Image formation and cognition*. Appleton-Century-Crofts. [GEM]
- Horst, S. (1996) *Symbols, computation and intentionality: A critique of the computational theory of mind*. University of California Press. [NJTT]
- Howard, I. P. (1982) *Human visual orientation*. Wiley. [aZWP]
- Howard, R. J., flytche, D. H., Barnes, J., McKeefry, D., Ha, Y., Woodruff, P. W., Bullmore, E. T., Simmons, A., Williams, S. C., David, A. S. & Brammer, M. (1998) The functional anatomy of imagining and perceiving colour. *NeuroReport* 9(6):1019–23. [aZWP]
- Humphrey, G. (1951) *Thinking: An introduction to its experimental psychology*. Methuen. [aZWP]
- Ingle, D. (1990) Spatial short-term memory: Evolutionary perspectives and discoveries from split-brain studies. *Behavioral and Brain Sciences* 15:760–62. [DI]
- Intons-Peterson, M. J. (1983) Imagery paradigms: How vulnerable are they to experimenters' expectations? *Journal of Experimental Psychology: Human Perception and Performance* 9(3):394–412. [aZWP]
- Intons-Peterson, M. J. & White, A. R. (1981) Experimenter naïvete and imaginal judgments. *Journal of Experimental Psychology: Human Perception and Performance* 7(4):833–43. [aZWP]

- Intraub, H. (1981) Identification and processing of briefly glimpsed visual scenes. In: *Eye movements: Cognition and visual perception*, ed. R. A. M. D. F. Fisher & J. W. Senders. Erlbaum. [aZWP]
- Irwin, D. E. (1991) Information integration across saccadic eye movements. *Cognitive Psychology* 23:420–56. [aZWP]
- (1993) Perceiving an integrated visual world. In: *Attention and performance XIV*, ed. D. E. Meyer & S. Kornblum. MIT Press. [aZWP]
- Ishai, A. & Sagi, D. (1995) Common mechanisms of visual imagery and perception. *Science* 268(5218):1772–74. [arZWP]
- Izmailov, C. A. & Sokolov, E. N. (1991) Spherical model of color and brightness discrimination. *Psychological Science* 2:249–59. [ENS]
- Jankowiak, J., Kinsbourne, M., Shalev, R. S. & Bachman, D. L. (1992) Preserved visual imagery and categorization in a case of associative visual agnosia. *Journal of Cognitive Neuroscience* 4(2):119–31. [aZWP]
- Janssen, J. J. & Sheikh, A. A. (1994) Enhancing athletic performance through imagery: An overview. In: *Imagery in sports and physical performance*, ed. A. A. Sheikh & E. R. Korn. Baywood. [MR]
- Jeannerod, M. (1995) Mental imagery in the motor context. *Neuropsychologia* 33(11):1419–32. [MOB]
- Johnson, D. M. (1939) Confidence and speed in the two-category judgment. *Archives of Psychology* 34:1–53. [WMP]
- Johnson, R. A. (1980) Sensory images in the absence of sight: Blind versus sighted adolescents. *Perceptual and Motor Skills* 51(1):177–78. [SM, aZWP]
- Johnson, S. H., Corballis, P. M. & Gazzaniga, M. S. (2001) Within grasp but out of reach: Evidence for a double dissociation between imagined hand and arm movements in the left cerebral hemisphere. *Neuropsychologia* 39:36–50. [MOB]
- Jonides, J., Kahn, R. & Rozin, P. (1975) Imagery instructions improve memory in blind subjects. *Bulletin of the Psychonomic Society* 5(5):424–26. [aZWP]
- Julesz, B. (1971) *Foundations of Cyclopean perception*. University of Chicago Press. [DCD]
- (1981) Textons, the elements of texture perception, and their interactions. *Nature* 290:91–97. [aZWP]
- Julstrom, B. A. & Baron, R. J. (1985) A model of mental imagery. *International Journal of Man-Machine Studies* 23:313–34. [NJTT]
- Just, M. A. & Carpenter, P. A. (1976) Eye fixations and cognitive processes. *Cognitive Psychology* 8(4):441–80. [aZWP]
- Kable, J., Spellmeyer-Lease, J. & Chatterjee, A. (2002) Neural substrates of action event knowledge. *Journal of Cognitive Neuroscience* 14:795–805. [AC]
- Kant, E. (1787/1964) *The critique of pure reason*. Macmillan. [EW]
- Kelly, F. & Grossberg, S. (2000) Neural dynamics of 3-D surface perception: figure-ground separation and lightness perception. *Perception and Psychophysics* 62:1596–618. [SG]
- Kelso, J. A. S., Cook, E., Olson, M. E. & Epstein, W. (1975) Allocation of attention and the locus of adaptation to displaced vision. *Journal of Experimental Psychology: Human Perception and Performance* 1:237–45. [aZWP]
- Kerr, N. H. (1983) The role of vision in “visual imagery” experiments: Evidence from the congenitally blind. *Journal of Experimental Psychology: General* 112(2):265–77. [aZWP]
- Kerst, S. M. & Howard, J. H. (1978) Memory psychophysics for visual area and length. *Memory and Cognition* 6:327–35. [WMP]
- Kessel, F. S. (1972) Imagery: A dimension of mind rediscovered. *British Journal of Psychology* 63:149–62. [NJTT]
- King, J. A., Burgess, N., Hartley, T., Vargha-Khadem, F. & O’Keefe, J. (2002) The human hippocampus and viewpoint dependence in spatial memory. *Hippocampus* 12(6). (in press). [NB]
- Klatzky, R. L. & Lederman, S. J. (1995) Identifying objects from a haptic glance. *Perception and Psychophysics* 57:1111–23. [MJ]
- (1999) The haptic glance: A route to rapid object identification and manipulation. *Attention and Performance* 17:165–96. [MJ]
- Klein, G. & Crandall, B. W. (1995) The role of mental simulation in problem solving and decision making. In: *Local applications of the ecological approach to human-machine systems. Vol. 2: Resources for ecological psychology*, ed. P. Hancock. Erlbaum. [aZWP]
- Klein, I., Dubois, J., Mangin, J. M., Kherif, F., Flandin, G., Poline, J. B., Denis, M., Kosslyn, S. M. & Le Bihan, D. (submitted) Retinotopic organization of visual mental images as revealed by functional Magnetic Resonance Imaging. Harvard University manuscript. [SMK]
- Kosslyn, S. M. (1978) Measuring the visual angle of the mind’s eye. *Cognitive Psychology* 10:356–89. [SMK, arZWP]
- (1980) *Image and mind*. Harvard University Press. [SMK, rZWP, NJTT, FvdV]
- (1981) The medium and the message in mental imagery: A theory. *Psychological Review* 88:46–66. [aZWP] Also in: *Imagery*, ed. N. Block. MIT Press, 1981. [VG]
- (1983) *Ghosts in the mind’s machine – creating and using images in the brain*. W. W. Norton. [GG]
- (1992) Mental imagery. In: *An invitation to cognitive science, vol. 2: Visual cognition and action*, ed. D. N. Osherson, S. M. Kosslyn & J. M. Hollerbach. MIT Press. [JAT]
- (1994) *Image and brain: The resolution of the imagery debate*. MIT Press. [PB, EdH, DCD, SMK, MOB, arZWP, MR, PPS, NJTT, JAT, EW]
- Kosslyn, S. M., Ball, T. M. & Reiser, B. J. (1978) Visual images preserve metric spatial information: Evidence from studies of image scanning. *Journal of Experimental Psychology: Human Perception and Performance* 4:46–60. [JRP, aZWP]
- Kosslyn, S. M., Ganis, G. & Thompson, W. L. (2001) Neural foundations of imagery. *Nature Reviews Neuroscience* 2:635–42. [SMK]
- Kosslyn, S. M., Pascual-Leone, A., Felician, O., Camposano, S., Keenan, J. P., Thompson, W. L., Ganis, G., Sukel, K. E. & Alpert, N. M. (1999a) The role of area 17 in visual imagery: Convergent evidence from PET and rTMS. *Science* 284(April 2):167–70. [EdH, CG, SMK, aZWP]
- Kosslyn, S. M., Pinker, S., Smith, G. & Shwartz, S. P. (1979) On the demystification of mental imagery. *Behavioral and Brain Science* 2:535–48. [aZWP, PPS]
- Kosslyn, S. M. & Pomerantz, J. R. (1977) Imagery, propositions, and the form of internal representations. *Cognitive Psychology* 9:52–76. [SMK]
- Kosslyn, S. M., Sukel, K. E. & Bly, B. M. (1999b) Squinting with the mind’s eye: Effects of stimulus resolution on imaginal and perceptual comparisons. *Memory and Cognition* 27(2):276–87. [aZWP]
- Kosslyn, S. M. & Sussman, A. L. (1995) Roles of imagery in perception: Or, There is no such thing as immaculate perception. In: *The cognitive neurosciences*, ed. M. S. Gazzaniga. MIT Press. [aZWP]
- Kosslyn, S. M. & Thompson, W. L. (2002) When is early visual cortex activated during visual mental imagery? *Psychological Bulletin*. (under review). [SMK]
- Kosslyn, S. M., Thompson, W. L., Kim, I. J. & Alpert, N. M. (1995) Topographical representations of mental images in primary visual cortex. *Nature* 378:496–98. [aZWP]
- Kosslyn, S. M., Thompson, W. L., Kim, I. J., Rauch, S. L. & Alpert, N. M. (1996) Individual differences in cerebral blood flow in area 17 predict the time to evaluate visualized letters. *Journal of Cognitive Neuroscience* 8:78–82. [SMK]
- Kowler, E. (1989) Cognitive expectations, not habits, control anticipatory smooth oculomotor pursuit. *Vision Research* 29:1049–57. [aZWP]
- (1990) The role of visual and cognitive processes in the control of eye movement. In: *Eye movements and their role in visual and cognitive processes*, ed. E. Kowler. Elsevier Science. [aZWP]
- Kreiman, G., Koch, C. & Itzhak, F. (2001) Imagery neurons in the human brain. *Nature* 408:357–61. [GEM]
- Kunen, S. & May, J. G. (1980) Spatial frequency content of visual imagery. *Perception and Psychophysics* 28(6):555–59. [aZWP]
- (1981) Imagery-induced McCollough effects: Real or imagined? *Perception and Psychophysics* 30(1):99–100. [aZWP]
- Kunzendorf, R. G., Spanos, N. P. & Wallace, B., eds. (1996) *Hypnosis and imagination*. Baywood. [aZWP]
- Ladd, G. T. (1894) *Psychology: Descriptive and explanatory*. Scribner’s. [JRP]
- Lairy, G. C., Barros de Ferreira, M. & Goldstein, L. (1968) Les phases intermédiaires du sommeil. In: *The abnormalities of sleep in man*, ed. H. Gastaut, E. Lugaresi, G. Berti Ceroni & G. Coccagna. Aulo Gaggi. [CG]
- Lamme, V. A. F., Rodríguez-Rodríguez, V. & Spekreijse, H. (1999) Separate processing dynamics for texture elements, boundaries, and surfaces in primary visual cortex of the macaque monkey. *Cerebral Cortex* 9:406–13. [SG]
- Lamme, V. A. F., Super, H. & Spekreijse, H. (1998) Feedforward, horizontal and feedback processing in the visual cortex. *Current Opinion in Neurobiology* 8:529–35. [SG]
- Landau, B. & Gleitman, L. R. (1985) *Language and experience: Evidence from the blind child*. Harvard University Press. [arZWP]
- Landau, P. & Schwartz, E. L. (1994) Subset warping: Rubber sheeting with cuts. *Computer Vision, Graphics and Image Processing* 56:247–66. [JP]
- Landis, T. (2000) Disruption of space perception due to cortical lesions. *Spatial Vision* 13(2–3):179–91. [rZWP]
- Lang, P. J. (1979) A bio-informational theory of emotional imagery. *Psychophysiology* 16:495–512. [MR]
- (1987) Imagery as action: A reply to Watts and Blackstock. *Cognition and Emotion* 1:407–26. [MR]
- Larson, J. D. & Foulkes, D. (1969) Electromyogram suppression during sleep, dream recall, and orientation time. *Psychophysiology* 5:548–55. [CG]
- Levesque, H. (1986) Making believers out of computers. *Artificial Intelligence* 30:81–108. [aZWP]
- Levesque, H. J. & Brachman, R. J. (1985) A fundamental tradeoff in knowledge representation and reasoning (revised version). In: *Readings in knowledge representation*, ed. H. J. Levesque & R. J. Brachman. Morgan Kaufmann. [arZWP]
- Lewis, D. (1971) Analog and digital. *Nous* 5(3):321–27. [rZWP]
- (1980) Veridical hallucination and prosthetic vision. *Australasian Journal of Philosophy* 58:239–49. [IG]

- Llinas, R. & Ribary, U. (1993) Coherent 40-Hz oscillation characterizes dream state in humans. *Proceedings of the National Academy of Science USA* 90:2078–81. [CG]
- Lövblad, K. O., Thomas, R., Jakob, P. M., Scammell, T., Bassetti, C., Griswold, M., Ives, J., Matheson, J., Edelman, R. R. & Warach, S. (1999) Silent functional magnetic resonance imaging demonstrates focal activation in rapid eye movement sleep. *Neurology* 53:2193–95. [CG]
- Luria, A. R. (1966) *Higher cortical functions in man*. Basic Books. [ENS]
- Lycan, W. (2000) Representational theories of consciousness. In: *Stanford encyclopedia of philosophy*, ed. E. N. Zalta. Stanford Center for the Study of Language and Information. [Online serial] <http://plato.stanford.edu/archives/win2001/entries/consciousness-representational/> [NJTT]
- Madsen, P. L., Holm, S., Vorstrup, S., Friberg, L., Lassen, N. A. & Wildschiodtz, G. (1991) Human regional cerebral blood flow during rapid-eye-movement sleep. *Journal of Cerebral Blood Flow Metabolism* 11:502–507. [CG]
- Maher, L., Chatterjee, A., Gonzales-Rothi, L. & Heilman, K. (1995) Agrammatic sentence production: The use of a temporal-spatial strategy. *Brain and Language* 49:105–24. [AC]
- Malebranche, N. (1712/1997) *The search after truth*, ed. and trans. T. M. Lennon & P. J. Olscamp. Cambridge University Press. [PPS]
- Marg, E. (1970) A neurologic approach to perceptual problems. In: *Early experience and visual information processing in perceptual and reading disorders*, ed. F. A. Young and D. B. Lindsley. National Academy of Science. [GEM]
- (1973) Recording from single cells in the human visual cortex. In: *Handbook of sensory physiology: Vol. VII/3B*, ed. R. Jung. Springer-Verlag. [GEM]
- Marg, E., Adams, J. E. & Rutkin, B. (1968) Receptive fields of cells in the human visual cortex. *Experience* 24:348–50. [GEM]
- Marks, D. F. (1973) Visual imagery differences in the recall of pictures. *British Journal of Psychology* 64(1):17–24. [aZWP]
- Marmor, G. S. & Zaback, L. A. (1976) Mental rotation by the blind: Does mental rotation depend on visual imagery? *Journal of Experimental Psychology: Human Perception and Performance* 2(4):515–21. [aZWP]
- Marr, D. (1982) *Vision: A computational investigation into the human representation and processing of visual information*. W. H. Freeman. [aZWP]
- Marr, D. & Nishihara, H. K. (1976) *Representation and recognition of spatial organization of three-dimensional shapes*. MIT A. I. Memo 377. [arZWP]
- (1978) Representation and recognition of spatial organization of three-dimensional shapes. *Proceedings of the Royal Society of London B* 200:269–94. [MJ, aZWP]
- Marshall, J. C. (2001) Auditory neglect and right parietal cortex. *Brain* 124(4):645–46. [rZWP]
- Mather, J. A. & Lackner, J. R. (1977) Adaptation to visual rearrangement: Role of sensory discordance. *Quarterly Journal of Experimental Psychology* 29(2):237–44. [aZWP]
- (1980) Visual tracking of active and passive movements of the hand. *Quarterly Journal of Experimental Psychology* 32(2):307–15. [aZWP]
- (1981) Adaptation to visual displacement: Contribution of proprioceptive, visual, and attentional factors. *Perception* 10(4):367–74. [aZWP]
- McCarley, R. W., Winkelman, J. W. & Dufy, F. H. (1983) Human cerebral potentials associated with REM sleep rapid eye movements: Link to PGO waves and waking potentials. *Brain Research* 274:359–64. [CG]
- McConkie, G. M. & Currie, C. B. (1996) Visual stability across saccades while viewing complex pictures. *Journal of Experimental Psychology: Human Perception and Performance* 22(3):563–81. [aZWP]
- McGlinchey-Berroth, R., Milberg, W. P., Verfaellie, M. & Grande, L. (1996) Semantic processing and orthographic specificity in hemispatial neglect. *Journal of Cognitive Neuroscience* 8(3):291–304. [rZWP]
- McManus, I. & Humphrey, N. (1973) Turning the left cheek. *Nature* 243:271–72. [AC]
- Meador, K. J., Loring, D. W., Bowers, D. & Heilman, K. M. (1987) Remote memory and neglect syndrome. *Neurology* 37:522–26. [PB]
- Mehta, Z. & Newcombe, F. (1996) Selective loss of verbal imagery. *Neuropsychologia* 34:441–47. [EdH]
- Mehta, Z., Newcombe, F. & De Haan, E. H. F. (1992) Selective loss of imagery in a case of visual agnosia. *Neuropsychologia* 30:645–55. [EdH]
- Mel, B. W. (1986) A connectionist learning model for 3-dimensional mental rotation, zoom, and pan. In: *Proceedings of the Eighth Annual Conference of the Cognitive Science Society*, pp. 562–71. Erlbaum. [NJTT]
- Mellet, E., Petit, L., Mazoyer, B., Denis, M. & Tzourio, N. (1998) Reopening the mental imagery debate: Lessons from functional anatomy. *Neuroimage* 8(2):129–39. [CG, aZWP]
- Mellet, E., Tzourio, N., Crivello, F., Joliot, M., Denis, M., & Mazoyer, B. (1996) Functional anatomy of spatial mental imagery generated from verbal instructions. *Journal of Neuroscience* 16(20):6504–12. [aZWP]
- Miashita, Y., Sakai, K., Higuchi, S.-I. & Masui, N. (1991) Localization of primal long-term memory in the primate temporal cortex. In: *Memory: Organization and locus of change*, ed. L. R. Squire, N. M. Weinberger, G. Lynch & L. McGough. Oxford University Press. [ENS]
- Millar, S. (1979) Utilization of shape and movement cues in simple spatial tasks by blind and sighted children. *Perception* 8:11–20. [SM]
- (1981) Self-referent and movement cues in coding spatial location by blind and sighted children. *Perception* 10:255–64. [SM]
- (1985) Movement cues and body orientation in recall of location by blind and sighted children. *Quarterly Journal of Experimental Psychology* 37:257–79. [SM]
- (1988) Models of sensory deprivation: The nature/nurture dichotomy and spatial representation in the blind. *International Journal of Behavioural Development* 11:69–87. [SM]
- (1994) *Understanding and representing space: Theory and evidence from studies with blind and sighted children*. Clarendon Press. [SM]
- Millar, S. & Al-Attar, Z. (2000) Vertical and bisection bias in active touch. *Perception* 29:481–500. [SM]
- (2001) Illusions in reading maps by touch: Reducing distance errors. *British Journal of Psychology* 92:643–57. [SM]
- (2002) Müller-Lyer illusions in touch and vision: Implications for multisensory processes. *Perception and Psychophysics* 64:353–65. [SM]
- Millar, S. & Ittyerah, M. (1991) Movement imagery in young and congenitally blind children: Mental practice without visuo-spatial information. *International Journal for the Study of Behavioral Development* 15(1):125–46. [SM]
- Millikan, R. G. (2000) *On clear and confused ideas: An essay about substance concepts*. Cambridge University Press. [KKN]
- Milner, A. D. & Goodale, M. A. (1995) *The visual brain in action*. Oxford University Press. [DI, aZWP]
- Milton, R. (1997) *Shattering the myths of Darwinism*. Park Street Press. [JAT]
- Minsky, M. L. (1975) A framework for representing knowledge. In: *The psychology of computer vision*, ed. P. H. Winston. McGraw-Hill. [rZWP]
- Minsky, M. L. & Papert, S. (1971) On some associative, parallel and analog computations. In: *Associative information techniques: Symposium at the General Motors Research Laboratories*, ed. E. L. Jacks. Elsevier. [rZWP]
- Mitchell, D. B. & Richman, C. L. (1980) Confirmed reservations: Mental travel. *Journal of Experimental Psychology: Human Perception and Performance* 6:58–66. [aZWP]
- Miyachi, S., Takino, R., Fukuda, H. & Torii, S. (1987) Electrophysiological evidence for dreaming: Human cerebral potentials associated with rapid eye movements during REM sleep. *Electroencephalography and Clinical Neurophysiology* 66:383–90. [CG]
- Moran, T. P. (1973) *The symbolic imagery hypothesis: A production system model*. Unpublished doctoral dissertation, Carnegie-Mellon University, Pittsburgh. (University Microfilms No. 74–14, 657.) [NJTT]
- Moyer, R. S. (1973) Comparing objects in memory: Evidence suggesting an internal psychophysics. *Perception and Psychophysics* 13:180–84. [WMP]
- Moyer, R. S. & Bayer, R. H. (1976) Mental comparison and the symbolic distance effect. *Cognitive Psychology* 8:228–46. [WMP]
- Moyer, R. S., Bradley, D. R., Sorensen, M. H., Whiting, J. C. & Mansfield, D. P. (1978) Psychophysical functions for perceived and remembered size. *Science* 200:330–32. [WMP]
- Moyer, R. S., Sklarew, P. & Whiting, J. C. (1982) Memory psychophysics. In: *Psychophysical judgment and the process of perception*, ed. H. G. Geissler & P. Petzold. VEB Deutscher Verlag der Wissenschaften. [WMP]
- Mueller, H. J. & Findlay, J. M. (1987) Sensitivity and criterion effects in the spatial cuing of visual attention. *Perception and Psychophysics* 42:383–99. [rZWP]
- Munzert, J. & Raab, M. (in preparation) Information processing in sports. *Encyclopedia for Psychology*. [MR]
- Nakajima, N. & Shimojo, S. (1981) Adaptation to the reversal of binocular depth cues: Effects of wearing left-right reversing spectacles on stereoscopic depth-perception. *Perception* 10:392–402. [EW]
- Neisser, U. (1976) *Cognition and reality*. W. H. Freeman. [NJTT]
- Newell, A. (1990) *Unified theories of cognition*. Harvard University Press. [aZWP]
- Newell, A. & Simon, H. A. (1976) Computer science as empirical inquiry. *Communications of the Association for Computing Machinery* 19(3):113–26. [aZWP]
- Nicod, J. (1970) *Geometry and induction*. University of California Press. [aZWP]
- Nielsen, T. (2000) Cognition in REM and NREM sleep: A review and possible reconciliation of two models of sleep mentation. *Behavioral and Brain Sciences* 23:851–66. [CG]
- Nijhawan, R. (1991) Three-dimensional Müller-Lyer Illusion. *Perception and Psychophysics* 49(9):333–41. [aZWP]
- (1994) Motion extrapolation in catching. *Nature* 370(6487):256–57. [aZWP]
- Ohkuma, Y. (1986) A comparison of image-induced and perceived Müller-Lyer illusion. *Journal of Mental Imagery* 10(4):31–38. [aZWP]
- O'Keefe, J. (1976) Place units in the hippocampus of the freely moving rat. *Experimental Neurology* 51(1):78–109. [NB]
- Olivetti Belardinelli, M., Del Gratta, C., Di Matteo, R., Ferretti, A. & Romani, G. L. (2001) Intermodal analysis of sensory image generation by means of fMRI. In: *Proceedings of the 8th European Workshop on Imagery and Cognition, Saint-Malo, France*, p. 84. Psychology Press. [MOB]

- O'Regan, J. K. (1992) Solving the "real" mysteries of visual perception: The world as an outside memory. *Canadian Journal of Psychology* 46:461–88. <http://nivea.psych.univ-paris5.fr/CanJ/CanJ.html> [arZWP, NJTT]
- O'Regan, J. K., Deubel, H., Clark, J. J. & Rensink, R. A. (2000) Picture changes during blinks: Looking without seeing and seeing without looking. *Visual Cognition* 7:191–212. [aZWP]
- O'Regan, J. K. & Lévy-Schoen, A. (1983) Integrating visual information from successive fixations: Does trans-saccadic fusion exist? *Vision Research* 23(8):765–68. [aZWP]
- O'Regan, J. K. & Noë, A. (2001) A sensorimotor account of vision and visual consciousness. *Behavioral and Brain Sciences* 24(5):939–73. [PB, arZWP, NJTT]
- Ortigue, S., Viaud-Delmon, I., Annoni, J. M., Landis, T., Michel, C., Blanke, O., Vuilleumier, P. & Mayer, E. (2001) Pure representational neglect after right thalamic lesion. *Annals of Neurology* 50(3):401–404. [PB]
- Oudejans, R. R. D., Michaels, C. F., Bakker, F. C. & Davids, K. (1999) Shedding some light on catching in the dark: Perceptual mechanisms for catching fly balls. *Journal of Experimental Psychology: Human Perception and Performance* 25(2):531–42. [MR]
- Paivio, A. (1971) *Imagery and verbal processes*. Holt, Reinhart, and Winston. [WMP, aZWP]
- (1975) Perceptual comparisons through the mind's eye. *Memory and Cognition* 3:635–47. [WMP]
- Palmer, S. E. (1977) Hierarchical structure in perceptual representation. *Cognitive Psychology* 9:441–74. [aZWP]
- (1978) Fundamental aspects of cognitive representation. In: *Cognition and categorization*, ed. E. Rosch and B. Lloyd. Erlbaum. [VG]
- Pani, J. R. (1996) Mental imagery as the adaptationist views it. *Consciousness and Cognition* 5:288–326. [JRP]
- Parsons, L. M., Gabrieli, J. D. E. & Corkin, S. (1987) Failure to improve a skill for mental rotation in global amnesia. Paper presented at the Annual Meeting of the Society for Neuroscience. [GEM]
- Pascual-Leone, A., Hamilton, R., Tormos, J. M., Keenan, J. P. & Catalá, M. D. (1999) Neuroplasticity in the adjustment to blindness. In: *Neuronal plasticity: Building a bridge from the laboratory to the clinic*, eds. J. Grafman & Y. Christen. Springer-Verlag.
- Patterson, J. & Deffenbacher, K. (1972) Haptic perception of the Mueller-Lyer illusion by the blind. *Perceptual and Motor Skills* 35(3):819–24. [aZWP]
- Pavani, F., Ladavas, E. & Driver, J. (2002) Selective deficit of auditory localisation in patients with visuospatial neglect. *Neuropsychologia* 40(3):291–301. [rZWP]
- Penfield, W. (1958) *The excitable cortex in conscious man*. Liverpool University Press. [ENS]
- Perky, C. W. (1910) An experimental study of imagination. *American Journal of Psychology* 21(3):422–52. [aZWP]
- Pessoa, L., Thompson, E. & Noë, A. (1998) Finding out about filling in: A guide to perceptual completion for visual science and the philosophy of perception. *Behavioral and Brain Sciences* 21(6):723–802. [arZWP]
- Peterson, M. A. (1993) The ambiguity of mental images: Insights regarding the structure of shape memory and its function in creativity. In: *Imagery, creativity, and discovery: A cognitive perspective. Advances in psychology, vol. 98*, ed. H. Roskos-Ewoldson, M. J. Intons-Peterson & R. E. Anderson. North Holland/Elsevier Science.
- Peterson, M. A., Kihlstrom, J. F., Rose, P. M. & Glisky, M. A. (1992) Mental images can be ambiguous: Reconstructions and reference-frame reversals. *Memory and Cognition* 20(2):107–23. [aZWP]
- Petrusic, W. M. (2001) Contextual effects and associative processes in comparative judgements with symbolic and perceptual stimuli. In: *Fechner Day 2001. Proceedings of the Sixteenth Annual Meeting of the International Society for Psychophysics*, pp. 75–80, ed. E. Sommerfeldt. The International Society for Psychophysics. [WMP]
- Petrusic, W. P., Baranski, J. V. & Aubin, P. H. (1998a) Comparing perceived and remembered magnitudes. In: *Fechner Day '98. Proceedings of the Fourteenth Annual Meeting of the International Society for Psychophysics*, pp. 44–49, ed. S. Grondin & Y. Lacouture. The International Society for Psychophysics. [WMP]
- Petrusic, W. P., Baranski, J. V. & Kennedy, R. (1998b) Similarity comparisons with remembered and perceived magnitudes: Memory psychophysics and fundamental measurement. *Memory and Cognition* 26:1041–55. [WMP]
- Piaget, J. (1970) *Genetic epistemology*. Columbia University Press. [EW]
- Pinker, S. (1980) Mental imagery and the third dimension. *Journal of Experimental Psychology: General* 109(3):354–71. [aZWP]
- (1984) Visual cognition: An introduction. *Cognition* 18:1–63. [GEM]
- Place, U. T. (1956) Is consciousness a brain process? *British Journal of Psychology*. Reprinted in: *The philosophy of mind: Classical problems/contemporary issues*, ed. B. Beakley & P. Ludlow. MIT Press. [PPS]
- Podgorny, P. & Shepard, R. N. (1978) Functional representations common to visual perception and imagination. *Journal of Experimental Psychology: Human Perception and Performance* 9:380–93. [JRP, aZWP]
- Poggio, T. & Edelman, S. (1990) A network that learns to recognize three-dimensional objects. *Nature* 343:263–66. [MJ]
- Pouget, A. & Sejnowski, T. J. (1999) A new view of hemineglect based on the response properties of parietal neurons. In: *The hippocampal and parietal foundations of spatial cognition*, ed. N. Burgess, K. J. Jeffery & J. O'Keefe. Oxford University Press. [NB]
- Predebon, J. & Wenderoth, P. (1985) Imagined stimuli: Imaginary effects? *Bulletin of the Psychonomic Society* 23(3):215–16. [aZWP]
- Pylyshyn, Z. W. (1973) What the mind's eye tells the mind's brain: A critique of mental imagery. *Psychological Bulletin* 80:1–24. [arZWP, EW]
- (1978) Imagery and Artificial Intelligence. In: *Perception and cognition: Issues in the foundations of psychology, vol. 9*, ed. C. W. Savage. University of Minnesota Press. [aZWP]
- (1979a) Do mental events have duration? *Behavioral and Brain Sciences* 2(2):277–78. [rZWP]
- (1979b) The rate of "mental rotation" of images: A test of a holistic analogue hypothesis. *Memory and Cognition* 7:19–28. [aZWP]
- (1979c) Validating computational models: A critique of Anderson's indeterminacy of representation claim. *Psychological Review* 86(4):383–94. [rZWP]
- (1980) Cognitive representation and the process-architecture distinction. *Behavioral and Brain Sciences* 3(1):154–69. [aZWP]
- (1981) The imagery debate: Analogue media versus tacit knowledge. *Psychological Review* 88:16–45. [SMK, arZWP]
- (1984) *Computation and cognition: Toward a foundation for cognitive science*. MIT Press. [GDB, KKN, arZWP]
- (1989) The role of location indexes in spatial perception: A sketch of the FINST spatial-index model. *Cognition* 32:65–97. [rZWP]
- (1991a) The role of cognitive architectures in theories of cognition. In: *Architectures for intelligence*, ed. K. VanLehn. Erlbaum. [aZWP]
- (1991b) Rules and representation: Chomsky and representational realism. In: *The Chomskian turn*, ed. A. Kashir. Basil Blackwell. [aZWP]
- (1994a) Mental pictures on the brain: Review of "Image and brain: The resolution of the imagery debate" by Stephen Kosslyn. *Nature* 372(17):289–90. [aZWP]
- (1994b) Some primitive mechanisms of spatial attention. *Cognition* 50:363–84. [aZWP]
- (1996) The study of cognitive architecture. In: *Mind matters: Contributions to cognitive science in honor of Allen Newell*, ed. D. Steier & T. Mitchell. Erlbaum. [aZWP]
- (1998) Visual indexes in spatial vision and imagery. In: *Visual attention*, ed. R. D. Wright. Oxford University Press. [arZWP]
- (1999) Is vision continuous with cognition? The case for cognitive impenetrability of visual perception. *Behavioral and Brain Sciences* 22(3):341–423. [VG, aZWP, MR]
- (2000) Situating vision in the world. *Trends in Cognitive Sciences* 4(5):197–207. [arZWP]
- (2001a) Connecting vision and the world: Tracking the missing link. In: *The foundations of cognitive science*, ed. J. Branquinho. Clarendon Press. [rZWP]
- (2001b) Seeing, acting, and knowing: Commentary on O'Regan & Noë. *Behavioral and Brain Sciences* 24(5):999. [rZWP]
- (2001c) Visual indexes, preconceptual objects, and situated vision. *Cognition* 80(1/2):127–58. [arZWP]
- (forthcoming) *Seeing and visualizing: It's not what you think*. MIT Press/Bradford Books. [arZWP]
- Pylyshyn, Z. W. & Cohen, J. (1999) Imagined extrapolation of uniform motion is not continuous. Paper presented at the Annual Conference of the Association for Research in Vision and Ophthalmology, Ft. Lauderdale, FL, May 1999. [aZWP]
- Raizada, R. D. S. & Grossberg, S. (2001) Context-sensitive binding by the laminar circuits of V1 and V2: A unified model of perceptual grouping, attention, and orientation contrast. *Visual Cognition* 8:431–66. [SG]
- Ranney, M. (1989) Internally represented forces may be cognitively penetrable: Comment on Freyd, Pantzer, and Cheng (1988). *Journal of Experimental Psychology: General* 118(4):399–402. [aZWP]
- Rao, S. C., Raier, G. & Miller, E. K. (1999) Integration of "what" and "where" in the primate prefrontal cortex. *Science* 276:821–24. [DI]
- Reed, S. K., Hock, H. S. & Lockhead, G. R. (1983) Tacit knowledge and the effect of pattern configuration on mental scanning. *Memory and Cognition* 11:137–43. [aZWP]
- Rehkaemper, K. (1991) *Sind mentale Bilder bildhaft? – Eine Frage zwischen Philosophie und Wissenschaft*. Doctoral dissertation, University of Hamburg, Germany. [VG]
- Reisberg, D. & Chambers, D. (1991) Neither pictures nor propositions: What can we learn from a mental image? *Canadian Journal of Psychology* 45(3):336–52. [aZWP]

- Reisberg, D. & Morris, A. (1985) Images contain what the imager put there: A nonreplication of illusions in imagery. *Bulletin of the Psychonomic Society* 23(6):493–96. [aZWP]
- Rensink, R. A. (2000a) The dynamic representation of scenes. *Visual Cognition* 7:17–42. [aZWP]
- (2000b) Visual search for change: A probe into the nature of attentional processing. *Visual Cognition* 7:345–76. [aZWP]
- Rensink, R. A., O'Regan, J. K. & Clark, J. J. (1997) To see or not to see: The need for attention to perceive changes in scenes. *Psychological Science* 8(5):368–73. [aZWP]
- (2000) On the failure to detect changes in scenes across brief interruptions. *Visual Cognition* 7:127–45. [aZWP]
- Rentschler, I., Caelli, T., Bischof, W. & Jüttner, M. (2000) Object recognition and image understanding: Theories of everything? *Spatial Vision* 13:129–35. [MJ]
- Rentschler, I., Jüttner, M., Osman, E., Miller, A. & Caelli, T. (submitted) Haptic exploration reinforces visual category learning for 3-D objects. [MJ]
- Rey, G. (1981) What are mental images? In: *Readings in the philosophy of psychology, vol. II*, ed. N. Block. Harvard University Press. [aZWP]
- Reynolds, J., Chelazzi, L. & Desimone, R. (1999) Competitive mechanisms subserve attention in macaque areas V2 and V4. *Journal of Neuroscience* 19:1736–53. [SG]
- Richman, C. L., Mitchell, D. B. & Reznick, J. S. (1979) Mental travel: Some reservations. *Journal of Experimental Psychology: Human Perception and Performance* 5:13–18. [aZWP]
- Richter, W., Somorjai, R., Summers, R., Jarmasz, M., Menon, R. S., Gati, J. S., Georgopoulos, A. P., Tegeler, C., Ugurbil, K. & Kim, S.-G. (2000) Motor area activity during mental rotation studied by time-resolved single-trial fMRI. *Journal of Cognitive Neuroscience* 12(2):310–20. [rZWP]
- Rieke, F., Warland, D., de Ruyter van Steveninck, R. & Bialek, W. (1998) *Spikes: Exploring the neural code*. MIT Press. [JP]
- Rode, G. & Perenin, M. T. (1994) Temporary remission of representational hemineglect through vestibular stimulation. *NeuroReport* 5:869–72. [PB]
- Rode, G., Rossetti, Y. & Boisson, D. (2001) Prism adaptation improves representational neglect. *Neuropsychologia* 39(11):1250–54. [PB, rZWP]
- Roffwarg, H. P., Dement, W. C., Muzio, J. N. & Fisher, C. (1962) Dream imagery: Relationship to rapid eye movement sleep. *Archives of General Psychiatry* 7:235–58. [CG]
- Royer, A. S. & Schwartz, E. L. (1990) Design consideration for a space-variant sensor with complex-logarithmic geometry. In: *Proceedings of the Tenth International Conference on Pattern Recognition, Atlantic City, NJ, vol. 2*, pp. 278–85. [JP]
- Roland, P. E. & Gulyas, B. (1994a) Beyond 'pet' methodologies to converging evidence: Reply. *Trends in Neurosciences* 17(12):515–16. [aZWP]
- (1994b) Visual imagery and visual representation. *Trends in Neurosciences* 17(7):281–87. [aZWP]
- (1995) Visual memory, visual imagery, and visual recognition of large field patterns by the human brain: Functional anatomy by positron emission tomography. *Cerebral Cortex* 5(1):79–93. [aZWP]
- Rolls, E. T. (2000) Memory systems in the brain. *Annual Review of Psychology* 51:599–630. [EdH]
- Romero, S. G., Manly, C. F. & Grafman, J. (in press) Investigating cognitive neuroplasticity in single cases: Lessons learned from applying functional neuroimaging techniques to the traditional neuropsychological case study framework. *Neurocase*, in press.
- Rorty, R. (1979) *Philosophy and the mirror of nature*. Princeton University Press. [PPS]
- Rossetti, Y., Rode, G., Pisella, L., Farnè, A., Li, L., Boisson, D. & Perenin, M. T. (1998) Prism adaptation to a rightward optical deviation rehabilitates left hemispatial neglect. *Nature* 395:166–69. [PB]
- Rouw, R., Kosslyn, S. M. & Hamell, R. (1997) Detecting high-level and low-level properties in percepts and mental images. *Cognition* 63:209–26. [SMK]
- (1998) Aspects of mental images: Is it possible to get the picture? *Cognition* 66:103–107. [SMK]
- Ryle, G. (1968) A puzzling element in the notion of thinking. In: *Studies in the philosophy of thought and action*, ed. P. F. Strawson. Oxford University Press. [PPS]
- Sadato, N., Pascual-Leone, A., Grafman, J., Ibañez, V., Deiber, M.-P., Dold, G. & Hallett, M. (1996) Activation of the primary visual cortex by Braille reading in blind subjects. *Nature* 380:526–528.
- Sarbin, T. R. & Juhasz, J. B. (1970) Toward a theory of imagination. *Journal of Personality* 38:52–76. [NJTT]
- Sartre, J.-P. (1940) *L'imaginaire*. Gallimard. [GDB]
- (1948) *The psychology of imagination*, trans. Bernard Frechtman. Washington Square Press/Philosophical Library. (Original French publication, 1940). [MEA, NJTT]
- Satoh, T. (1971) Direct cortical response and PGO spike during paradoxical sleep of the cat. *Brain Research* 28:576–78. [CG]
- Schiller, P. H. (1994) Area V4 of the primate visual cortex. *Current Directions in Psychological Science* 3:89–92. [SG]
- Schlaegel, T. F. (1953) The dominant method of imagery in blind compared to sighted adolescents. *Journal of Genetic Psychology* 83:265–77. [SM]
- Schwartz, E. L. (1980) Computational anatomy and functional architecture of striate cortex: A spatial mapping approach to perceptual coding. *Vision Research* 20:645–69. [JP]
- (1994) Computational studies of the spatial architecture of primate visual cortex: Columns, maps, and protomaps. In: *Primary visual cortex in primates. Cerebral cortex, vol. 10*, ed. A. Peters & K. Rocklind. Plenum Press. [JP]
- Schwartz, E. L. & Rojer, A. S. (1991) Cortical hypercolumns and the topology of random orientation maps. Technical Report 593, Courant Institute of Mathematical Sciences, New York University. [JP]
- Schweinberger, S. R. & Stief, V. (2001) Implicit perception in patients with visual neglect: Lexical specificity in repetition priming. *Neuropsychologia* 39(4):420–29. [rZWP]
- Segal, S. J. & Fusella, V. (1969) Effects of imaging on signal-to-noise ratio, with varying signal conditions. *British Journal of Psychology* 60(4):459–64. [arZWP]
- (1970) Influence of imaged pictures and sounds on detection of visual and auditory signals. *Journal of Experimental Psychology* 83(3):458–64. [arZWP]
- Sellars, R. W. (1932) *The philosophy of physical realism*. Macmillan. [EW]
- Sereno, M. I., Dale, A. M., Reppas, J. B., Kwong, K. K., Belliveau, J. W., Brady, T. J., Rosen, B. R. & Tootell, R. B. (1995) Borders of multiple visual areas in humans revealed by functional magnetic resonance imaging. *Science* 268:889–93. [SMK]
- Sereno, M. I., Pitzalis, S. & Martinez, A. (2001) Mapping of contralateral space in retinotopic coordinates by a parietal cortical area in humans. *Science* 294:1350–56. [DI]
- Servos, P. & Goodale, M. A. (1995) Preserved visual imagery in visual form agnosia. *Neuropsychologia* 33:883–94. [aZWP]
- Shaki, S. & Algom, D. (2002) The locus and nature of semantic congruity in symbolic comparison: Evidence from the Stroop effect. *Memory and Cognition* 30:3–17. [WMP]
- Sheehan, P. W. (1967) A shortened form of Betts' questionnaire upon mental imagery. *Journal of Clinical Psychology* 23:386–89. [MOB]
- Sheingold, K. & Tenney, Y. J. (1982) Memory for a salient childhood event. In: *Memory observed*, ed. U. Neisser. W. H. Freeman. [aZWP]
- Shepard, R. N. (1964) Attention and the metrical structure of the similarity space. *Journal of Mathematical Psychology* 1:54–87. [rZWP]
- Shepard, R. N. & Chipman, S. (1970) Second-order isomorphism of internal representations: Shapes of states. *Cognitive Psychology* 1:1–17. [HA, rZWP]
- Shepard, R. N. & Feng, C. (1972) A chronometric study of mental paper folding. *Cognitive Psychology* 3:228–43. [JRP, aZWP]
- Shepard, R. N. & Metzler, J. (1971) Mental rotation of three-dimensional objects. *Science* 171:701–703. [NB, aZWP]
- Shuren, J. E., Brott, T. G., Schefft, B. K. & Houston, W. (1996) Preserved color imagery in an achromatopsic. *Neuropsychologia* 34(6):485–89. [aZWP]
- Silbersweig, D. A. & Stern, E. (1998) Towards a functional neuroanatomy of conscious perception and its modulation by volition: Implications of human auditory neuroimaging studies. *Philosophical Transactions of the Royal Society of London B: Biological Sciences* 353(1377):1883–88. [aZWP]
- Sillito, A. M., Jones, H. E., Gerstein, G. L. & West, D. C. (1994) Feature-linked synchronization of thalamic relay cell firing induced by feedback from the visual cortex. *Nature* 369:479–82. [SG]
- Simons, D. J. (1996) In sight, out of mind: When object representations fail. *Psychological Science* 7(5):301–305. [aZWP]
- Simons, D. J. & Levin, D. T. (1997) Change blindness. *Trends in Cognitive Sciences* 1:261–67. [aZWP]
- Simons, D. J. & Wang, R. F. (1998) Perceiving real-world viewpoint changes. *Psychological Science* 9:315–20. [NB]
- Skinner, B. F. (1948) "Superstition" in the pigeon. *Journal of Experimental Psychology* 38:168–72. [FG]
- (1957) *Verbal behavior*. Appleton-Century-Crofts. [EW]
- (1963) Behaviorism at fifty. *Science* 140:951–58. Reprinted in: *The selection of behavior: The operant behaviorism of B. F. Skinner*, ed. A. C. Catania & S. Harnad. Cambridge University Press. [PPS]
- Slezak, P. (1991) Can images be rotated and inspected? A test of the pictorial medium theory. Paper presented at the Thirteenth Annual meeting of the Cognitive Science Society. [aZWP]
- (1992) *When can images be reinterpreted: Non-chronometric tests of pictorialism*. Paper presented at the Fourteenth Conference of the Cognitive Science Society. [aZWP]
- (1995) The 'philosophical' case against visual imagery. In: *Perspective on cognitive science: Theories, experiments and foundations*, ed. P. Slezak & T. Caelli & R. Clark. Ablex. [aZWP]
- (2002a) The tripartite model of representation. *Philosophical Psychology* 15(3) (in press). [aZWP]

- (2002b) The world gone wrong? Images, illusions, mistakes and misrepresentations. In: *Representation in mind*, ed. P. Staines, H. Clapin & P. Slezak. Praeger/Greenwood (forthcoming). [PPS]
- (2002c) Thinking about thinking: Language, thought and introspection. *Language and Communication* 22:353–73. [PPS]
- Slovan, A. (1971) Interactions between philosophy and artificial intelligence: The role of intuition and non-logical reasoning in intelligence. *Artificial Intelligence* 2:209–25. [aZWP]
- Smania, N., Bazoli, F., Piva, D. & Guidetti, G. (1997) Visuomotor imagery and rehabilitation of neglect. *Archives of Physical Medicine and Rehabilitation* 78(4):430–36. [PB]
- Snyder, L. H., Grieve, K. L., Brotchie, P. & Andersen, R. A. (1998) Separate body- and world-referenced representations of visual space in parietal cortex. *Nature* 394(6696):887–91. [NB]
- Sokolov, E. N. & Boucsein, W. (2000) A psychophysiological model of emotion space. *Integrative Physiological and Behavioral Science* 35:81–119. [ENS]
- Sparing, R., Mottaghy, F. M., Ganis, G., Thompson, W. L., Töpper, R., Kosslyn, S. M. & Pascual-Leone, A. (2002) Visual cortex excitability increases during visual imagery – a TMS study in healthy human subjects. *Brain Research* 938:92–97. [SMK]
- Sperling, G. (1960) The information available in brief visual presentations. *Psychological Monographs* 74 (Whole No. 11). [rZWP]
- Squire, L. R. & Slater, P. C. (1975) Forgetting in very long-term memory as assessed by an improved questionnaire technique. *Journal of Experimental Psychology: Human Perception and Performance* 10:4:50–54. [aZWP]
- Steinbach, M. J. (1976) Pursuing the perceptual rather than the retinal stimulus. *Vision Research* 16:1371–76. [aZWP]
- Steriade, M., Paré, D., Bouhassira, D., Deschênes, M. & Oakson, G. (1989) Phasic activation of lateral geniculate and pregeniculate thalamic neurons during sleep with ponto-geniculo-occipital waves. *Journal of Neuroscience* 9:2215–29. [CG]
- Stoerig, P. (1996) Varieties of vision: From blind responses to conscious recognition. *Trends in Neurosciences* 19(9):401–406. [aZWP]
- Stucki, D. J. & Pollack, J. B. (1992) Fractal (reconstructive analogue) memory. In: *Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society, Bloomington, IN*. Erlbaum. [NJTT]
- Takeuchi, T., Miyasita, A., Inugami, M. & Yamamoto, Y. (2001) Intrinsic dreams are not produced without REM sleep mechanisms: Evidence through elicitation of sleep onset REM periods. *Journal of Sleep Research* 10:43–52. [CG]
- Talmy, L. (2000) *Towards a cognitive semantics: Concept structuring systems*. MIT Press. [AC]
- Tarr, M. J. & Pinker, S. (1989) Mental rotation and orientation-dependence in shape recognition. *Cognitive Psychology* 21:233–82. [MJ]
- Tarr, M. J., Williams, P., Hayward, W. G. & Gauthier, I. (1998) Three-dimensional object recognition is viewpoint dependent. *Nature Neuroscience* 1:275–77. [MJ]
- Taube, J. S., Muller, R. U. & Ranck, J. B., Jr. (1990) Head-direction cells recorded from the postsubiculum in freely moving rats. I. Description and quantitative analysis. *Journal of Neuroscience* 10(2):420–35. [NB]
- Taylor, M. M. (1961) Effect of anchoring and distance perception on the reproduction of forms. *Perception and Motor Skills* 12:203–30. [aZWP]
- Thomas, N. J. T. (1999) Are theories of imagery theories of imagination? An active perception approach to conscious mental content. *Cognitive Science* 23(2):207–45. <http://www.calstatela.edu/faculty/nthomas/im-im/im-im.htm> [PB, aZWP, NJTT]
- (2001) Color realism: Toward a solution to the “hard problem.” *Consciousness and Cognition* 10:140–45. <http://www.calstatela.edu/faculty/nthomas/col-real.htm> [NJTT]
- (in press) Mental imagery, philosophical issues about. In: *Encyclopedia of cognitive science*, ed. L. Nadel. Macmillan/Nature Publishing. <http://www.calstatela.edu/faculty/nthomas/mipia.htm> [rZWP, NJTT]
- Thomas, W. L. & Kosslyn, S. M. (2000) Neural systems activated during visual mental imagery: A review and meta-analysis. In: *Brain mapping II: The systems*, ed. A. W. Toga & J. C. Mazziotta. Academic Press. [SMK]
- Tipper, S. P. & Behrmann, M. (1996) Object-centered not scene-based visual neglect. *Journal of Experimental Psychology: Human Perception and Performance* 22(5):1261–78. [rZWP]
- Titchener, E. B. (1909) *Lectures on the experimental psychology of thought*. Macmillan. [SM]
- Tlauka, M. & McKenna, F. P. (1998) Mental imagery yields stimulus-response compatibility. *Acta Psychologica* 98(1):67–79. [aZWP]
- Tootell, R. B. H., Silverman, M. S., Switkes, E. & de Valois, R. L. (1982) Deoxyglucose analysis of retinotopic organization in primate striate cortex. *Science* 218(4575):902–904. [SMK, arZWP]
- Toth, J. A. (in preparation) Research ethics, method, and best practices in the Department of Defense. Technical Report, Institute for Defense Analysis, Alexandria, VA, USA. [JAT]
- Tracy, R. L. & Tracy, L. N. (1974) Reports of mental activity from sleep stages 2 and 4. *Perceptual and Motor Skills* 38:647–48. [CG]
- Tranel, D., Damasio, H. & Damasio, A. R. (1997) A neural basis for the retrieval of conceptual knowledge. *Neuropsychologia* 35:1319–27. [EdH]
- Treisman, A. M. & Schmidt, H. (1982) Illusory conjunctions in the perception of objects. *Cognitive Psychology* 14:107–41. [MEA]
- Trojano, L., Grossi, D., Linden, D. E. J., Formisano, E., Hacker, H., Zanella, F. E., Goebel, R. & Di Salle, F. (2000) Matching two imagined clocks: The functional anatomy of spatial analysis in the absence of visual stimulation. *Cerebral Cortex* 10:473–81. [EdH]
- Tye, M. (1991) *The imagery debate*. MIT Press. [SMK, NJTT]
- Uhlarik, J. J. (1973) Role of cognitive factors on adaptation to prismatic displacement. *Journal of Experimental Psychology* 98:223–32. [aZWP]
- Ungerleider, L. G. & Mishkin, M. (1982) Two cortical systems. In: *Analysis of visual behavior*, ed. D. Ingle, M. A. Goodale & R. Mansfield. MIT Press. [DI]
- Vallar, G. (1993) The anatomical basis of spatial hemineglect in humans. In: *Unilateral neglect: Clinical and experimental studies*, ed. I. H. Robertson & J. C. Marshall. Erlbaum. [PB]
- Van der Velde, F. (1997) On the use of computation in modelling behavior. *Network: Computation in Neural Systems* 8:1–32. [FvdV]
- Van der Velde, F. & de Kamps, M. (2001) From knowing what to knowing where: Modeling object-based attention with feedback disinhibition of activation. *Journal of Cognitive Neuroscience* 13(4):479–91. [FvdV]
- Van Essen, D. C., Lewis, J. W., Drury, H. A., Hadjikhani, N., Tootell, R. B., Bakircioglu, M. & Miller, M. I. (2001) Mapping visual cortex in monkeys and humans using surface-based atlases. *Vision Research* 41:1359–78. [SMK]
- Vanni-Mercier, G., Debilly, G., Lin, J. S. & Pelissou, D. (1996) The caudo ventral pontine tegmentum is involved in the generation of high velocity eye saccades in bursts during paradoxical sleep. *Neuroscience Letters* 213:127–31. [CG]
- Virsu, V. (1971) Tendencies to eye movement, and misperception of curvature, direction, and length. *Perception and Psychophysics* 9(1-B):65–72. [aZWP]
- Vuilleumier, P. & Rafal, R. (1999) “Both” means more than “two”: Localizing and counting in patients with visuospatial neglect. *Nature Neuroscience* 2(9):783–84. [rZWP]
- Wallace, B. (1984a) Apparent equivalence between perception and imagery in the production of various visual illusions. *Memory and Cognition* 12(2):156–62. [aZWP]
- (1984b) Creation of the horizontal-vertical illusion through imagery. *Bulletin of the Psychonomic Society* 22(1):9–11. [aZWP]
- Wang, R. F. & Simons, D. J. (1999) Active and passive scene recognition across views. *Cognition* 70(2):191–210. [NB]
- Watanabe, K. & Shimojo, S. (1998) Attentional modulation in perception of visual motion events. *Perception* 27(9):1041–54. [aZWP]
- Weinberg, R. S. & Gould, D. (1995) *Foundations of sport psychology*. Human Kinetics. [MR]
- Wexler, M., Kosslyn, S. M. & Berthoz, A. (1998) Motor processes in mental rotation. *Cognition* 68(1):77–94. [rZWP]
- Wiener, N. (1958) *Nonlinear problems in random theory*. Wiley. [FG]
- Wiesel, T. N. & Hubel, D. H. (1963) Single-cell responses in striate cortex of kittens deprived of vision in one eye. *Journal of Neurophysiology* 26:1003–17. [RN]
- Wilson, C. L., Babb, T. L., Halgren, E. & Crandall, P. H. (1983) Visual receptive fields and response properties of neurons in human temporal lobe and visual pathways. *Brain* 106:473–502. [GEM]
- Wilson, M. A. & McNaughton, B. L. (1993) Dynamics of the hippocampal ensemble code for space. *Science* 261(5124):1055–58. [NB]
- Wittgenstein, L. (1953) *Philosophical investigations* [Philosophische Untersuchungen]. Basil Blackwell. [aZWP, PPS]
- Wood, R. & Schwartz, E. L. (1999) Topographic shear and the relationship of ocular dominance columns to orientation columns in monkey and cat visual cortex. *Neural Networks* 12:205–10. [JP]
- Wraga, M., Creem, S. H. & Proffitt, D. R. (2000) Updating displays after imagined object and viewer rotations. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 26(1):151–68. [NB]
- Wright, E. L. (1983) Inspecting images. *Philosophy* 58(223):57–72. [EW]
- (1992) The Entity Fallacy in epistemology. *Philosophy* 67(259) 33–50. [EW]
- (1996) What it isn't like. *American Philosophical Quarterly* 33(1):23–42. [EW]
- Wyatt, H. J. & Pola, J. (1979) The role of perceived motion in smooth pursuit eye movements. *Vision Research* 19:613–18. [aZWP]
- Yeshurun, Y. & Schwartz, E. L. (1989) Cepstral filtering on a columnar image architecture: A fast algorithm for binocular stereo segmentation. *IEEE Transactions: Pattern Analysis and Machine Intelligence* 11(7):759–67. [JP]
- Young, A. W., Humphreys, G. W., Riddoch, M. J., Hellawell, D. J. & de Haan, E. H. F. (1994) Recognition impairments and face imagery. *Neuropsychologia* 32:693–702. [EdH]
- Young, M. P. (2000) The architecture of visual cortex and inferential processes in vision. *Spatial Vision* 13(2–3):137–46. [rZWP]
- Zald, D. H. & Pardo, J. V. (2000) Functional neuroimaging of the olfactory system in humans. *International Journal of Psychophysiology* 36:165–81. [MOB]

- Zatorre, R. J., Halpern, A. R., Perry, D. W., Meyer, E. & Evans, A. C. (1996) Hearing in the mind's ear: A PET investigation of musical imagery and perception. *Journal of Cognitive Neuroscience* 8(1):29–46. [MOB]
- Zeki, S. (1991) Cerebral akinetopsia (visual motion blindness): A review. *Brain* 114(Pt 2):811–24. [rZWP]
- Zhang, K. (1996) Representation of spatial orientation by the intrinsic dynamics of the head-direction cell ensemble: A theory. *Journal of Neuroscience* 16(6):2112–26. [NB]
- Zhou, H. & May, J. G. (1993) Effects of spatial filtering and lack of effects of visual imagery on pattern-contingent color aftereffects. *Perception and Psychophysics* 53.:145–49. [aZWP]
- Zimler, J. & Keenan, J. M. (1983) Imagery in the congenitally blind: How visual are visual images? *Journal of Experimental Psychology: Learning, Memory, and Cognition* 9(2):269–82. [aZWP]