

VISUAL IMAGERY—VISUAL FORMAT OR VISUAL CONTENT?

Abstract

It is clear that visual imagery is somehow significantly visual. Some theorists, like Kosslyn, claim that the visual nature of visualisations derives from features of the neural processes which underlie those episodes. Pylyshyn claims, however, that it may merely reflect special features of the contents which we grasp when we visualise things. This paper discusses and rejects Pylyshyn's own attempts to identify the respects in which the contents of visualisations are notably visual. It then offers a novel and very different account of what is distinctively sensory about the contents of sensory images. The paper's alternative account is used in explaining various pieces of phenomenological and behavioural data concerning visualisation. Finally, it is tentatively suggested that the proposed account of the contents of sensory images may also shed light upon some of the neurological data involving visualisation and sensory imagery more generally.

1. Introduction

The burst of experimental studies of mental imagery over the last few decades has led to a host of debates about the relationships between visualisation¹ and vision. Kosslyn and

Many thanks to Rosanna Keefe for helpful discussion of the ideas in this paper; and many thanks to two anonymous referees for this journal for their sharp and helpful comments.

Address for correspondence: Department of Philosophy, University of Sheffield, 45 Victoria Street, Sheffield, S3 7QB.

Email: d.gregory@sheffield.ac.uk

his associates have argued, for instance, that visual images are realised in our brains using a roughly pictorial format which exploits specifically visual neurological resources (e.g. Kosslyn, 1980 and 1981; Kosslyn et al, 1979; Kosslyn et al, 2003), provoking Pylyshyn's vigorous dissent (e.g. Pylyshyn, 1981, 2002, 2003 and 2006).² But, even if one is wary of drawing conclusions about the format of the mental representations which feature in visualisations, one might suspect that there is at least something especially visual about the contents which we grasp during those episodes.

So, Pylyshyn notes that '[i]mage representations contain information about the [visual] appearance of things', that '[i]mage representations contain information about the relative location of things' and that the information which image representations contain is logically simple in certain respects ('[i]mage representations lack explicit quantifiers', for example) (Pylyshyn, 2006, pp. 423 – 4). Analogues of those points hold true for pictures. A snapshot of some people having a picnic, say, will tell us about the visual appearances and relative locations of the diners, but the photo doesn't do that using counterparts of 'all' and 'some'. That the noted points hold true of pictures and visualisations seems, too, to be connected to the particularly visual nature of what is expressed by both pictures and visual images: the remarked facts mirror the way in which

¹ Visual images feature in many sorts of mental episodes. In a particularly simple sort of case, we might use a visual image in just picturing how things look from somewhere; but, in a more complex type of instance, we might produce a visual image and then make the supposition that the image captures how things look to, say, a brain in a vat. Visualisations, as we ordinarily think of them, as they are generally assumed to be in the philosophical, psychological and neurological literature, and as I will assume them to be below, correspond to those especially simple cases in which we use visual images merely to show how things look from somewhere. (See Gregory, Forthcoming a for discussion of the relationships between visualisations of the type discussed below and another interesting range of cases, in which we imagine visual sensations.)

² In discussing the relationships between visual images and pictures, people have tended to ignore some kinds of pictures that do not seem to have that much in common with visual images at all, such as certain of the characteristic types of pictures which are produced by young children; I will do likewise.

our visual sensations tell us about visual appearances and relative locations, in a logically simple manner.

Some hard questions face the thought that especially visual contents are grasped during visualisations and picture-viewings, however. What is ‘visual’ about those contents? Do visualisations stand to picture-viewings as they also stand to typical verbal descriptions about what things look like, say, or are the contents of pictures and visual images more strongly visual than that? This paper will answer those questions and numerous others, by identifying a broad class of distinctively sensory representational contents. The members of the relevant class naturally divide up into further groups of contents which are tied to specific sensory modalities; and the resulting visual group of contents includes those which belong to visual images and pictures.

As we will see, the ideas elaborated below identify one clear way in which the particularly visual nature of visualisations does indeed derive from the nature of the contents which belong to visual images. The resulting position thus sheds light on our intuitive sense that visualisations are somehow like picture viewings, while yet helping us to understand the sources of various introspectively apparent limits on that natural analogy. And it can also be used in accounting for many of the experimental results concerning visual imagery which are contained within the psychological literature, as well as providing an interesting slant on the accumulating pile of relevant neurological data.

So, the framework developed below will largely concern the contents which we grasp during visualisations. But philosophical and psychological researchers on visual imagery have tended to be exercised by questions about the format of the mental

representations which figure in visualisations; by, for example, the question whether those representations marshal specifically visual neural resources. Hence it might be suspected that, even if the semantic framework that I am about to develop is a sturdy one, it will nonetheless be irrelevant to the concerns which have driven recent debates over visual imagery.

Recent work on visualisation has indeed revolved around theories that are quite different to the one that will be articulated below. In particular, though, it has commonly been claimed that various pieces of data are best to be explained by citing hypotheses relating to the format of visual images. An issue that is therefore very relevant to the recent debates is whether alternative explanations of the data are available, ones which do not employ assumptions about the format of visual images but rather assumptions about their contents. (Pylyshyn has repeatedly emphasised this point; see, for example, his remarks about the ‘null hypothesis’ strategy (Pylyshyn, 2006, p. 334).)

Yet if we are to address that issue, we need to have both an adequate account of the nature of the contents of visual images and a decent sense of that semantic account’s explanatory powers. Those two things are part of what this paper aims to provide. More generally, the task of articulating a satisfying theory of the contents which we grasp during visualisations is a crucial one because, unless it is undertaken, we will not have a sufficiently clear grasp on whether various aspects of visualisations arise from features of the *format* of the mental representations which are involved in those episodes or from mere features of those representations’ *contents*.

The thought that assumptions about the contents of visual images can do a lot of the explanatory work which people have instead sought to do using hypotheses about

their format is not a new one; it has long been a central part of Pylyshyn's views. But as we will see in the next two sections, while Pylyshyn's ideas illustrate the promise borne by the suggestion that visual images have especially visual contents, they also indicate difficulties which will need to be overcome if that suggestion is to be articulated properly.

2. Pylyshyn on Visualisation

Verbal representations can characterise what they represent in all sorts of ways. They can describe visible features of things, for instance, but they can also specify merely audible properties, as well as ones whose presence cannot be registered using any of our senses. The representational powers of visual images are more narrowly circumscribed; they can only represent things by 'capturing visual appearances', as we might naturally say.³ And a similar point applies to pictures.⁴

Pylyshyn 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' (Pylyshyn, 2002, p. 158). But visual images seem to be especially *visual*; and the characteristically visual specialness of visual images also belongs to pictures.

Pylyshyn's remarks therefore suggest that the especially visual nature of the contents of

³ The idea that the contents of visual images are somehow bound to visual appearances is a commonplace in the literature. See e.g. Finke, 1989, p. 2; Kosslyn, 1981, p. 216; Pylyshyn, 2002, p. 158; and Tye, 1991, pp. 90 – 1.

⁴ Alberti famously remarked, for example, that 'the painter is concerned solely with representing what can be seen' (Alberti, 1436/1966, p. 43), while Lessing stated that 'bodies with their visible properties are the peculiar subjects of painting' (Lessing, 1766/2005), p. 91); Picasso, too, once spoke of how certain painters, poisoned by 'the spirit of research', had sought to 'paint the invisible and, therefore, the unpaintable' (Fry, 1966, p. 166).

visual images and pictures resides in the fact that those contents are both subject to the same sort of sight-related restrictions.⁵

Where, in the case of visual imagery, might those restrictions come from? Pylyshyn holds that visualisings are episodes in which one ‘simulates’ seeing something (e.g. Pylyshyn, 2006, p. 300, p. 301 and p. 325)⁶, or in which one produces a simulation of what things would look like if one were to see them (Pylyshyn, 2006, p. 325). If he is correct, the representational restrictions noted in the previous paragraph look to be easily explicable. For our simulations of seeings will presumably be informed by our awareness of what real seeings are like; we will only simulate seeing the sorts of scenes which we we could apparently see. Hence, for example, the only features of things which we will pretend to see when we simulate seeings will be ones which we might seem to see.

Pylyshyn offers to explain some of the best-known experimental results in the literature using the thesis that they arise from the ‘use of tacit knowledge [by experimental subjects] to simulate aspects of real-world events, as they would appear if [the subjects] were to see them unfold’ (Pylyshyn, 2006, p. 325; see also p. 306). For example, Finke and Kosslyn used visualisation-involving counterparts of certain standard tests for visual acuity ‘to compare directly the fields of resolution in imagery and

⁵ Note that Pylyshyn might allow that pictures are especially visual in other ways, ones which do not derive from special features of their contents. In particular, he might think that there is something especially visual about the resources upon which we need to call if we are to understand pictures correctly. (The literature on mental imagery contains numerous writers who assume that pictures express their contents on account of visible resemblances between those pictures and what they depict, for instance; see for example Dennett, 1969, p. 52 and Fodor, 1975, p. 78.)

⁶ This is very much the dominant approach to visualisation in Pylyshyn’s writings; see also Pylyshyn, 1981, p. 182, p. 189 and pp. 191 – 2. (Pylyshyn also speaks of ‘pretendings’ in this connection; see Pylyshyn, 2006, p. 306.) It should be noted that Pylyshyn’s thesis that visualisations amount to simulated seeings is not the same as the ‘simulationist’ view explored in Currie, 1995. For Currie explicitly assumes that visual images result from the ‘off-line’ use of portions of the post-retinal visual system, while Pylyshyn’s notion of simulation is much more permissive; Pylyshyn writes, for instance, that ‘the scanning experiment and many other such experiments involving mental images ... [invite] observers to pretend (*in whatever way they can*) that they are looking at some situation ...’ (Pylyshyn, 2006, p. 306; emphasis added).

perception' (Finke, 1989, p. 32, citing Finke and Kosslyn, 1980). Their experiments led them to conclude that 'constraints on visual resolution in imagery are similar to those in perception' (Finke, 1989, p. 34). Pylyshyn proposes, however, that the cited results merely manifest the fact that the experimental subjects, in simulating visual sensations, drew upon their tacit knowledge of 'roughly how far into the periphery of their visual field things can be before they cease to be discriminable' (Pylyshyn, 2006, p. 312).

Now, regardless of the virtues or vices of Pylyshyn's specific account of visualisation, his approach to Finke and Kosslyn's experiments encapsulates the following important general point.

Introspection tells us that there is something importantly visual about visualisation; and it tells us that, in that respect, visualisations are akin to picture-viewings. There is also a lot of experimental data which shows that our performance in visualisation tasks is similar to our performance in tasks in which we really see things, including pictures. But one might well wonder how those introspective and experimental considerations could immediately support conclusions about the *format* of the mental representations which figure in visualisations, as they have often been supposed to do. For our introspective judgements about what visualisation is like, and our performance in the relevant visualisation tasks, seem to be dictated by our comprehension of features of the *contents* which we consciously entertain when visualising, rather than by introspectively inaccessible facts about the nature of the representational media which supply the vehicles for those contents.

Pylyshyn's views thus point up a very general issue concerning the proper morals which may be drawn from the behavioural and phenomenological data concerning

visualisation. But do Pylyshyn's own suggestions—that visualisations amount to simulated seeings, or to simulations of how things would look—provide us with an adequate account of the contents of visual images? In particular, has Pylyshyn done enough to account for the especially visual nature of the contents which we grasp in the course of visualisations, an especially visual nature which those contents seem also to share with the contents of pictures?

3. Discussion of Pylyshyn

As we have seen, Pylyshyn calls on a couple of accounts of what visualisations involve. He sometimes assumes, first, that one who visualises a scene simulates seeing the scene. And, second, he sometimes takes it that one who visualises a scene simulates how things would look to her if she were to see the scene. Each of those accounts evidently connects the contents which are grasped during visualisations to sight. But the putative connections are weaker than they need to be.

Imagine, for instance, a blind person who lacks any awareness of the subjective character of visual sensations, although she possesses a decent fund of factual knowledge about vision. That person may be able to simulate seeing a certain sort of scene, or to simulate how things would look if she were to see them. She may engage, for example, in an elaborate imagining of what a scene looks like to her, by entertaining verbal descriptions of what she supposedly sees ('It's a bright day, bright enough to make me squint a bit. I can see two young girls. One of them has long hair and is wearing trousers ...').

The envisaged blind person's ignorance of what it is like to see things immediately tells us, however, that her simulated seeings, and her simulations of how things would look to her if she were to see them, do not amount to *visualisings*. For one who performs a visualising is thereby aware of how the accompanying visual images show things as looking, in the sense that the visualiser is aware of what it is like for things to look those ways. The especially visual nature of the contents of visual images thus derives, at least partly, from the fact that our ability to grasp those contents is bound to our appreciation of the subjective character of vision. Pylyshyn's position ignores that fact.

Similar points apply to the contents of pictures. One who comprehends a picture thereby appreciates how the picture shows things as looking, in the sense that the picture-viewer knows what it is like for things to look those ways. So pictorial contents are bound to the subjective character of sight in just the way that the contents of visual images are. The similarities between the contents which we grasp through visualisations and picture-viewings therefore don't merely depend upon the fact that visual images and pictures are restricted to characterising apparently visible features of scenes. For the fact that the representations of some kind are subject to that restriction is clearly insufficient to ensure that the contents which belong to those representations are linked to the subjective character of visual sensations in the way that the contents of pictures and visual images are.

The next section will introduce some ideas that will be essential to this paper's alternative account of the contents of visual images. Section 5 will use those ideas to introduce the broad class of distinctively sensory contents to which the contents of both

visual images and pictures belong. (According to the resulting view, the contents which we grasp during visualisations are much more deeply optical than the verbally expressed contents grasped by the hypothetical blind person in the case just discussed in relation to Pylyshyn.) I will then use the resulting framework in examining a range of pieces of phenomenological and experimental data bearing on visualisation.

4. Ways That Things Look or Sound or ... From Perspectives

Perform a visualisation. Then you have thereby pictured how things look from some perspective; your visual image characterises a certain way for things to look as being a way that things look from a perspective.⁷ But what is it for a way for things to look—a type of visual sensations—to be ‘a way that things look from’ a given perspective?

Consider the way that things look to you right now; consider, that is, the visual sensation-type *Your View* whose instances are precisely those visual sensations which are subjectively indiscernible from the visual sensation that you’re currently having. Your *Your View*-sensation involves certain visual appearances whose content you can capture ostensively, by saying that it looks to you as though things are *thus*. But things will look, to anyone who has a *Your View*-sensation, as they now look to you to be. So it will look to anyone who has a *Your View*-sensation as though things are *thus*.

More generally, suppose that certain sensory appearances are common to any sensation of a given type. Then the sensation-type’s *phenomenological content* is the common way that things seem sensorily to be to those subjects who have sensations of

⁷ The points of view from which visual images show how things look can obviously be nonactual. This fact raises a host of hard philosophical problems which I won’t explore in this paper.

the relevant kind. So, for example, the sensation-type *Your View*'s phenomenological content is *things being thus*, because it looks to anyone who has a *Your View*-sensation as though things are precisely *thus*.

Suppose that you are the victim of a *malin génie*, for example, one who has rigged things so that how things look to you to be isn't how they are in front of you at all. Now consider a perspective *a* just off to your right; consider, that is, a bundle of contextual features centred on a nearby place on your right, relative to which it makes sense to consider whether or not things are *thus*. Assume, finally, that your *malin génie* is more timid than they generally are; for your *Your View*-sensations do capture how things are around *a*, as things are in fact *thus* around *a*. Then it comes very naturally to us to say that the way that things look to you is a way that things look from *a*. Why so?

The answer is clear enough. The way that things look to you is a way that things look from *a* because the visual appearances which are common to *Your View*-sensations capture how the world is laid out around *a*. That is, to use the jargon recently introduced, the way that things look to you is a way that things look from *a* because *Your View*'s phenomenological content is true relative to *a*; *Your View* is a way that things look from *a* because, at *a*, things are *thus*.

More generally, suppose that the phenomenological content of a given sensation-type is true relative to some perspective. Then the relevant sensation-type is a way that things look or sound or ... from the perspective. That point is a large part of what is needed for this paper's promised account of the especially visual contents of pictures and visual images.

5. Distinctively Sensory Contents

Visual images and pictures represent scenes in broadly the same manner; they represent scenes by showing how things look from perspectives. The visual nature of visual images and pictures is obviously an important factor in the similarities between them, but those cases also feel notably similar to many nonvisual representations. Thus, we often regard playbacks of sound recordings as auditory analogues of pictures, and auditory images are clearly counterparts of visual images, as indeed are tactile and kinaesthetic images. How are those nonvisual cases related to the visual ones?

Playbacks of sound recordings do not represent scenes by showing how things look, but they may do something related, something that auditory images may also do: they may characterise scenes by capturing how things *sound* from perspectives. Similarly, tactile and kinaesthetic images may characterise scenes by capturing how things *feel* at certain places.

The various nonvisual examples just cited have a further feature in common with their visual counterparts. Recall that, as we saw when discussing Pylyshyn's approach to visualisations, visualisations have a significant subjective aspect, in that one who visualises things as looking a certain way is thereby aware of what it is like for things to look that way. Recollect, too, that the contents of pictures have the same subjective aspect.

Now entertain an auditory image of a scene. Your image captures how things sound from some perspective and your production of it supplies you with an awareness—'from the inside'—of how the image represents things as sounding. Similarly, one who

appreciates the representational significance of a playback that captures things as sounding some way from a perspective is thereby aware of what it is like for things to sound that way; one who grasps the content of a tactile image which represents things as feeling a certain way at a perspective is thereby aware of the subjective character of sensations in which things feel that way; and so on.

Let's say that a representational content is *distinctively sensory* just in case the following obtains: first, the content characterises things as looking or sounding or ... some way from some perspective; and, second, one who grasps that content is thereby aware of what it is like for things to look etc. as the content characterises them as looking etc. Let's also say that a representation is *distinctively sensory* just in case it has a distinctively sensory content. Visual images and pictures are distinctively sensory representations, for example, as are auditory images and many playbacks of sound recordings, and as also are tactile and kinaesthetic images.⁸

The previous paragraph's description of distinctively sensory contents needs fleshing out further, however. For what is it for a content to characterise things as looking etc. some way from some perspective? The ideas articulated in the previous section allow us to answer that question.

Consider a visual image of a scene. That image's content captures how things look from a certain perspective p . That is, the image's content identifies a certain sensation-type—let's call it T —as showing how things look from p . But a sensation-type

⁸ The above definition is meant to allow for the possibility of distinctively sensory representations which bear other sorts of content in addition to their distinctively sensory contents, although I shall simplify things below by speaking as though distinctively sensory representations only possess distinctively sensory contents. The definition is also meant to allow for the possibility of distinctively sensory contents which characterise things as looking etc. more than one way from some perspective, as well as for ones which characterise how things look etc. from more than one perspective.

captures how things look etc. from a perspective just in case the type's phenomenological content is true relative to that perspective. More fully, then, the image's content characterises T 's phenomenological content as being true relative to p .

A sensation-type's phenomenological content captures the common way which things look etc. to those who have sensations of that type. Your image's content thus specifies a certain scene, by identifying the common way which things look to people who have T -sensations as being such that it captures how things are laid out around p . So suppose, for example, that it looks to anyone who has a T -sensation as though a green balloon is next to a pink balloon. Then your visual image, in characterising T as being how things look from p , represents the presence of a green balloon next to a pink balloon.

Generalising all that, we get the following. Distinctively sensory contents characterise the phenomenological contents of sensation-types as being true relative to perspectives. The scenes which are captured by those contents then derive from the specifics of the phenomenological contents which they characterise as perspectively true. Their characterisations of the phenomenological contents of sensation-types as perspectively true are, lastly, *subjectively informative*, in the following sense: one who grasps a distinctively sensory content is thereby aware of what it is like for things to look etc. as the content specifies things as looking etc.

Any two distinctively sensory representations are similar in a certain way; they both have distinctively sensory contents. The class of distinctively representations nonetheless divides up naturally into further subclasses. There is a subclass of distinctively *visual* representations, whose members have contents which characterise the phenomenological contents of *visual* sensation-types as perspectively true, and that

subclass includes both visual images and pictures; there is a subclass of distinctively *auditory* representations, whose members have contents which characterise the phenomenological contents of *auditory* sensation-types as perspectively true, and that subclass includes auditory images and many playbacks of audio recordings; and so on.

Section 1 raised the question whether visualisations stand to picture-viewings as they also stand to merely verbal thoughts about what things look like. The contents of ordinary merely verbal descriptions of ‘how things look’ from somewhere won’t in fact typically count as distinctively visual. The notion of a distinctively sensory content therefore buttresses our natural sense that verbal descriptions of how things look are, in general, fundamentally different from representations like pictures and visual images.

Consider, for example, the following description of how things look to me now: ‘there is a computer screen in front, resting on top of a brown wooden desk, with a wall behind it; there are some little greyish speakers standing on either side of the screen ...’. If that description is to count as a distinctively visual representation, it must characterise the phenomenological content of some visual sensation-type as perspectively true, and in such a way that one who understands the description is thereby aware of what it is like to have sensations of the relevant type.

But recall that the phenomenological contents of visual sensation-types capture the shared ways that things look to be to those who have visual sensations of that kind. A visual sensation-type will therefore have a phenomenological content only if the sensations of that kind literally involve a common way for things to look. The previous description will therefore have a distinctively visual content only if its content’s specification of a scene proceeds by means of the identification of a common way for

things to look, one of whose subjective character we are made aware in grasping that content.

The description does not do that, though. There is, for one thing, no genuine common look associated with in the visual sensation-type which is identified by the description's content; the range of visual appearances in which it 'looks to a subject as though there is a computer screen in front, resting on top of a brown wooden desk ...' do not really involve a single way for things to look to a subject. Thus think for a moment about some of the very different sensations which may feature appearances of that type; consider, for example, the variations in the apparent distances between screen and speakers, or in the overall appearances of the speakers ... which are encompassed by visual appearances which fit that description.⁹

Some rather special verbal descriptions of how things look from somewhere may count as distinctively visual, however. So suppose that somebody shows you a photograph of a scene and informs that, somewhere, things look like *that*. Then the content of that verbal description is perhaps distinctively visual. For the description's content characterises things as looking a certain way from somewhere; and, it is at least natural to think, one can only understand the demonstrative featuring in the description if one is aware of what it is like for things to look like *that*. In this case, though, the description's content is obviously parasitic on the distinctively visual content which belongs to the indicated photograph.

⁹ This does not mean, though, that the contents of distinctively visual representations must be as rich as the contents of ordinary visual appearances. For the phenomenological contents of visual sensation-types may be fairly schematic; there may a common way that things look to the subjects of certain visual sensations which are, subjectively, quite varied. (Consider, for instance, the common way that things look to anyone who seems to see something with a certain specific outline shape.) A distinctively sensory content, in characterising the phenomenological contents of some sensation-types as being perspectively true, may therefore provide a relatively bare specification of a scene.

6. Phenomenological Data

The profusion of experimental results relating to mental imagery has conferred theoretical respectability on the ideas that visual images are inherently visual and that they are, in particular, importantly like pictures. But the introspective sources of those ideas provide them with their most immediately powerful supports. ‘Looking’ at a visual image with one’s ‘mind’s eye’ just does somehow feel like looking at a picture; and our capacity for visualisation feels like something that is intimately related to our ability to appreciate the subjective character of visual sensations.

The idea that visual images and pictures are alike in having distinctively visual contents can help us to understand all that.

Entertain a visual image of, say, a house. As noted earlier, that visual image has a host of general features in common with pictures of houses and also with genuine sightings of houses—for example, it can only represent *visible* aspects of the house and its surroundings. And it is natural to think that the image possesses those features because it, like a picture of a house and indeed like a sighting of a house, shows you what things look like. More generally, looking at a visual image in your mind’s eye is like looking at a picture, because in both cases you end up with a grasp of how something looks.

But suppose that both your visual image of a house and pictures of houses are distinctively visual representations. Then your visual image’s characterisation of its house will depend upon how it shows things as looking, and it will indeed only be able to represent visible aspects of the house and its surroundings; the same things will also hold

true of pictures of houses. For the nature of the scene which is represented by a distinctively visual representation depends on the specifics of the phenomenological contents which the representation's distinctively visual content characterises as perspectively true, where those phenomenological contents capture how things may look.

The assumption that visual images and pictures have distinctively visual contents therefore allows us to explain various specific similarities between visualisations, picture-viewings and seeings.¹⁰ More generally, it substantiates our intuitive sense that picture-viewings and visualisations are alike in culminating in a grasp of how things are showed as looking from somewhere.

The idea that visual images have distinctively visual contents also provides for our sense that the capacity for visualisation is bound to a subject's ability to appreciate the subjective nature of visual sensations. For distinctively sensory contents are subjectively informative; in particular, one can only grasp a distinctively visual content which characterises things as looking some way from somewhere if one can be aware of what it is like for things to look that way. So one can only grasp distinctively visual contents in general if one can be aware of what it is like to see things. And hence the capacity for visualisation is indeed tied up with an awareness of the subjective character of sight.

When describing Pylyshyn's views on visualisation, in section 2, I noted that his ideas dramatise an important point: it is hard to see how the various pieces of phenomenological and experimental data relating to visualisation can straightforwardly support conclusions about the format of the mental representations which figure in

¹⁰ See Gregory, Forthcoming b for a more extended illustration of this point, concentrating upon some of the philosophical issues raised by pictorial representation.

visualisations, because that data seems to derive in the first place from facts about the contents which we grasp when visualising. That point is particularly pertinent in the current context, as suppositions about the format of the representations which figure in visualisations seem to be strikingly irrelevant to the phenomenological data.

Someone might hope to explain our introspective sense that visualisations are like picture-viewings by, for instance, positing analogies between the processes which, in the course of visualisations and picture-viewings, underlie our grasp of certain contents. In particular, someone might hold that visualisations are realised using neural occurrences which feature the deciphering of picture-like representational vehicles which are ‘viewed’ by some sort of inner ‘eye’; and that, in addition, our sense of the similarities between visualisations and picture-viewings arises from our consciousness of the various parallel steps in the processes which underlie visualisations and picture-viewings.

But that is too strong. The very experience of visualising a scene doesn’t explicitly involve an encounter with a representational vehicle, any more than the experience of seeming to see a scene involves an apparent encounter with a representation of that scene. Rather, one who visualises a scene is thereby aware of things as looking a certain way from somewhere without the apparent help of any representing intermediaries, just as we seem simply to be presented with our surroundings in visual experiences. Visualising a scene is, in many respects, more like looking at a *world* with one’s mind’s eye than it is like looking at a *picture*.¹¹

Our intuitive sense that visualisations are like picture-viewings thus does not support the view that the neural processes which realise visualisations are analogous to

¹¹ Note that the ease with which we talk about ‘visual images’ does not clash with these points, for visual images are more like inner analogues of visual sensations than they are like analogues of inner pictures.

picture-viewings.¹² By contrast, the mere idea that visual images and pictures are alike in having distinctively visual contents does not require that visualisations should seem to us to involve representational intermediaries. So, if one identifies the common presence of distinctively visual contents in visualisations and picture-viewings as the basis of the phenomenological similarities between those episodes, one posits just the right degree of introspectively available resemblance between visualisations and picture-viewings.

The thesis that we grasp distinctively visual contents in the course of both visualisations and picture-viewings thus illuminates nicely our introspective sense that visualisation is intimately related both to seeing the world and to comprehending pictures. And the point generalises beyond the mere case of visualisation: the view that sensory mental images of scenes have distinctively sensory contents makes very good sense of our natural thought that the episodes in which we use sensory imagery to ‘picture’ things to ourselves are importantly analogous both to episodes in which we evidently encounter representations which show us how things look etc. and to ones in which we seem to see etc. the world around us.

7. Experimental Data

Introspective considerations make it very hard to believe that there isn’t something especially visual about visualisation. But introspective data have often been regarded as

¹² The above discussion is not an attack on the view that visualisations are in fact realised using neural processes of the type just considered. (The general specification of those processes is loosely modelled on Kosslyn’s idea that visualisations involve scans of picture-like representations laid out on the surface of the ‘visual buffer’; see for example Kosslyn, 1980, p. 6.) Indeed, nothing that I will say in this paper is inconsistent with the view that visualisations are realised using neural processes which have that form. The current point is, rather, that even if visualisations are realised by means of such processes, that fact alone does not provide us a decent explanation of the special phenomenological features of visualisations.

of limited use in assessing hypotheses about the natures of the processes involved in visualisation. Experimental data has therefore been gathered, and lots of it. The view that visual images have distinctively visual contents may help us to understand many of the phenomenological features of visualisation; but can it shed any light on the large amount of data concerning visualisation that has accumulated within the psychological literature, or on the impressive body of neurological results which has more recently been amassed?

The amount of research reporting experiments relating to the performance of visual imagery tasks is huge and I could never have hoped to provide an exhaustive survey of it here. My aim below is therefore more modest: to show that the view that visual images have distinctively visual contents does in fact relate interestingly to a few of the most striking and well-known empirical results concerning visualisation. My examination of pertinent neurological data will be particularly speculative, but I hope that the resulting discussion will at least persuade the reader to see merit in this one thought: that the earlier account of the contents of visual images may let us finally build a bridge between views about the especially visual contents of visual images and ones concerning the especially visual nature of the neural formats which our brains employ in visualisations.

a. Consider, to begin, the imagery scanning results that have been obtained by Kosslyn and others (e.g. Kosslyn, 1973; see also chapter 3 of Finke, 1989 and Pylyshyn, 2006, pp. 303 – 10 for surveys of relevant literature). In the relevant experiments, subjects are made to produce specific visual images—by memorising a map, for instance. The subjects are then instructed to focus on one of the image’s points—the point on the visualised map at

which a tree is marked, say. They are then asked to move their focus to a different point on the image, indicating the moment at which their attention rests on the specified end-point. It has been found, time and time again, that the time-lag between the initial focusing and the subsequent reported one is proportional to the spatial distance between the counterparts of the two relevant focal points which would lie upon a faithful physical realisation of what the subject has visualised—that the delay corresponds to, for example, the distance between the start and end-points on the very map which the subject had to memorise.

These results have led many researchers to the conclusion that '*[t]he spatial arrangement of the elements of a mental image corresponds to the way objects or their parts are arranged on actual physical surfaces or in an actual physical space*' (Finke, 1989, p. 61; see also e.g. Kosslyn et al, 1979, p. 138; Kosslyn, 1980, p. 51). But Pylyshyn thinks differently. He has tried to explain the scanning results using the materials which he also mobilises in response to the experimental data concerning the acuity of the mind's eye (see section 2 above): '[m]y proposal regarding the scanning experiment and many other such experiments involving mental images is that the task of imagining invites observers to pretend (in whatever way they can) that they are looking at some situation and then to use whatever knowledge they have and whatever problem-solving technique seems to them most relevant, to generate the appropriate sequence of mental states' (Pylyshyn, 2006, p. 306).

Pylyshyn's resources may be adequate to the task of explaining the raw data concerning response times. But, when faced by Pylyshyn's proposal, it is surely very hard to resist the pull of the phenomenological considerations which were cited in section 3's

discussion of Pylyshyn's ideas. For Pylyshyn's account of the scanning results does nothing to distinguish between why we—who visualise—are subject to the scanning results and why some other hypothetical beings—who have no awareness of what vision is like and who consequently cannot visualise, but whose nonexperiential knowledge about vision guides their response times in a suitable manner—would be so.

It is nonetheless easy to share Pylyshyn's unease at the way in which the scanning results have led researchers to conclusions about the pictorial format of the representations which figure in visualisations. Our performance on scanning tasks seems to reflect, in the first place, facts about our grasp of the contents of visual images rather than facts about the representational media that our brains use throughout those tasks. This can be illustrated by considering an elaboration of some fairly obvious intuitive ideas concerning the scanning results.

Suppose that one had to perform a typical scanning task. Then having, say, memorised a certain map, one would produce a visual image which focuses on the relevant starting point. Next, one would 'scan' from that initial point to the specified end-point. That is, one would form a series of visual images whose shifting points of focus led smoothly from the initial focal point to the specified end-point. But the length of that series of visual images would depend upon the distance between the two relevant points on the map. The amount of time which it would take for one to scan from the first point to the last would thus correspond to the distance separating those two points on the map itself.

One might think, accordingly, that the measurements obtained in the course of standard scanning experiments are owed to the fact that the experimental subjects

performed the relevant scanning tasks in the general way just described. That explanation of the results assumes, first, that the subjects produced a series of visual images. It assumes, also, that the response times measured in the tasks corresponded to the lengths of the produced sequences of images. Finally, it assumes that the lengths of the produced sequences of images corresponded to the physical distance lying between, say, the start and end-points on the memorised map. Can those assumptions be supported without the use of any claims about the formats of the mental representations which feature in visualisations?

They surely can be. In particular, those assumptions follow naturally once we flesh out the proposed explanation of the scanning results using the idea that visual images have distinctively visual contents.

So suppose, first, that the subjects initially produced a visual image which showed how a map looked from somewhere, and where the specified starting-point provided the focal point of the way that the map was shown as looking in that initial image. And suppose that the subjects then produced a series of further visual images which showed how the map looked, but where the focal points of those ensuing images smoothly traced a path from the starting-point to the end-point. Then we have our sequence of visual images, to whose length the response-times were proportional. And a longer path between the start and end-points will have led to a longer sequence of visual images.¹³

¹³ Similarly, the framework developed earlier lets us stick very closely to the most natural explanations of, say, the mental rotation results first reported in Shepard and Metzler, 1971. The subjects studied by Shepard and Metzler were presented with pairs of drawings of three-dimensional shapes, where each of the pairs consisted either of representations of a single shape or of shapes which were three-dimensional mirror-images ('enantiomorphs') of each other. They then had to assess whether the represented shapes merely differed with regard to their orientations. It was found that the time taken by the subjects to judge whether the matching pairs were indeed representations of a single shape was proportional to the degree of angular displacement between the represented shapes. We can then use the idea that visual images have distinctively visual contents in explaining that finding: when presented with one of the pairs of drawings, a

One of the virtues of the account of distinctively visual contents formulated earlier is that it allows us to make sense of some features of the explanation just presented which might otherwise seem mysterious. In particular, the previous explanation assumed that the ways which visual images show things as looking may involve ‘focal points’. But that assumption might seem to call for the introduction of an ‘inner eye’ which scans representations whose format is pictorial, and which focuses on portions of that inner picture. The framework developed earlier clears the murk, however.

So consider some visual sensation-type. That type may involve a single point of focus, in that the sensations of that type all have subjectively indiscernible centres of visual attention. But, in showing how things look from somewhere, distinctively visual contents characterise visual sensation-types as being such that their phenomenological contents are perspectively true. If a distinctively visual content’s characterisation features a visual sensation-type which involves a single point of focus in the sense just identified, however, then the way that the content captures things as looking from somewhere thereby involves that point of focus too.

In sum, the view that visual images have distinctively visual contents allows us to explain the facts about response times that have been logged in scanning experiments, and to do so in a manner which is faithful to our most natural thoughts about why those results hold. Yet the resulting explanation respects Pylyshyn’s important general insight about the limitations of the behavioural data bearing on visualisation. For the explanation

given subject produced a series of visual images, starting with one capturing how one of the drawings showed its represented shape as looking from somewhere and continuing with ones showing how that shape looked as it rotated smoothly in relation to the same perspective, until the subject reached an image which allowed her to decide whether or not the visualised shape was the same as the one pictured in the remaining drawing; but the length of the resulting sequence of images will have depended upon the degree of angular displacement holding between the shapes depicted in the drawings; so the length of time required to return the requested decision will have been likewise proportional to that degree of angular displacement.

never adverts to introspectively inaccessible features of the representational media which figure in visualisations. Analogous points apply to a host of further well-known experimental results concerning visualisation.

Hence, although recent debates over visualisation have been driven by concerns relating to the format of the mental representations involved in visualisations, the idea that visual images have distinctively visual contents is nonetheless immediately relevant to those debates. For the explanatory uses of that idea just remarked show that some of the major work to which people have sought to put assumptions about the format of visual images can instead be performed using hypotheses solely concerning the contents of visual images. The explanatory utility of this paper's semantic framework therefore substantiates the suspicion that some of 'what is special about imaginal thinking might simply be the fact that mental imagery is associated with certain contents ... [rather than] the way the information is encoded' (Pylyshyn, 2006, p. 289).

b. Section 2 cited Finke and Kosslyn's use of counterparts of standard tests for visual acuity 'to compare directly the fields of resolution in imagery and perception' (Finke, 1989, p. 32), along with their conclusion that 'constraints on visual resolution in imagery are similar to those in perception' (Finke, 1989, p. 34). Pylyshyn seeks to explain those results as merely reflecting how our simulations of seeings are guided by our awareness of constraints on genuine visual resolution, while Kosslyn and Finke take it 'to support the notion that the display medium itself is ... [such that] resolution within it decreases gradually toward the periphery' (Kosslyn, 1980, p. 269). Does the view that visual images have distinctively visual contents shed any light on these issues?

Let's begin by considering the very idea of a visual image's having varying resolution. Kosslyn and Finke's way of articulating that idea regards such varying levels of resolution as deriving from variable properties of the characteristic format in which the mental representations featuring in visualisations are realised. And it may be that certain variations in the introspectively inaccessible 'display medium' utilised in visualisations are causally responsible for the acuity results concerning visual imagery. But the format-based approach does not identify what it is about our conscious experiences of visual images which generate the relevant acuity results.

Think about the real visual case. The varying levels of visual resolution that are manifested by our visual sensations amount to variations of certain sorts in the ways that things look to us; in particular, the information captured by the more peripheral parts of our visual fields is less discriminating than that captured by the more central parts. The varying levels of resolution apparently manifested by visual images are similar: there is a distinction between the more and less peripheral portions of the ways that visual images show things as looking, and the acuity results suggest that the more peripheral regions capture less discriminating information about visualised scenes.

We can use the idea that visual images have distinctively visual contents to elaborate nicely that initial account of what it is for a visual image to have variable resolution. Suppose that some visual image's distinctively visual content characterises a certain visual sensation-type as showing how things look from a perspective. And suppose that, in addition, any sensation of that type involves the instantiation of a certain pattern of varying degrees of visual resolution across the parts of the subject's visual

field. Then the way that the visual image shows things as looking will thereby also involve that very pattern of variable levels of visual resolution.

The pattern of visual resolutions which belongs to a visual image thus corresponds to the common pattern of visual resolutions instantiated by those visual sensations in which things look as the image shows them as looking. The acuity results for visual imagery therefore suggest the following. Suppose that one performs a visualisation, thereby grasping a certain distinctively visual content which characterises some visual sensation-type as being a way that things look from a perspective. Then the visual sensations of that type will involve a common pattern of visual resolutions, one that approximates (e.g. Finke (1989), pp. 34 – 5) the patterns of visual resolutions that are manifested in the visual sensations which we actually have.

What is responsible for that approximation? The mere idea that visual images have distinctively sensory contents does not suggest an answer to that question. So one might think, to adapt Pylyshyn's ideas, that the acuity results merely reflect ways in which our tacit knowledge of what seeing is really like somehow places restrictions on the range of distinctively visual contents which we come to grasp through visualisations. In the absence of some fuller account of why our tacit knowledge should have that effect, however, I would like to suggest another and perhaps more natural approach, an approach which will lead naturally into a brief discussion of the remaining significant category of experimental data which bears upon visualisation.

The variations in visual acuity across regions of our visual fields correlate, initially, with retinal differences: the region of highest visual acuity corresponds to the fovea, the entirely cone-packed central region of the retina in which the photoreceptor to

ganglion cell connection ratio is also lowest; and the increasing loss of visual acuity away from that most acute region corresponds to an increasing decline, away from the fovea, in the ratio of cones to rods and to a rise in the photoreceptor to ganglion cell connection ratio. But those retinal differences also correlate with subsequent cortical ones.

For when one maps the cells on the surface of the striate visual cortex onto the retinal regions to whose patterns of activity they respond, one gets a striking ‘topographic’, and more fully ‘retinotopic’, map of the visual field—adjacent parts of the cortical surface are responsive to adjacent regions of the retina.¹⁴ That map is very distorted, however; it manifests the widespread phenomenon of ‘cortical magnification’—‘[a]bout half of the cortical surface is devoted to the central 10° of the visual field, which occupies only 1% of the visual field’ (Mather, 2006, p. 201). And those distortions correspond nicely to acuity differences, in that the more acute regions of the visual field appear to map onto larger areas of the striate visual cortex than those less acute regions of the visual field which are nonetheless the same size (e.g. Cowey and Rolls, 1974 and Duncan and Boynton, 2003).

Faced by all that, one might well wonder whether the acuity results for visual imagery somehow derive from the neurological circumstances which are responsible for the strikingly related variations in visual acuity across our real visual fields. One might wonder, for example, whether the cortical magnification exhibited by the retinotopic map lying across the surface of the striate visual cortex is not somehow partly responsible for the acuity results concerning visual imagery, or whether the work is not at least done by some later visual areas of the brain, ones whose organisation reflects in some relevant

¹⁴ See, for example, the famous image of the flattened surface of the left side of a macaque’s visual cortex initially published in Tootell et al, 1982 (as reproduced in e.g. Pylyshyn, 2006, p. 392 and Mather, 2006, p. 15).

way the cortical magnification occurring across the striate visual cortex. Those thoughts finally link this paper's discussion to the substantial corpus of research exploring the connections between neurology and visualisation.

Pylyshyn states that '[t]here appears to be some evidence that mental imagery involves areas of striate cortex associated with vision' (Pylyshyn, 2006, p. 394), although the extent of that involvement is controversial.¹⁵ Questions about the precise identities of the visual areas of the brain which are activated during visualisations are within neither my remit nor my competence, but the fact that especially visual neurological areas light up when we produce visual images does pose an obvious question about the relationship between that fact and the approach to the contents of visual images developed in this paper. Do the remarked patterns of activity somehow derive from the distinctively visual nature of the contents which we grasp during visualisations?

It will be useful, at this point, to revisit briefly the analogies and contrasts between visualisations and picture-viewings.

According to what has been stated above, a major point of contact between visualisations and picture-viewings is that those episodes both involve the grasp of distinctively visual contents: contents which characterise visual sensation-types as showing how things look from somewhere, in—and this is the feature which is about to assume prominence—a subjectively informative manner, so that one who grasps such a content is thereby aware of what it is like for things to look those ways. But one crucial

¹⁵ Some research suggests that visualisations involve early visual parts of the brain (see Pylyshyn, 2006, p. 394 for some relevant references; see also Mazard et al, 2004) while other research suggests 'that only later visual areas, the so-called visual association areas, are active in mental imagery' (Pylyshyn, 2006, p. 394 again; see the same page for a substantial list of relevant references).

and very obvious difference between visualisations and picture-viewings relates to the ways in which we come to grasp distinctively visual contents during them.

During picture-viewings, we grasp distinctively visual contents by *looking* at representational vehicles—pictures. It seems easy enough for us to understand, at the barest level of detail, how the process of looking at a picture could provide us with an awareness of the subjective character of those visual sensations in which things look as the picture shows them as looking: if the picture looks to us suitably like how the picture shows things as looking, the subjective character of the visual sensations which we have upon looking at the picture may lead us to an awareness of the subjective character of those visual sensations in which things look as the pictures shows them as looking.¹⁶

Visualisations are, by contrast, wholly internal episodes, ones in which our grasp of distinctively visual contents isn't mediated by proximate uses of our eyes. So how do we then come by an awareness of the subjective characters of visual sensations?

Certain suppositions about broad features of the relevant mechanisms seem plausible in the light of current thinking about the relationships between the visual brain and visual consciousness. The subjective characters of real visual sensations presumably derive, at least in large measure, from neurological events which occur within our visual cortices in response to retinal stimulations. So one might hazard that episodes in which we are aware of the subjective characters of visual sensations, even though we are not

¹⁶ The idea that pictures express their contents by means of resemblances between the experiences which we have upon viewing those pictures and appropriately corresponding real experiences is controversial (see e.g. Wollheim, 1987, pp. 76 – 7; versions of the position are defended by Budd, 1993; Hopkins, 1995 and 1998; and Peacocke, 1987). I think that the view that pictures have distinctively visual contents can be used to provide substantial support for the idea that they at least sometimes express their contents using such resemblances, however. For one of the most simple ways in which to express a distinctively visual content using an external representational vehicle whose content is to be grasped through being viewed is, surely, through the exploitation of similarities between the subjective characters of those visual sensations which one has upon viewing the vehicle and those visual sensations in which things look as the vehicle characterises them as looking.

having visual sensations of a relevantly similar type, also feature activity in the visual parts of one's brain, activity which overlaps what would have occurred in those neural areas if one had really enjoyed a visual sensation having the relevant subjective character.¹⁷

The conjectures in the last paragraph thus suggest that any episode in which one grasps a distinctively visual content, in the absence of suitable prompting visual sensations, will involve activity in neural visual areas, on account of the subjectively informative nature of distinctively visual contents. But we have already seen that visualisations do indeed appear to feature visual neurological activity, and further cases provide that suggestion with yet more support. Visual memories are another category of wholly internal episodes in which subjects grasp distinctively visual contents, for instance, and visual rememberings do seem to feature the activation of visual neural areas.¹⁸

The fact that distinctively visual contents are subjectively informative thus promises finally to build a bridge from theses about the visual nature of the *contents* which are grasped during visualisations to theses about visual features of the representational *formats* which our brains exploit during those episodes. For the subjectively informative nature of distinctively visual contents may mean that those parts

¹⁷ Tononi and Koch, 2008, p. 248 cite evidence which suggests that, while activity in the primary visual cortex is relevant to visual attention, it does not fix the subjective character of our visual sensations themselves. If that is right, and given the supposition in the text, it may therefore be that only activity in relatively late visual areas is required for us to be aware of the subjective characters of visual sensations. The idea that the involvement of visual parts of the brain in visualisations reflects the subjectively informative nature of distinctively visual contents therefore does not force us towards the controversial view that early visual areas should be activated during visualisations.

¹⁸ See Damasio et al, 1993 and Le Bihan et al, 1993 for evidence that visual rememberings activate early visual areas. Handy et al, 2004, p. 632 state, with accompanying references, that 'the emerging consensus is that the retrieval of visual representations from memory leads to the reactivation of cortical areas that were initially activated during the perceptual encoding of those representations'.

of our brains which enable us to enjoy genuine visual sensations also play a role in the representation of distinctively visual contents during visualisations. While the distinctively visual nature of the contents which figure in visualisations may thus be what explains why visualisations are significantly visual in one important sense, the very distinctively visual nature of those contents may mean that visualisations are visual in other important senses as well.¹⁹

Note, moreover, that the previous conjectures generalise immediately to cover those episodes in which we produce nonvisual sensory imagery, because distinctively sensory contents are always subjectively informative. And sensory but nonvisual neural activities are in fact implicated in the production of nonvisual sensory images (see e.g. Bunzeck et al, 2005, Kraemer et al, 2005 and Zatorre and Halpern, 2005 for discussion of auditory imagery; see also Decety, 2003 for discussion of motor imagery). The view that sensory but nonvisual images have distinctively sensory but nonvisual contents may therefore allow us to explain why, just as those images are especially sensory on account of the contents which they have, so they are also especially sensory on account of the neurological activities which take place when we produce them.

¹⁹According to the view mooted in the text, the subjectively informative nature of the distinctively visual contents which we grasp during visualisations is to be explained by the fact that visualisations involve appropriate visual neural areas. It should be noted, though, that the proposed position apparently leaves open the precise amount of further explanatory weight that is to be borne by the assumption that visual parts of the brain are involved in visualisations, because it only posits a fairly specific role for visual parts of the brain in relation to visualisation. So, assuming the view stated in the text, it might be that the acuity results for visual imagery are best explained by the involvement of visual neural areas in visualisations, and in particular by the fact that relevant retinotopic visual neural areas exhibit cortical magnification. Yet, again assuming that view, it might also be that Pylyshyn is correct in holding that the best explanations of many other pieces of experimental data, like for instance facts concerning how people perform in the ‘additive colour mixing’ tasks discussed on Pylyshyn, 2006, pp. 297 – 9 involve the assumption that the performances of the experimental subjects are guided by their tacit beliefs concerning visual matters. (Kosslyn, 1981, p. 229 acknowledges that tacit beliefs probably play a central role in the colour-mixing tasks.) (Many thanks to an anonymous referee for raising these issues.)

8. Conclusion

It is evident to everyone that visualisations are somehow visual. Pylyshyn proposes that the especially visual character of visualisation is owed to features of the contents which we grasp when visualising. His account of the relevant ‘visual’ features of the contents of visual images is deficient, however. For visualisations, unlike simulations of seeings but again like real seeings themselves, carry with them an awareness of what it is subjectively like for things to look certain ways.

The visual character of the contents which we grasp during visualisations can be captured, though. For the view that visual images have distinctively sensory contents—and distinctively *visual* ones in particular—lacks both of the problems from which Pylyshyn’s views suffer, while yet doing justice to our introspective sense that visualisations are essentially visual. And as pictures also have distinctively visual contents, the view also caters for our feeling that visualisations are importantly akin to picture-viewings, without assimilating episodes of those two sorts more closely than the phenomenological data allows.

The view that visual images have distinctively visual contents also sheds light on the experimental data on visualisation deriving from behavioural tests on subjects performing visual imagery tasks. The most natural explanations for many of those experimental results rely upon the fact that, in performing the relevant tasks, subjects will produce sequences of representations of how things look, where those sequences of representations have certain specific features. And by suitably employing the notion of a distinctively visual content, which isolates the precise way in which the relevant

representations ‘show how things look’, we can elaborate those intuitive explanations in theoretically satisfying ways.

The view that visual images have distinctively visual contents does not imply that the especially visual character of visualisations must reside solely in the visual nature of the contents which we grasp during visualisations, however. For the subjective character of our visual sensations themselves is presumably at least partly owed to activities in specifically visual neurological areas. Hence it may be that the subjectively informative nature of distinctively visual contents requires the activation of properly visual parts of the brain during visualisations. The thought that visual images have distinctively visual contents thus promises to help us make a connection between the especially visual nature of the contents which we grasp during visualisations and the especially visual nature of the neural events which take place during those episodes.

Finally, the generality of the notion of a distinctively sensory content means that analogues of all of the above points relating to visualisation apply to those episodes in which we produce nonvisual sensory images, once it is assumed that nonvisual sensory images have suitable sorts of distinctively sensory contents. The notion of a distinctively sensory content therefore does not just illuminate the especially visual nature of visualisation; it can help us to understand what is special about all sensory mental images.

REFERENCES

Alberti, L.B. 1435 – 6/1966: *On Painting*, trans. John R Spencer. New Haven, Connecticut: Yale University Press.

- Budd, M. 1993: 'How pictures look'. In D. Knowles and J. Skorupski (eds), *Virtues and Taste: Essays on Politics, Ethics and Aesthetics*. Blackwell: Oxford.
- Bunzeck, N., Wuestenberg, T., Lutz, K., Heinze, H.J. and Jancke, L. 2005: Scanning silence: mental imagery of complex sounds. *Neuroimage*, 26, 1119 – 27.
- Cowey, A. and Rolls, E.T. 1974: Human cortical magnification factor and its relation to visual acuity. *Experimental Brain Research* 21, 447 – 54.
- Damasio, H., Grabowski, T.J., Damasio, A., Tranel, D., Boles-Ponto, L., Watkins, G.L. and Hickwa, R.D. 1993: Visual recall with eyes closed and covered activates early visual cortex, *Neuroscience abstracts* 19, 1603.
- Decety, J. 2002: Is there such a thing as functional equivalence between imagined, observed, and executed action? In Meltzoff and Prinz (eds.) 2002: *The Imitative Mind: Development, Evolution, and Brain Bases*. Cambridge: Cambridge University Press.
- Denis, M., Mellet, E. and Kosslyn, S.M. (eds) 2006: *Neuroimaging of Mental Imagery: A Special Issue of The European Journal of Cognitive Psychology*. Hove: Psychology Press.
- Dennett, D. 1969: The nature of images and the introspective trap. In Block, N. (ed.) 1981: *Imagery*. Cambridge, Mass.: MIT Press.
- Duncan, R. O. and Boynton, G.M. 2003: Cortical magnification within human primary visual cortex correlates with acuity thresholds. *Neuron* 38, 659 – 71.
- Finke, R.A. 1989: *Principles of Mental Imagery*. Cambridge, Mass.: MIT Press.

- Finke, R.A. and Kosslyn, S.M. 1980: Mental imagery acuity in the peripheral visual field. *Journal of Experimental Psychology: Human Perception and Performance* 6, 244 – 64.
- Fodor, J. 1975: Imagistic representation. In Block, N. (ed.) 1981: *Imagery*. Cambridge, Mass.: MIT Press.
- Fry, E. F. (ed.) 1966: *Cubism*. London: Thames and Hudson.
- Gregory, D. Forthcoming a: Imagery, the imagination and experience. To appear in the *Philosophical Quarterly*, DOI: 10.1111/j.1467-9213.2009.644.x.
- Gregory, D. Forthcoming b: Pictures, pictorial contents and vision. To appear in the *British Journal of Aesthetics*, DOI: 10.1093/aesthj/ayp041.
- Handy, T.C., Miller, M.B., Schott, B., Shroff, N.M., Janata, P., Van Horn, J.D., Inati, S., Grafton, S.T., and Gazzaniga, M.S. 2006: Visual imagery and memory: do retrieval strategies affect what the mind's eye sees? In Denis, M., Mellet, E. and Kosslyn, S.M. (eds) 2006: *Neuroimaging of Mental Imagery: A Special Issue of The European Journal of Cognitive Psychology*. Hove: Psychology Press.
- Hopkins, R. 1995: Explaining depiction. *Philosophical Review* 104, 425 – 55.
- Hopkins, R. 1998: *Picture, Image and Experience*. Cambridge: Cambridge University Press.
- Kosslyn, S.M. 1980: *Image and Mind*. Cambridge, Mass.: Harvard University Press.
- Kosslyn, S.M., Ganis, G. and Thompson, W.L. 2003: Mental imagery: against the nihilistic hypothesis. *Trends in Cognitive Sciences* 7, 109 -11.
- Kosslyn, S.M., Pinker, S., Smith, G.E. and Schwartz, S.P. 1979: On the demystification of mental imagery. In Block, N. (ed.) 1981: *Imagery*. Cambridge, Mass.: MIT Press.

- Kosslyn, S.M., 1981: The medium and the message in mental imagery. In Block, N. (ed.) 1981: *Imagery*. Cambridge, Mass.: MIT Press.
- Kramer, D.J.M., Macrae, C.N., Green, A.E. and Kelley, W.M. 2005: Sound of silence activates auditory cortex. *Nature* 434, 158.
- LeBihan, D., Turner, R., Zeffiro, T.A., Cuénod, C., Jezzard, P. and Bonnerot, V. 1993: Activation of human primary visual cortex during visual recall: a magnetic resonance imaging study. *Proceedings of the National Academy of Sciences U.S.A.* 90, 11802 – 5.
- Lessing, G.E. 1766/2005: *Laocoon: An Essay upon the Limits of Painting and Poetry*, trans. E. Frothingham. Mineola, New York: Dover Publications.
- Mather, G. 2006: *Foundations of Perception*. Hove: Psychology Press.
- Mazard, A., Tzourio-Mazoyer, N., Crivello, F., Mazoyer, B. and Mellet, E. 2006: A PET meta-analysis of object and spatial mental imagery. In Denis, M., Mellet, E. and Kosslyn, S.M. (eds) 2006: *Neuroimaging of Mental Imagery: A Special Issue of The European Journal of Cognitive Psychology*. Hove: Psychology Press.
- Pylyshyn, Z.W. 1981: The imagery debate. In Block, N. (ed.) 1981: *Imagery*. Cambridge, Mass.: MIT Press.
- Pylyshyn, Z.W. 2002: Mental imagery: in search of a theory. *Behavioural and Brain Sciences* 25, 157 – 37.
- Pylyshyn, Z.W. 2003: Return of the mental image: are there really pictures in the brain? *Trends in Cognitive Sciences* 7, 113 – 8.
- Pylyshyn, Z.W. 2006: *Seeing and Visualizing: It's Not What You Think*. Cambridge, Mass.: MIT Press.

Shepard, R.N. and Metzler, J. 1971: Mental rotation of three dimensional objects. *Science* 171, 701 – 3.

Tononi, G. and Koch, C. 2008: The neural correlates of consciousness: an update. *Annals of the New York Academy of Sciences* 1124, 239 – 61.

Tootell, R.B., Silverman, M.S., Switkes, E. and de Valois, R.L. 1982: Deoxyglucose analysis of retinotopic organization in primate striate cortex. *Science* 218, 902 – 4.

Tye, M. 1991: *The Imagery Debate*. Cambridge, Mass.: MIT Press.

Zatorre, R.J. and Halpern, A.R. 2005: Mental concerts: musical imagery and auditory cortex. *Neuron* 47, 9 – 12.

Department of Philosophy

University of Sheffield

Word count: 11818