

APPEARANCE AND REPRESENTATION

by

MAARTEN STEENHAGEN

Thesis submitted in partial fulfilment of the requirements for
the degree of Doctor of Philosophy

Department of Philosophy

University College London

2015

Abstract

At the intersection of aesthetics and the philosophy of perception lies a problem about representational images. When you look at Vermeer's *View of Delft*, do you in fact get to see Delft? It would be nice if we could answer in the affirmative, as it would so neatly explain many of our practices in engaging with images. Be it in churches, advertising, or psychology labs, we typically use images as substitutes for the immediate perception of things. Here is what I claim: Images make it possible to see the objects they represent, and they do so because they can function in vision as perceptual mediators. A perceptual mediator is an object of perception by means of which we can perceive something wholly distinct from it. I show that, while such perceptual mediators may be exceptional in vision, in other sensory modalities they occur universally. Our sensory lives would be impoverished without the possibility of perceiving things mediately.

In chapter 1 I identify a fundamental dispute about images. The dispute is about whether images represent by making absent scenes appear to a viewer, or by merely copying the way an object or event looks. Many philosophers favour the idea that images merely copy how things look. In chapter 2 I explain my dissatisfaction with the copy view. I lay out my main positive proposal in chapter 3. How does the idea that images make absent objects visible fit with our current concept of perception? I formulate my answer through a study of the other sensory modalities. In addition, it is no mystery that sounds can make other things heard. My arguments forge a connection between images and sounds that has not been explored in recent

work on representation or perception. I show how it offers an attractive solution to our problem. In chapter 4, I argue that moreover, we may regard mediate perception as a purely sensory mode of perception. There is no need to rely on non-sensory or quasi-sensory effects of thought or cognition. Instead, coming to see a painted surface as an image consists in no more than discovering a previously unseen aspect of our visible surroundings. In chapter 5 I consider at a more general level how visual images contribute to visual perception. Overall, I establish how visual representation has a structural place in our visual world.

Contents

Contents	7
List of Figures	9
Acknowledgements	13
Introduction	17
Images in practice	19
Images in perception	23
The argument	25
1 The problem of images	29
1.1 The problem of vision	37
1.2 Active and passive perception	43
1.3 More than meets the eye?	50
1.4 Understanding visual representation	56
1.5 The way ahead	60
2 Copying appearances	63
2.1 Ringers and reproductions	67
2.2 Sight's motley crew	71
2.3 Ringers in the realm of vision	77
2.4 The visible past	86
2.5 Traces and symbols	93
2.6 Images without copying	99

3	The eye's contraband	101
3.1	Sounds	103
3.2	Beyond the world of sound	106
3.3	Representation in perception	113
3.4	Images	119
3.5	The problem of visual opacity	127
3.6	Representational appearances	135
4	Coming into view	139
4.1	Hylas' theory of visual representation	140
4.2	Suggestion and sensory variation	147
4.3	Perception's partial character	155
4.4	Penetrating vision	160
4.5	Difference and identity	167
4.6	Discovering an image	172
5	Structures of appearing	175
5.1	Appearance and representation	179
5.2	The image's silence	186
5.3	The accretion of standards	191
5.4	Misperception	198
5.5	Conclusion and summary	205
	Bibliography	211

List of Figures

- 1 Johannes Vermeer, *Gezicht op Delft* (View of Delft), 1660–1661. Mauritshuis, The Hague 18
- 2 Infants’ manual investigation of pictures. Reproduced from DeLoache, Pierroutsakos, and Uttal 2003. 20
- 3 Six “real-world objects”. Reproduced from Brady et al. 2009, also available on <http://cvcl.mit.edu/MM/> 22
- 4 The Trinity. In: Hildegard of Bingen, *Scivias*, II.2: Fol. 47r, 1151-1152. Rupertsberg MS (facsimile) 31
- 5 Yūsai Toshiaki, *Kinshū Daioshōzan sekkō funsen no zu* (Japanese scouts fighting at Daheshang Mountain near Jinzhou), 1895. British Museum 72
- 6 Umberto Boccioni, *Dinamismo di un ciclista* (Dynamism of a Cyclist), 1913. Gianni Mattioli Collection 79
- 7 Philippe Halsman, *Salvador Dalí at a book signing*, 1963. 91
- 8 Pierre-Joseph Redouté, *Le ne m’oubliez pas* (Forget-me-not). In: *Choix des Plus Belles Fleurs*, 1827-1833. Paris 120
- 9 Paul Cézanne, *Les Grandes Baigneuses* (The Bathers), 1898–1905. Philadelphia Museum of Art 125
- 10 Peter Paul Rubens, *The Emperor Julius Caesar*, ca. 1626. Stephen Mazoh, New York 141
- 11 Necker cube. First published in 1832 by Louis Albert Necker 151
- 12 Currier and Ives, *The Puzzled Fox*, 1872. New York 152
- 13 R.C. James, *Dog Picture*, 1965. LIFE magazine 154

14	Interrobang, <i>Palatino</i> typeface	175
15	Carl Mydans, <i>Art students at the painting class</i> , 1939. Texas	180
16	Reproduction of a late bronze age seal-impression	185
17	Domenico Ghirlandaio, <i>An Old Man and his Grandson</i> , ca. 1490. Louvre, Paris	193
18	Giorgio Morandi, <i>Still Life</i> , 1962	195
19	Card no. 1 in the Rorschach test, developed by Hermann Rorschach	196
20	Gerard ter Borch II, <i>Gallant Conversation</i> , also known as 'The Paternal Admonition', c. 1654. Rijksmuseum, Am- sterdam	199
21	Gerard ter Borch II, <i>Gallant Conversation</i> (Detail)	203

Acknowledgements

Most of this work was written between January 2014 and May 2015. Several of the chapters draw on earlier material, some of which I discussed with audiences at conferences and workshops.

I am most indebted to Mark Eli Kalderon; for his sound advice and guidance, for his questions and criticism, and for showing me how philosophy can be done profitably. Not only has he implanted in me a healthy suspicion of jargon and footnotes, I have also come to share with him a deep admiration for the history of philosophy, and an appreciation of the many ways in which the past can turn out to be more original than the present.

During the final stages of writing I worked in the Department of Philosophy at the University of Toronto. It was a welcoming place. I am especially grateful to Mohan Matthen, with whom I discussed many of the key ideas in the chapters that follow. A patient and acute interlocutor, his prods for clarification and requests for specification have improved this work in countless ways. I am also thankful to him for sharing some of his own work before it appeared in print.

In working through the extensive literature on perception and representation, the published contributions of three philosophers stood out immediately for their originality and their grasp on the questions that matter. A reader familiar with their writing will, even where I do not explicitly refer to them, find in the pages that follow many echoes of their work. What is most interesting is that they—Richard Wollheim, Brian O’Shaughnessy and John Hyman—prove that even the

most sensitive of observers can end up advocating starkly different views about how to make sense of observations that, on the face of it, appeared plain and simple.

While writing I have benefited from suggestions and comments of a number of people. They include Chris Meyns, Mike Martin, Chris Peacocke, Anthony Savile, Caitlin Dolan, Paloma Atencia-Linares, Gideon Rosen, Elisa Caldarola, Lambert Wiesing, Mary Carman, Bence Nanay, Roberta Locatelli, Martin Pickavé, Jennifer Hornsby, Benj Hellie, Léa Salje, Rob van Gerwen and Mark Fortney.

Introduction

At the intersection of aesthetics and the philosophy of perception lies a problem about representational images. When you look at Vermeer's *View of Delft* (Fig. 1), then do you in fact get to see Delft? Flint Schier characterises our experience of the painter's cityscape. He writes that "we seem to see through his canvas to a small Dutch town, its dark reflection shimmering in the river" (Schier, 1986, p.1). It would be nice if we could answer that we indeed get to look at and see Delft. This is because it would so neatly explain so many of the ways we already engage with images. Be it in advertising, in churches, or in psychology labs, people typically and knowingly use images as substitutes for immediate perception.

Because of its reluctance to grant the possibility of perceiving what is not present to the senses, modern philosophy has been unable to give a satisfactory story about how images figure in our everyday visual world.

But is this reluctance at all justified? It is undeniable that prephilosophically, people tend to think of images as making absent things visible. Moreover, they tend to think of the way such absent things appear to them when they see images as perfectly normal; as consonant with how they know things to be in their actual surroundings. It is natural to say that a portrait of a warrior who died in battle makes that warrior visible. This is why we take the painting to be a representation. When Philostratus the Younger describes a painting of Achilles on Scyros, he writes about two of the heroes in the visually represen-



Figure 1: Johannes Vermeer, *Gezicht op Delft* (View of Delft), 1660–1661. Mauritshuis, The Hague

ted narrative, Diomedes and Odysseus, that

you see them both, one keeping the glance of his eyes sunk low by reason, [...] the other, Tydeus' son, prudent, ready in counsel and intent on the task before him. (*Imagines*, I, 393 K)

Here, as on numerous other occasions, Philostratus seems to find it unobjectionable to assume that a painted image makes two half-gods visible. You *see* (*ὁρᾶς*) them both, he writes. Philostratus is hardly alone in taking it for granted that paintings allow us to see figures that are not present to sight.

It is not unreasonable to place this common understanding of images at or close to the root of the concept of representation as we find it in the western tradition. The Latin word '*representare*,' from which

our ‘representation’ derives via old French, means the making present again of something (previously) absent (cf. Pitkin, 1967, Appendix). Throughout the tradition authors take painted portraits of someone deceased or in some other way no longer among us as a typical example.

To bring out that we cannot dismiss this as just an ancient way of conceiving of images, let me expand in some detail on two ways in which this understanding of visual representation still forms the background against which we engage with images, both implicitly and explicitly.

Images in practice

Our everyday understanding of what images are is deeply entrenched in our early attempts to make sense of the visual world. Try to reflect back on how you have come to form a sense of what images are. It is likely you will think of picture books, family photographs, or television. All of these were simply part of the visible world around you. It is striking how the way we learn to engage with and look at images is largely continuous with the way we learn about our world of vision.

A team of developmental psychologists at the University of Virginia led by Judy S. DeLoache has conducted a number of studies on how children develop a tacit grasp of what an image is. They allowed the infants they observed to explore pictures of familiar objects. In the meantime, an experimenter documented and rated their behaviour (DeLoache & Burns, 1994; DeLoache et al., 2003). The young participants sit in a high chair. The psychologist presents them with a book, placed on a tray in front of them. The book is a picture book, and it contains realistic colour photographs of various objects (Fig. 2). The experimenters observed over and over again how especially very young infants touch, rub, and strike at the surface of the images in the book. The young children “frequently make grasping motions

as if trying to pick up the depicted objects” (DeLoache et al., 2003, p.115).




Figure 2: Infants’ manual investigation of pictures. Reproduced from DeLoache, Pierroutsakos, and Uttal 2003.

It is natural to interpret these observations in terms of what these children take themselves to see. When confronted with an image of a ball or a cup, their response shows that they take a ball or a cup to be visible before their eyes. Given the basic knowledge of the visual world they have, they assume they are able to touch and grasp the object. It is only typically after nineteen months that infants learn to treat images of familiar items as representations of those items. But it is interesting to see what changes at that point.

To capture what they learn when infants form a tacit understanding of what an image is, DeLoache and colleagues introduce the concept of pictorial competence. As they define it, it is the ability to perceive, interpret, and understand the nature and use of pictures. They explicitly conceive of pictorial competence as a mixed bag of abilities. Some of these abilities rely on or exploit cultural conventions and the pragmatic use of images.

When infants become pictorially competent and so form a tacit understanding of what an image is, they do not lose their interest in images. Neither do they cease actively to engage with them, nor lose their ability to acquire visual knowledge that applies to the objects represented. In particular, when children acquire pictorial competence they do not at all seem to give up their strong perceptual engagement with the objects represented by the images. The only salient behavioural change these studies reveal is that they cease to look at these

represented items as if there were actually there. By nineteen months, grasping typically makes way for pointing. The experimenters take this as an indication of “a nascent awareness of how pictures differ from the real entities they represent” (Simcock & DeLoache, 2006, p.1352). Although they will not be able to put it into words, the behavioural evidence suggests that infants at that age begin to think of an image as something that makes some other thing visible, but does so in its absence.

Of course, such observational evidence is not conclusive. But it at least makes plausible the idea that our everyday understanding of what images are finds its roots in our early explorations of the visible world. From early on, most of us have grown used to seeing an image as a visible thing that puts observers in a position to see something else, even though that other thing is not actually there. The developmental findings at least suggests a natural explanation of the way adults continue to engage with images throughout their lives.

At another front the idea that images make absent things visible is present in our practices explicitly, and in a way that seems hard to eradicate. In psychological experiments images are widely used as substitutes for perception. It has become a robust methodological principle that seeing an image of an item puts observers in a position to see that item. This is by far the best explanation of why our capacity to perceive objects and scenes, and our capacities to integrate such perceptions into further psychological functioning, can be so fruitfully investigated by presenting participants with images.

We find a good example of this in the experimental paradigms of short- and long-term memory studies. Over the last few decades psychologists have studied the relation between our visual perception of items and our short- and long-term abilities to remember the items we saw. These studies have revealed crucial facts about both vision and memory. Yet, typically their experiments rely on imagery (see for instance the extensive review in Brady et al., 2011). Experimenters as-



Figure 3: Six “real-world objects”. Reproduced from Brady et al. 2009, also available on <http://cvcl.mit.edu/MM/>

sume that presenting observers with an image of a bowling peg, say, puts them in a position to form a visual memory of a bowling peg. Visual memory is a kind of perceptual memory; a kind of memory that presupposes that the item remembered once was an object of perception. For this reason, studying visual memory of items by presenting participants with images of those items presupposes that seeing an image of an item puts you in a position to see that item.

To illustrate the point, I want to consider the way Timothy F. Brady et al. describe the method of such an experiment:

We presented observers ($N = 6$) with six real-world objects arrayed in two rows of three [Fig. 3]. The objects were taken from the test pairs used in our previous study of long-term visual memory. On each trial, an onscreen message informed viewers of how long the objects would appear on that trial (1.2, 6 or 18 seconds). Observers then pressed a key and the six objects appeared for the specified amount of time. Then the objects disappeared for

1 second, after which a single object reappeared and observers had to indicate whether it was the same exact object that had previously occupied that location. (Brady et al., 2009, p.1)

The actual procedure of the experiment involves presenting observers with images of six real-world objects. Yet as their description makes clear, the experimenters do not hesitate to assume that by means of such images they can make the objects they represent appear and disappear to an observer. They assume those images suffice to test their ability to form visual memories of those objects. I do not want to claim these experimenters confuse representation and represented. Neither is their reliance on images a mere oversight or product of sloppiness. In other places Brady et al. make explicit that they work with images instead of the items themselves. Yet they suggest that such images have proven to be reliable in investigating the workings of visual object memory. Regardless of what our philosophical theories imply about the role images can play in vision, it seems clear that these images in fact allow empirical study of the way we form visual memories of the items they represent.

Images in perception

Philosophers have been struggling to give mundane observations like the ones I discussed a place. The difficulty arises because of assumptions they make about visual perception. When things appear to sight, it is typically assumed, they need to be present. If we seem to see something that is not present to sight, we are undergoing some kind of illusion. But images hardly ever present us with illusions, so much should be clear. Although they perhaps in some way trick infants, there is good reason to think that as soon as a child acquires a grasp of what an image is—a tacit grasp of the way images can figure in perception—images hardly ever fool their eyes.

The standard way to resolve this is to suggest that images merely copy the visual appearance of the scenes they represent, or more precisely, the way those scenes would look when seen from a specific point of view. I think such a copy theory leaves much to be desired. There have been some philosophers who have suggested an alternative to the copy theory. They maintain that the image gives rise to some special form of quasi-perception. The image Vermeer painted, they claim, invites us to imagine seeing the cityscape of Delft (Walton, 1990); or, they claim, it gives rise to a *sui generis* visual episode unique to our engagement with images (Wollheim, 1980; Wollheim, 2003); or again, they claim stubbornly, our experience of them is illusory or ‘non-veridical’ (Gombrich, 1968; Newall, 2011). In light of our actual practices and attitudes towards images, none of these solutions seem satisfactory. Yet authors repeatedly take them to be the only moves available (see Kulvicki, 2014a for a critical discussion).

In the chapters to come I develop a novel way to elucidate how images represent. I aim to show that our prephilosophical sense of images is accurate. They indeed enable us to perceive absent things, and do so in an entirely mundane and ordinary way. On my view images enable us to perceive things not present to us. My main suggestion is that images are a species of perceptual mediator. This means that they stand on a par with completely familiar objects of perception such as sounds and smells. Just as we can hear things by means of sounds, we can see things by means of images.

My argument for this conclusion is especially relevant for philosophical aesthetics and for the philosophy of perception. For aesthetics, because the discussion about images in that field has had a tendency to focus on the intentions or aims behind images. To my mind this overplays the role of the historical contexts in which images figure. Although such historical factors give rise to important questions on their own, it has led philosophers writing about images to misunderstand the way visual representation is at root a concept that

belongs to our understanding of perception. On the other hand, my argument will be relevant for the philosophy of perception, because recent discussions about perception—vision in particular—have tended to rely so heavily on the concept of representation. Yet they tend to do so without ever elucidating that concept as such (a point emphasised in Travis, 2004). There is no single way of conceiving of representation. I am aware of the large variety of ways philosophers have employed the concept to solve specific puzzles (see Rowlands, 2015 for a recent attempt at an inventory). My aim in what follows is to identify and elucidate a particular kind of perceptual representation, the kind exemplified by images.

As will have become clear from the examples discussed above, when I talk about images, I talk about sensible, visible items located in our environment. An image typically comprises a surface that is differentially marked, stained or otherwise covered or overlaid with an array of colours and shapes. Think of paintings, photographs, drawings, visual projections, photocopies, and so forth. For the moment such examples give us a sufficient grasp on the focus of my argument. I prefer ‘image’ over ‘picture’ or ‘depiction’. This is because the latter two terms are too closely tied to manual painting and drawing (they derive from the Latin *pingere*, to paint). I am confident that not all things that are representational in the way characteristic of painting and drawing are manually produced. As I will suggest later on, some of them may not even be produced at all.

The argument

I will begin in chapter 1 by discussing the problem about images in more detail. Images seem to present us with the visual appearance of the scenes they represent. But does this mean we should we conceive of them as making those scenes visible? Philosophers nowadays typically assume that the answer must be negative. This because they

think it is impossible to see something that is not present to the senses. For this reason the problem is typically not seriously considered. Yet that it is impossible to see something that is not present to the senses is correct only on some conceptions of what vision is. Late medieval traditions in the philosophy of perception, for example, adopted a different stance. They left room for visions of the absent. This is manifest in the visual culture of the period. Recent philosophers have dismissed a potentially fruitful understanding of images without good reason.

The most prevalent view is that images copy the way things look. In chapter 2 I will explain my dissatisfaction with this way of thinking about visual representation. There are two versions of the copy theory. The one, which focuses on visual resemblances, is unattractive because underestimates the range of things a painter can represent. The other version of the theory, which focuses on visual reproduction, is unattractive because it lets an image's history determine what it represents. In that way it is unable to secure the way images represent specifically to the eye. The unattractiveness of the copy theory in both its guises gives sufficient reason to consider a novel alternative.

I will lay out my main proposal in chapter 3. How does the idea that images make absent objects visible fit with our current conception of perception? In audition, the role of sounds is to make other things heard. They function as mediators in perception. Because hearing the sound of a coach does not depend on the actual presence of the coach itself, we have a form of representation—of 'presence in absence'. Representational hearing (by means of sounds) and representational seeing (by means of images) are both species of the genus of representational perception. Philosophers have not explored this connection between images and sounds in recent work on representation or perception, yet it offers an attractive solution to our problem.

George Berkeley maintained that a mediated perception of something is not genuinely sensory. He thinks it crucially involves a non-

sensory act of cognition or thought. If he is right, images can never genuinely make something wholly distinct from them appear to sight. In chapter 4, I show that we may resist Berkeley's conclusion. His argument ignores the way aspects of the visible world around us may require perceptual skill and sophistication to discern and notice. Instead of relying on non-sensory or quasi-sensory effects of thought or cognition, coming to see a painted surface as an image consists in no more than discovering a previously unseen aspect of the visible world.

My main claim is that images represent by making absent things visible to sight. They can do this because they function as perceptual mediators in vision. In chapter 5 I consider one implication that this claim has for our understanding of the world of vision. In particular, I will show how images introduce into the visible world around us a distinct kind of visible aspect. We may identify the representational aspects of images with visible complexes of more elementary qualities, and instantiating the right visible qualities suffices for a surface to be an image. This firmly roots representation in our visible surroundings.

If what I will argue is along the right lines, then we no longer need to consider the idea that images make the scenes they represent visible as incoherent or philosophically naive. Visual representation simply belongs to the fabric of our visual world.

1

The problem of images

The painter has traditionally been understood as an artist who is able, just by skilfully smearing paints or pigments on a suitable surface, to render the visual appearance of whatever she sets her mind to. This conception has since antiquity informed attitudes towards the art of painting, and has shaped the way philosophers currently conceive of visual imagery.

The painter's ability to create appearances is both a source of admiration, and a cause of philosophical puzzlement. On the one hand we may, as Ernst Gombrich did, marvel at the sublime wizardry of Rembrandt. With a single brushstroke he was able to present viewers of one of his most famous paintings with the visual appearance of the gold braid on Jan Six' cloak (Gombrich, 1968, pp.280–81). It seems apt to say that Rembrandt's brushwork presents us with the visual appearance of things or, in more recent times, that the images produced by photographic equipment do so. On the other hand, it is not obvious what exactly this comes down to. Philosophers have been puzzled by the idea. What is it for an image to present us with the visual appearance of a figure or scene?

In the Middle Ages in the West we find a distinctive answer to this question. During the time in which a scattered bunch of religious cults developed into a Christian church, the painter was thought to

possess the skill to make saintly figures, heavenly scenes and sacred narratives appear to sight. Her brushwork or drawing could make absent scenes visible to the eyes of a viewer. This conception of the painter's art was not so much heralded or defended as a prominent philosophical theory. Instead, it was part of the 'manifest image' of the time, as Wilfred Sellars (1962) has called it—the framework in terms of which people understood their own place in the world. Here that concerns in particular their relationship to visual imagery.

Only in recent years have scholars started to unearth seriously the details of visual culture in the later Middle Ages (see e.g. Alloa, 2013; James, 2004; Scribner, 2001). The resulting literature has revealed a prephilosophical understanding of visual experience that is quite different from the more scientifically informed framework popularised during the Renaissance. Whereas modern people came to think of vision as a purely optical affair, in the pre-modern period—roughly between the early eleventh and late fifteenth century—vision was commonly conceived as in part a spiritual activity. That vision was commonly conceived as in part spiritual, explains the special status of images in both public and private settings in medieval culture. Images were treasured for their role as visual aids. They were thought to offer a way of connecting spiritually with a higher, more perfect reality.

At the time, many people believed that in visionary experience one could get closer to the divine. Such divine visions could in principle be attained purely in the mind's eye, without any visual aids. In solitary contemplation, trained monks could become enraptured in visionary states in which they believed to be in spiritual contact with a saint or deity. However, such an ability was deemed rare. Fortunately, spiritual visions were also believed to be facilitated by the images produced by the painters (McGinn, 2006, p.187). No doubt such material images were less pure, as they required a corporeal eye. Nonetheless people accepted that the works of the painter were at least a good way

to train one's spiritual vision. People were told that, by focusing their bodily eyes on a painting of a religious scene, they could enter into a direct, affective bond with the sacred figures they attended to. They could put themselves in the same position that a trained monk was able to attain unaided (Meiss, 1951, p.145).



Figure 4: The Trinity. In: Hildegard of Bingen, *Scivias*, II.2: Fol. 47r, 1151-1152. Rupertsberg MS (facsimile)

A clear example of this role for the painted image is found in the Rupertsberg manuscript of the *Scivias*, a work by the twelfth century philosopher and mystic Hildegard of Bingen (1098-1179). In this work (sadly lost during the second World War) we find illuminations that almost certainly were based on Hildegard's own vivid descriptions and sketches of visionary experience (Caviness, 1998). The illuminated Rupertsberg manuscript gives evidence of how lay people of the time optimistically attempted to attain a glimpse of sacred figures as they had appeared to mystics like Hildegard (Leitch, 2001, p.43). In describing her vision of the Trinity, Hildegard writes how she saw a radiant light (the Father) that suffused a sapphire-blue figure (the Son), who itself blazes with a glowing fire (the Holy Spirit) (Newman, 2005, pp.156–57). The painter tried to recreate the appearance of this Trinity quite literally, by producing an image that allows a viewer to see this figure as it appeared to Hildegard, though now aided by the bodily eye (Fig. 4). Images like this were conceived as instruments of visionary experience and, as Jeffrey Hamburger explains, were intended to “induce, channel, and focus that experience” (Hamburger, 1989, p.174).

Extant explicit discussions of painting in medieval writings are still scarce (Marshall, 1981, p.170). Nonetheless, there is ample cultural evidence that during the later Middle Ages paintings and illuminations were generally believed to expand the sphere of the visible. As Stephanie Leitch writes,

many objects created for private Christian devotion mediate the act of vision. They functioned as aids to conjuring visions, but did so in order to achieve a direct encounter with God (Leitch, 2001, p.46).

This mediating function of painted images was not limited to clerical or even primarily religious circles, but spread throughout the vernacular cultures of countries such as Italy, France, Germany and The

Netherlands (for detailed discussions, see Scribner 2000 and Biernoff 2002).

If this is right, then during the later Middle Ages images were believed to have a power to make absent scenes visible. To modern ears this idea is likely to sound odd. It does not seem possible for a painted image to make an absent scene visible, because it does not seem possible to see what is not present to the senses. From a modern frame of mind it is natural to assume that at best images can afford us the *illusion* of seeing; yet it is patently clear that such illusions are extremely rare. Looking at Rembrandt's portrait of Jan Six, or glancing at the illustrations while browsing a fourteenth century Book of Hours, it is manifest to us that what we see is images of those scenes, and equally clear that we are not presented with those scenes themselves. Despite their splendour, the painter's efforts hardly ever fool the eye.

This leaves us with a tension. On the one hand we are happy to accept that images present us with the visual appearance of a figure or scene. Yet on the other we may remain hesitant or explicitly reluctant to accept the conclusions the medievals seem to have drawn from this observation about images and the painter's art. What is behind this tension? I suggest the tension arises because of a crucial ambiguity in the concept of a visual appearance. With our descriptions of appearances, we sometimes aim to characterise perceptual encounters as such, as when I say the lorry appeared in the distance. At other times we wish to characterise the sensible qualities of the objects or scenes we so encounter, for example when I say of your car that it has a great appearance. I take this difference to reflect a distinction between two ways of thinking about appearances.

In his essay 'Primary and Secondary Qualities', published posthumously as a 'tentative investigation', John Cook Wilson highlights that perceptual acts such as seeing a lighthouse, feeling the wood, or hearing a barking dog may equally well be conceived as appearances of these objects—a lighthouse, wood, a barking dog—to a perceiver

(Cook-Wilson, 1926, p.796). In that case, we are consider the same occurrences from the side of their objects. Here we are strictly only concerned with a particular perceptual event, Cook Wilson emphasises—the *appearing* of the object to a perceiver. The event of an object’s appearing to a perceiver is in this case not itself an object of perception. It is the event of the object’s being seen, heard or smelt by a perceiver. It is not ‘*sensible*’ but ‘*sensory*’ as I will call it (I borrow the distinction from Byrne, 2009). A sensory appearance of the ceiling when I open my eyes in the morning is an event of my perceiving that ceiling.

Sensible appearances, by contrast, are those appearances that are themselves objects of perception. H.H. Price brings this out several decades later, when he discusses a viable but different way of thinking about appearances. Conceived in this way, appearances can themselves be perceived. In his paper ‘Appearing and Appearances’ Price observes that we describe objects of perception as “appearing so-and-so” or as having a “such-and-such appearance”, where ‘so-and-so’ or ‘such-and-such’ stand for predicates such as ‘red’, ‘rotten’, or ‘coated’ (Price, 1964, pp.5–6). Price thinks we describe appearances in this sense when we use specific locutions such as “the apple looks red”, “...smells rotten”, “...feels coated” etc. Here we attribute a possibly complex sensible quality to an object of perception. It is a quality that in part determines how the object strikes us in perception. On this understanding, the apple’s appearance is not an event but a sensible quality of the apple, and as such may itself be perceived. A sensible appearance is a possibly complex visible quality that an object has. As a visible quality, a sensible appearance is something that an object can share with other objects (a point emphasised by Brian O’Shaughnessy as well; 2000, p.571).

When discussing appearances, we should take care to distinguish between sensory and sensible appearances. A potential ambiguity in what we say and think about appearances helps us make sense

of what I take to be the central philosophical dispute about images. When we see an image, we are made aware of the visual appearance of some figure or scene. That is obvious—it is by rendering such appearances that the painter pursues her art. But how are we to understand what goes on when this happens? Medieval authors accepted that the painter's images make for a sensory appearance of what they represent—that they enable an event in which absent scenes visually appear to a viewer. The view popularised in the Renaissance was that the painted image merely copies the sensible appearance of a tridimensional scene—it copies the visible qualities that determine the way a scene looks. Which of these views, if any, is correct?

The view that the painted image merely copies the visible qualities that determine the way a scene looks or would look, popularised in the Renaissance but often traced back to Plato, may be regarded the standard view in current philosophy. Recent authors accept that the art of painting is more or less exclusively concerned with appearances. They understand this to mean that the painter copies the look of things through the application and two-dimensional arrangement of pigments on a flat bearer (cf. O'Shaughnessy, 2000, p.530,571; Hyman, 2006). Even philosophers who do not accord with this standard view are typically reluctant to treat images as genuinely sensory aids. They suggest that, although images do not represent by copying sensible appearances, it is equally false that they represent by making absent scenes visible. Instead, an image represents by way of a quasi-sensory appearance: images are visual illusions or cause non-veridical perceptions (Gombrich, 1968; Newall, 2011); they invite us to imagine seeing those scenes (Walton, 1990); they give rise to an experience unique to image-seeing (Wollheim, 1980; Wollheim, 2003). This has pushed the central question, namely whether images can make the scenes they represent genuinely visible, into the background. Yet, to my mind the issue is far from resolved.

It may be objected that the question whether images can make

scenes visible is not raised because the question is confused. The medieval conception of images, as I have presented it, rests on no more than a wishful fantasy, the objection runs. It is perhaps understandable that a deeply religious culture dreams up a connection with a reality entirely outside the grasp of ordinary people. Yet as a philosophical theory, this idea gives way under even the slightest pressure. Only things that are actually present to the senses can be perceived. Therefore, the suggestion that images afford us visions of scenes not present to the senses is not a philosophical answer in the current debate (cf. Hyman, 1989, p.22; see also Hopkins, 2003, p.157). It does not explain what it is for a painter to render the visual appearance of things. At best it is an expression of philosophical puzzlement. More problematically, in the hands of medieval authors it is likely to be evidence of animism, superstition, or plain wishful thinking. The real challenge is to explain why, in a situation where we *cannot* see a scene, we still find it so natural to say that we do, and exercise a variety of visual abilities successfully.

Although familiar enough, this diagnosis rests on a mistake. The reasoning lets a piece of dogma get in the way of a fruitful approach to images. As is widely known by scholars writing about the period, the medieval attitude to images is not just some instance of animism, superstition, or wishful thinking. Medieval visual culture manifests itself in the way it does because of specific philosophical assumptions about vision (Miles, 1983; Biernoff, 2005, p.44). Although their attitude towards images was not a topic of explicit philosophical theorising, medieval authors certainly did engage in philosophical disputes about what vision is.

In the next two sections I want to show in what way the medieval attitude towards visual images was made possible by specific assumptions about vision. I will first show how the philosophical dispute about images is not independent of another source of philosophical puzzlement, namely puzzlement over vision as such. Following, I ex-

plain how a specific resolve of the puzzlement about vision enabled philosophers to allow for an understanding in which vision can be of more than meets the eye. No doubt the medieval view seems controversial to contemporary philosophers of perception. Yet, as I hope to show, it at least contains the germ of an important correction to our own understanding of the role images play in vision.

1.1 The problem of vision

Implicit in late medieval religious practices, or more broadly in aspects of the manifest image of the time, we find a stance towards a crucial philosophical dispute about images. At the time, many took it for granted that images were able to make visible the scenes and figures they represented. I noted how this idea to many today will sound strange and superstitious. How could they have thought that objects not actually present to us could nonetheless be seen? There may not be a single answer to this question. But I want to suggest that the conviction cannot be dismissed simply as a piece of animistic lore. I will show that it quite naturally flows from a specific solution to a far more general philosophical puzzle.

Before I continue, however, I want to clear up one persistent misconception. Not infrequently it is assumed that image practices throughout the Middle Ages are rooted in the almost magical belief that an image of a saint is identical to that saint. Only because they believe that representation and represented are identical, the thought runs, do people during the Middle Ages engage with images in the way they do. Some authors ascribe this view tacitly to medieval practitioners, for instance when they assume that when viewers were said to gaze at a represented saint they were taken to be gazing at no more than a painting of that saint (e.g. García, 2011). Others, such as Belting (1994) and Freedberg (1989), attribute the view explicitly, and even celebrate it for its alleged profundity (for a recent criticism

of this way of ‘mythologising’ the image, see Wiesing, 2013). If these authors are right, then medieval image practices were clearly rooted in a confusion.

Such an interpretation is untenable, however. No doubt plenty of people will indeed have believed something like this (Antonova, 2010, p.170). Plenty of people nowadays believe they were abducted by extraterrestrials. Yet among the learned who took seriously the way images afforded visions of saints and religious narratives, there were many who distanced themselves quite explicitly from any such crude animism. Official guidelines on how to use images in clear terms dictate that the “sacramental gaze” of the pious onlooker should go out to the saint visible by means of the image, and not to the painting itself (Scribner, 2001, p.95). Moreover, the distinction between representation and represented seems to have been accepted even among lay communities. Eugène Honée describes how ordinary people were undertaking pilgrimages to alleged ‘miraculous’ images—images about which people said that the represented saint was actually present in them (Honée, 1994, p.158). Those pilgrimages may have been misguided, but the fact that those images were deemed *miraculous* proves that even lay people did not identify images as such with what they represented.

Instead, the crux is this. The claim that we can see things that are not present to sight is only puzzling on the assumption that vision requires its objects to be present. It is a widespread view among philosophers, both of the past and of today, that the objects of perception must somehow be in contact with or otherwise act upon a perceiver. This is because perception is passive—it is something we undergo. Yet, precisely during the later Middle Ages, this assumption itself stood under pressure. A significant tradition was sceptical of the suggestion that perception requires contact. They questioned the idea that perception is a passive response to a stimulus. The pressure in part turned forceful because of a general puzzlement about

vision.

Looking out of his office window in the Warburg Institute in Byng Place in London, Ernst Gombrich was able to see the towering building of Senate House on Malet Street. A considerable distance separates the respective locations of these buildings. Sight has the crucial characteristic of being one of the distal senses. Objects we see seem to be located at a distance from us. If objects are located at a distance, then they are not in contact with our power of sight. To be in contact, two things must be spatially simultaneous in some way. Contact does not allow for gaps. Hence, it emerges that in visual perception the perceiver is not in contact with the perceived. Yet at the same time, it is also natural to think that perception requires contact between perceiver and perceived. This is because in perception objects act upon our senses. If there is no such thing as action at a distance, then we are in contact with the objects acting upon our senses. Someone may throw a long ball in your direction. The ball can only act upon your cutaneous receptors if it actually makes contact with your skin. You only feel the ball when it actually hits you.

Puzzlement about vision arises because two assumptions seem very natural. On the one hand, that perception requires contact between perceiver and perceived. On the other hand, that in visual perception the perceiver is not in contact with the perceived.

1. Perception requires contact between perceiver and perceived
2. In visual perception perceiver and perceived are not in contact

The first assumption is a general thesis about perception. The second is an empirical observation about the conditions of visual perception, an assumption that any sighted individual can try to verify for themselves. (To follow a suggestion of Aristotle, one could try to bring an object in contact with the eye and assess whether one can still see it.) Something cannot both be and not be in contact with another thing. To resolve the puzzle, one or the other assumption has to give.

The puzzlement about vision has shaped discussions in the philosophy of perception since ancient times (see Lindberg, 1976; Tachau, 2006; Tachau, 1988). To many the assumption that perception requires contact between perceiver and perceived seemed unshakable. Perception was generally understood to be a kind of alteration. One perceives only when one's sensory receptors are altered by something else, an idea we find in Aristotle. As Aristotle emphasised, every case of alteration requires spatial simultaneity, that is, contact (*Phys.*vii.2, cf. Wardy, 1990). Reflection on the sense of touch confirms his general theory of perception as requiring alteration of sensory receptors through contact. As David Lindberg observes, in the ancient world

the analogy of perception by contact in the sense of touch seemed to establish to nearly everybody's satisfaction that contact was tantamount to sensation, and it was not apparent that further explanation was required. (Lindberg, 1976, p.39)

For a long time touch was regarded as the clearest model for perception as such, in part because it exemplified so clearly the idea that perception is a kind of alteration. For this reason, the most familiar and influential way to resolve the above puzzlement about vision was to reject the second assumption. Despite appearances, the idea would be, in visual perception the perceiver is in contact with the perceived.

Philosophers who accepted that perception requires contact inferred that there must be some some unnoticed object of perception, some *tertium quid* besides the perceiver and the distal object, that is able to act upon the perceiver as a proxy. This hypothesis gave rise to a sophisticated theory of vision, one that conceived of vision on the model of the ingestion or intromission of visual forms. The intromissionist solution to the problem of vision goes back at least as far as the early atomists, who proposed that distal objects emit thin films in all directions, films to which only our eyes are sensitive.

Frequently, these films were conceived as likenesses of the objects that shared their visible form. By means of the shape and colour of the forms that reach our eye the soul is able to interpret them as belonging to this or that distal object.

Here we have the emergence of a distinctive and controversial theory of vision, a hypothesis that seeks to explain how it is possible to see objects at all. Vision is possible, because objects can send out visible films that reach the eye. The distal object is causally responsible for bringing about an alteration in the perceiver, by propagating forms that impinge upon the perceiver. This secures the presence of the perceived object, be it by proxy. The object can be perceived because it emits visible forms, which in turn are in contact with the perceiver. In other words, it provides a way to hold on to the assumption that perception requires contact, and that we can only perceive what is present to sense.

In recent times, intromissionist theory has emerged as an apparent victor. With some modifications, it outlines the standard modern way of thinking about vision, as a wholly passive bombardment of the eye. The thin films have given way to rays of light, and the reliance on likenesses has been given up in favour of a confidence in the principles of modern optics and neuroscience. When we reach modern philosophy, the story we get told starts to sound dangerously mundane to our ears, obscuring its theoretical ingenuity: light is reflected by the object, it can be divided in rays, each travelling in a straight line from a part of the object to our eyes, causing subtle, systematic alterations in the organ of sight—distant objects reflect light that can act on the eye and its retina, enabling vision of those objects. Many philosophers still take this as an adequate explanation of what it is to see things at a distance (e.g. Burge, 2010; see Travis, 2013 for dissent).

However, the intromissionist view was not always dominant. In ancient and medieval times the intromissionist view was heavily criticised. Many writers took the view to be simply incoherent. When

Gombrich looked out of his window and saw Senate House's concrete façade, he was able to see one of the largest buildings in Bloomsbury. The intromissionist now makes it seem as if a visible film or copy of the entire building must have somehow fitted through the tiny apertures of Gombrich's eyes. This seems absurd (Lindberg, 1976, p.53). Especially before the view could be given an optical exposition, many philosophers took the intromissionist solution to leave the actual process of visual perception more mysterious than any of its rivals.

And rivals there were. At least in the later Middle Ages, the intromissionist answer to the puzzle of vision was taken to be far from obvious. Instead, an Augustinian understanding of psychology and perception, strongly based in a neo-Platonic tradition deriving from Plotinus, was frequently taken for granted (Miles, 1979, ch.2 helpfully compares Plotinus and Augustine; see also Emilsson, 2008). A driving thought of Augustinian psychology was that perception of the material world does not require any contact between it and the soul. Instead, perception is a spiritual activity directed at the world. Although Augustine's remarks on the nature of sight are scattered, it is clear that he maintained that we can only see the material world around us with our eyes because of the way we act and direct ourselves towards that world (Silva, 2014a; Miles, 1983, p.139). Seeing, according to Augustine, is an activity of the soul. It is the activity of directing one's visual attention outwards, of visually reaching out for and spiritually 'touching' the object. Without this activity of the soul, the eye could be open, but it would not be able to get anything into view.

A late exponent of a view of vision inspired by Augustine is the eighteenth century German mystic Emanuel Swedenborg. Swedenborg writes that "[t]he sight of the eye, strictly speaking, is nothing but the sight of the spirit produced outwards" (cited in Kant, 1900). In the eighteenth century, Swedenborg's adherence to an Augustinian psychology was deeply unfashionable. Yet in the later Middle Ages a version of the Augustinian theory of vision had the status of the re-

ceived view. It was only with the rediscovery of the works of Aristotle, via the commentators and the Arabic tradition of optics, that the active dimension to vision came to stand in need of defence.

1.2 Active and passive perception

Not infrequently, the heart of a decade-long war can be understood by grasping the key manoeuvres taken on its most crucial battle grounds. Late medieval philosophy saw a revived controversy over the nature of sight, resulting from a clash between an established Augustinian picture of psychology and a renewed interest in an Aristotelian science of the mind. In this dispute, an important ground of skirmish was the wax analogy—the simile that likened visual perception to the impression of a shaped ring on a wax tablet.

The wax analogy in the context of vision is, if not introduced, then at least most famously used by Aristotle when he writes that

In general concerning all perception, it must be grasped that a sense is a capacity of receiving the sensible forms without matter, as the wax receives the mark of the signet ring without the iron or gold; it takes the golden or brazen mark, but not as gold or bronze. (Aristotle, *De Anima* 424a17–24)

What is it for the soul to be impressed by an external form as the wax receives the mark of a signet ring? According to the theory of vision offered in *De Anima*, the sensible form of the distant object is transmitted through the medium of air, allowing the perceptible qualities to alter the state of the eye, without bringing the eye in contact with the object's matter (see Kalderon, 2015 for a compelling reconstruction of Aristotle's view). By its very nature, Aristotle's use of the wax analogy focuses attention on the passivity of perception. The soul,

like the wax, is impressed by the qualitative form of an external object.

From the second half of the twelfth century, the writings of Aristotle started to be rediscovered by philosophers and theologians in the Latin West. The resulting interest in Aristotelian philosophy, together with original contributions to optics, medicine and natural philosophy made by Arabic philosophers such as Avicenna, Averroes, and Alhazen, sparked off a scientific revolution. Aristotle's claim that perception, hence vision, occurs in a way akin to how a wax tablet receives a form, when combined with Alhazen's intromission theory grounded in ray geometry, made for a powerful research programme. It also installed a sudden point of friction with the more traditional visual theory as suggested by Augustine, and more or less taken for granted by the learned communities of the day—a theory that assumed that vision did not require contact or external impressions (Silva, 2014b, p.118).

Many medieval commentators took Aristotle's theory of vision to be broadly in agreement with the intromissionist account advanced by the atomists. Several influential authors took Aristotle to mean that vision consists in the sensory soul's being acted upon by visible forms or 'species', which can act upon the power of sight by being propagated through the air towards a viewer. The intellect can then attribute the resulting visual impression to the distal objects that seem to have emitted those species. A clear example of this reading is found in the writings of William of Auvergne (1180–1249), one of the earliest Aristotelian theologians. Auvergne takes himself to follow Aristotle when he writes that

the act of seeing consists of two things (...) namely the impression or reception of a visible modification that is produced in the eye and the cognition or judgment by which the visible thing is known and is judged with regard to its color and shape. (*DA* 5.6, 121, transl. Teske, 199. Quoted

in Silva, 2014b, p.122)

Auvergne's emphasis on an impression produced in the eye echoes Aristotle's mark of the signet ring. Yet if the claim here is that the eye is merely acted upon, and that it requires a non-sensory cognitive act or judgement to apprehend our distal environment, then this is not what Aristotle had in mind (unsurprisingly, the idea *is* prominent in Alhazen; see Hatfield & Epstein, 1979, p.369). Aristotle's core insight was that the visual power located in the eye is *reactive*: vision is a way the soul reacts to sensible forms acting upon it. This means that for Aristotle perception is both active and passive and not, as Auvergne seems to assume, wholly passive (in addition, medieval authors typically construe Aristotle as suggesting that visible forms are *propagated* through a medium. Also this seems inaccurate, cf. Kalderon, 2015, p.56). What Auvergne, and the medieval Aristotelians following him, do correctly take from Aristotle's view is an emphasis on the passive dimension of sight. In seeing, the soul is being acted upon by the forms or qualities of an external body, and it is this passive dimension that accounts for the sense's power to perceive at a distance. This passive dimension is captured well in the image of a signet ring impressing itself on a piece of wax, no doubt leading Aristotle to reach for the analogy.

As I described earlier, puzzlement about vision emerges because perception seems to require contact between perceiver and perceived, while in visual perception the perceiver seems not in contact with the perceived (cf. assumptions 1 and 2 above). Neo-Aristotelians typically followed the ancient intromissionist solution to this puzzlement. They maintained that, despite appearances, also visual perception involves contact with what we perceive, or at least with the species that propagate from it. In this way distal objects are able, via signet-ring-like species, to make contact with the power of sight.

By contrast, the older Augustinian tradition is based on the wholesale denial of the claim that the material world can act upon the soul.

Developing a neo-Platonist heritage, Augustine maintained that the soul's relative perfection makes it and its acts immune to being causally acted upon by the material world. Augustine was among the most influential Christian authors from the medieval period to the reformation, and his ideas about body and mind shaped late medieval theories of psychology (Miles, 1985, p.7). We find the specific assumption about the inability of the less perfect to alter the perfect even made by René Descartes when, in his ontological proof in the Third Meditation, he contends that the more perfect cannot be the effect of the less perfect. (See Menn, 1998 for a study of how Descartes was indebted to Augustine.) Because of the metaphysical commitment to a certain account of the relation between mind and body, many philosophers who followed Augustine were bound to deny that perception requires contact between perceiver and perceived. Instead, they held that all vision requires a viewer to engage in a spiritual activity.

The Augustinians denied exactly the feature that seemed to make the wax analogy so apt: that in seeing, the soul is acted upon by the form or qualities of an external body. New developments in optics and a revived confidence in the claim that perception is an intromission of visible forms shifted explanatory demands. What needed explanation was how, despite appearances, perception does not require contact.

A representative example of such a defence is offered by the Dominican Robert Kilwardby (c. 1215–1279). Kilwardby was a significant figure both in the church and the universities in England. His writings on sense perception date from the period in which he was first lecturer and then regent master of theology at the university of Oxford (Lagerlund & Thom, 2013, p.3). Throughout his career, later as Archbishop of Canterbury, Kilwardby concerned himself with the increasing tension between the church and the universities. It is unsurprising, then, that his philosophical works try to harmonise the Augustinian philosophy adhered to in theological circles, and the newly

rediscovered natural philosophy of Aristotle that became ever more prominent in the universities. Specifically, it is this project that explains the distinctive form of Kilwardby's solution to the problem of vision. He manages to ambush the opponent by questioning the limitations of their model. He tries to elucidate what it is to see by means of the traditional wax analogy, with a crucial twist. Kilwardby's move is to vivify the wax. He writes that

if you place a seal before wax so that it touches it, and you assume the wax has a life by which it turns itself towards the seal and by striking against it comes to be like it, by turning its eye upon itself it sees in itself the image of the seal. For in this way the sensory soul, by turning itself more attentively to its sense organ which has been informed by a sensible species, makes itself like the species, and by turning its own eye upon itself it sees that it is like the species. And thus it senses the object outside by means of the image which it has formed in itself. (Kilwardby, 1993, p.103)

Kilwardby levels the playing field by turning the assumed passive wax into an active substance. In this way he turns an analogy that tacitly assumes that matter is able to act upon the soul into one that is compatible with the Augustinian psychology. The wax no longer stands for the soul, but for the union of body and soul. The body is acted upon by the seal, but it is the soul of the wax which, by virtue of its awareness of the bodily change, forms in itself a likeness of the impression.

All this means that it is the active soul that determines perception. What we get to see is what we succeed in looking at. Kilwardby explains how, despite appearances, perception does not require contact. The view he defends is Augustinian in spirit, but Kilwardby's presentation brings it more in tune with the new Aristotelian fashion. Every

act of perception comprises two contrary motions: one a transmission from an external object to the organ of sense; the other an attentive act of the soul which responds to the affection of sense organ (Silva & Toivanen, 2010, p.296, cf. Plato, *Tim.*). With neo-Aristotelians Kilwardby accepts that distal vision is possible in virtue of a process of intromission of species: to perceive, the sensory organs must undergo some alteration. Yet he explicitly rejects the idea that this alteration is all there is to vision, or that it causes vision.

The core of Kilwardby's theory of vision is that the active attention senses the object outside by means of a purely spiritual image which it has formed in itself, but does so via a third image. This is a material image formed on the retina of the eye. As Kilwardby explains, we also see the retinal image by means of the spiritual image. This means that in order to see the distal object, the active soul attends to the spiritual image it forms in itself, comes to perceive the material image formed in retina, and in turn by means of the retinal image perceive an distal object or scene in the external world. Vision is a mode of attention that works its way outward following a chain of images (for a recent version of this idea that gives up Kilwardby's commitment to the activity of sight, see O'Shaughnessy, 1985).

Not all neo-Augustinians followed Kilwardby's baroque framework. Many did not take all the acts of corporeal vision that Kilwardby described, mediated by several stages of images, to be required. According to Peter John Olivi (1248-1298), the distal world is opened up to the attentive gaze immediately, and does not depend even on a passive alteration in the body (although he did acknowledge that bodily change is involved in perception in some way, see Toivanen, 2013, p.146). Olivi endorses the Augustinian idea that the soul is ontologically superior to material, worldly things, as well as the view that perception is active in nature. Powers of the soul never passively receive the less perfect external stimuli, but instead they are active. He writes that "we expressly perceive that our acts of seeing

and cognising come intrinsically from, or are produced by, the innermost part of us” (*Summa* II.Q.58; quoted in Toivanen, 2013, p.143). On these matters Olivi clearly sides with Kilwardby (he even transforms the wax analogy in the way Kilwardby does). However, Olivi disagrees with Kilwardby about the role of images in sensory acts. Whereas Kilwardby thinks that we first and foremost attend to the sensible qualities of the image in the eye, and only by means of those come to attend to external, distal objects, Olivi maintains that in vision we can attend to the distal world immediately (Adriaenssen, 2011; Toivanen, 2013, p.148). He sees no reason to suppose that our sensory activity must always proceed via an image formed in the organ of sight. This does not mean that we could never attend to images. Even Olivi allowed that on occasion, our sensory activity focuses on an image that mediates the attention to some object beyond it. This becomes clear when he discusses the role of images in imaginative vision (Olivi showed a keen interest in spiritual vision, which he thought required images, cf. Burr, 1985).

Kilwardby and Olivi’s theories of vision are representative of the neo-Augustinian doctrines of their day. Despite the differences in their theories of vision, they agree that what we see is determined by what we succeed in looking at. Vision is not restricted to whatever propagates visual forms towards the eye. This leaves room for things not present to the senses to appear to a perceiver. Kilwardby’s theory of vision even quite explicitly relies on the possibility that one object of vision acts as a mediator for another. The later Middle Ages are relevant for the philosophy of perception because, under pressure of the rediscovery of Aristotle and the emergence of a respectable theory of light, hitherto tacitly assumed views on vision were given explicit elaborations.

1.3 More than meets the eye?

A neo-Augustinian theory of vision as exemplified by Kilwardby and Olivi does not limit vision to what can be in contact with the sense organs or sensory power. Given this, it presents no *prima facie* obstacle for things not present to the senses to come into view. What appears to sight is at least in part determined by the active operation of the spiritual power of the soul. Both Kilwardby and Olivi conceive of all visual perception as first and foremost an activity of looking. It is our perceptual activity that determines the objects we perceive.

This philosophical emphasis on the activity of perception reflected common sense of the time. Margaret Miles has studied extensively how vision played a role in everyday Western visual culture. As she notes, there was no need to remind a medieval spectator of the active component to visual perception,

They recognized fully the extent to which *what* one sees is dependent on one's visual training, spiritual preparation, and active engagement with the object of vision. (Miles, 1985, p.65; see also Baxandall, 1972)

As I noted earlier (p. 31), where the objects people engaged with were images, it was considered possible for vision to move beyond the material world present to sense. In being presented with the image of the Trinity in the Rupertsberg manuscript, it was thought to be possible for visual attention to carry itself beyond the image, directing itself to the holy figure represented. It was thought to constitute a visual encounter that was at least to that extent non-corporeal. So understood, visual representation depended on a collaboration between image and viewer.

Let us not assume from the outset that the medieval attitudes towards images were rooted in superstition and animistic beliefs. If we avoid that hasty path, we can instead begin to make sense of late

medieval visual practices by seeing how they could rely on a sophisticated theory of vision; a theory rooted in an encompassing neo-Augustinian conception of psychology that was driven by considerations about the precarious relation between mind and body. Based on this specific conception of vision, the painter could be understood as capable of making biblical figures, heavenly scenes, and sacred narratives visible by means of her brushwork or drawing (although I have here spelled out the connection in more detail than usual, the point is emphasised by several authors, e.g. Miles, 1985; Deshman, 1997; Biernoff, 2002, chap.5).

The more the Augustinian framework comes under pressure from neo-Aristotelian conceptions of psychology, the more this specific understanding of the role and art of the painter starts to crumble. The relatively quick move away from a heavily spiritualised, active conception of vision and towards an understanding of vision as a purely optical process, had a visible effect on attitudes towards imagery and painting. As anyone familiar with the history of Western painting can attest, the transition from the art of the Middle Ages to the art of the Renaissance is dramatic.

The closer we move towards the Renaissance, the more we see a shift in conception of what the painter does. Alberti's famous treatise *On Painting* exemplifies the relatively sudden emphasis on and interest in recreating the optical conditions of sight (Edgerton, 2009; Dennery, 2005). The painter presents the eye with the same array of colours and shapes a landscape, a human body or building a would present it with. In this way, looking at a painting puts us in a visual situation that is like the visual situation we would be in were we to look at the represented item, given that the pattern of light reflected from the painting and received by the eye is formally similar. Instead of requiring a collaboration between image and viewer, visual representation becomes a purely optical affair. As Jeffrey Hamburger brings out, this leads to a growing uncertainty about the status of images (2000, p.50). If

what images can do is no more than put us in optical conditions that are merely like the optical conditions of vision of the real, the earlier conviction that images could bring us in attentive or spiritual contact with a saint or a religious scene was met with increasing scepticism. No doubt the painter could still be regarded a ‘visualiser’, but now only in the sense that they imitated or copied the sensible appearance of things. As is well known, painters embarked on this project with mathematical precision, yet in doing so transformed the direction of their art.

I have reached a diagnosis about two cultural frameworks, merely based on a comparison of what they thought about vision. It would be too simplistic to suggest that the shift in theory of vision explains the historical revolution that led to the painterly tradition of the Renaissance. That would pave over important additional factors, obscuring many internal disagreements, even among philosophers in the Augustinian tradition. It also would not yet account for the way Augustinian ideas about vision remained influential well into the eighteenth century. When Sellars lifted out the manifest image of humans as an object suited for philosophical reflection, he knew full well he presented an idealisation (Sellars, 1962, p.56). Still, bringing out how once unquestioned assumptions transform under novel scientific, cultural or religious pressures can reveal much about the roots of a philosophical controversy.

Cynthia Hahn writes that an “attempt to clarify dominating theories of vision and the parameters of their change is important to a historical understanding of image making and image reception” (Hahn, 2000, pp.169–70). My discussion in this chapter has laid out part of the history of thinking about images. Technically, both the conception of images prevalent in the later Middle Ages and the perspectivalist doctrines of the Renaissance proclaim that visual images present us with the visual appearance of what they represent. However, as I brought out earlier, that uniform description allows for disparate,

incompatible interpretations. We can now see how the late medieval and the Renaissance view discussed here each pick up on one way in which paintings can present us with the appearances of things. Recall the distinction between sensory and sensible appearances mentioned earlier—between appearances as perceptual events, and appearances as complex perceptible qualities. I suggest that precisely this distinction can capture how both lines of thought are concerned with appearances, but each in a different way. In the later Middle Ages we see painting being valued for its ability to make absent scenes appear to sight, while during the Renaissance painters are admired for their skill in producing surfaces that have the same optical properties as the scenes they represent, in short, for copying the way things look. Both sides of the discussion offer competing understandings of the claim, each opting for a distinct way of conceiving of the appearances that the painter renders with her brush.

This historical diagnosis teaches us something about our own conceptions. Retracing this aspect of the history of visual representation can uncover the factors that shaped views popular today (cf. Williams, 2002, chap.2). It brings out that how we have come to conceive of images and the way they represent is not independent of the way we conceive of vision as such. It teaches us that a dismissal of the claim that images make things visible as incoherent or mere wishful thinking cannot stand on its own. It is only warranted if we make substantive assumptions about the nature of sight—assumptions that have themselves repeatedly been the subject of philosophical disputes.

We can now see two distinct and well-grounded conceptions of visual representation. On the one hand the conception that images make the scenes they represent visible to a viewer, on the other hand the conception that they do not. It is important to see that their disagreement lies in a specific dispute over the active nature of vision. This result does not always stand out in the philosophical literature about this dispute. For example, John Hyman writes that

from Plato to the present day, there are two main contending doctrines about depiction. The first, which Plato states in the *Cratylus*, says that a picture represents an object by copying its form and color. The second, which appears much later in the history of philosophy and as a reaction to the first, says that a picture represents an object by producing a special kind of experience in a spectator's mind. The original source of the second doctrine is in Descartes's *Optics*. Pictures, Descartes argues, do not cause us to see the things they represent because they are likenesses. On the contrary, we call them likenesses because they cause us to see these things. (Hyman, 2006, p.2)

Following Hyman, it may seem that the central debate about visual representation is about whether or not representational properties are response-dependent (this indeed is Hyman's point). Of course, it is correct that we find those competing doctrines about images and visual representation in the Western tradition. One says that a picture represents an object by copying its form and colour, and indeed has often been attributed to Plato on the basis of what he suggests in the *Cratylus* and especially the *Republic* (though I will offer some reason to resist that interpretation in the next chapter). The other alternative identified here is the Cartesian suggestion that images *cause in us* a perceptual response (specifically, a judgement; Hyman's 'experience' is somewhat too general), a response that defines representation.

Yet this is not the only nor, if I am right, the most fundamental way to cast the dispute. For juxtaposing these two positions obscures a basic contrast with the medieval tradition that I considered earlier. Whatever their differences, the copy theory and the Cartesian model still agree on the claim that representation is independent of the perceptual activity of the viewer. As I have shown, a well established tradition that comes into prominence during the later Middle Ages and is well motivated by a specific theory of vision denies precisely

this. It views image perception as active, and opposes the idea that the visible world ever causes visual perceptions in us. On this late medieval view, a picture represents an object by playing a mediating function in a viewer's perceptual activity, enabling the viewer to look at the object by means of the image. If we take Plato and Descartes to be representatives of the two main doctrines about visual representation, we hide a deeper contrast between both their preferred theories and the medieval view. Then we are bound to neglect the view that pictures enable vision, not by causing any special experience, but by fulfilling a role in an ordinary perceptual process—a process that is in part due to our perceptual activity and instead of being a response to an image, is a perceptual collaboration with it.

Standing in a museum gallery, we ask a friend whether they see the small village on the horizon when they look at a painting by Raphael that hangs on the wall. Our question is natural, yet it presupposes that the settlement is visible in the context of the gallery. Do we really want to say this? Consider, merely by having seen the exhibits in the National Portrait Gallery one typically feels no entitlement to brag about having seen the royal family during one's visit to the United Kingdom. Do images make what they represent visible or not? The dispute about images I have identified in this chapter still has a tendency to pull us into two directions, as soon as we start reflecting on it (Alva Noë highlights this tension as well, see Noë, 2012, pp.83–84). Modern philosophy tells us that perception is wholly passive, and requires its objects to be present. This has led some to voice incredulity about the very idea of visions of the absent. Yet as soon as we survey more carefully how people have conceived of images and their relation to visibility, it becomes clear that such incredulity is at best an expression of parochial commitments.

1.4 Understanding visual representation

So far I identified two distinct and well-grounded conceptions of visual representation. On the one hand there is the conception that images make the scenes they represent visible to a viewer, and on the other hand the conception that they do not. I showed that which of these options one accepts is strongly influenced by one's standing background conceptions of vision as such, in particular by one's stance on the alleged passivity of sight. This makes for a fundamental point of dispute about the nature of images. (With 'fundamental' here I mean that this potential disagreement needs to be resolved before we can fruitfully ask further questions about images, such as questions about the kinds of things they can represent, about how their representational properties are determined, and so forth.) Resolving whether or not images make the scenes they represent visible requires a stance on the role of images in vision—a project that lies at the intersection of aesthetics and the philosophy of perception.

That the dispute about the role of images in perception is fundamental is already clear from practices of engaging with images. Visual representation is to a large extent an everyday phenomenon. People use and interact with images all of the time, whether for practical, playful, or scientific purposes. These practices exploit a prephilosophical grasp of images. More strongly, it is only because we already have a concept of an image or of visual representation that can we ask more specifically what images are. As such, any philosophical inquiry into visual representation must take this prephilosophical grasp into account. Now I have suggested, both in the Introduction (p. 19ff) and in this chapter, that our everyday engagement with images presupposes that images, as representations, fundamentally enable some form of visibility. In other words, the idea that images make the scenes they represent visible is a core dimension of our prephilosophical grasp of what images are. Some may deem this 'grasp' rather primitive, a mistake that a true philosophical account should help correct. I have

already expressed my doubts about such scepticism, but let that rest for now. The point is that the centrality of these everyday practices requires even the greatest sceptic to consider whether or not images make the scenes they represent visible.

There is another reason why the question about images is fundamental. Some of the theoretical concepts often used in merely characterising visual representation already presuppose an answer to the question whether images make the scenes they represent visible—or at least they come dangerously close to doing so. Let me highlight one prominent instance of this. Some philosophers beg the question against their opponents unwittingly, by taking it for granted that apprehending the representational properties of images is a form of *understanding*.

Flint Schier's book *Deeper into pictures* starts out from the assumption that by apprehending an image's representational properties we come to understand that image. Schier aims to give a fully general theory of visual representation, one that captures both high art and the images we find in newspapers, on cave walls, and in children's early attempts at drawing. The questions philosophy should ask about all of these images, he thinks, is what kind of representations they are, and how we understand them (Schier, 1986, p.1). Schier holds that the answers to these questions are connected: what is distinctive of images as representations (he calls them *iconic* representations) is the particular way they can be understood. Schier assumes that any philosopher interested in visual representation should start out from the question how we understand such iconic representations.

His answer is that our understanding of images is *naturally generated*. Understanding what an image represents and understanding what a written or spoken utterance means contrast in that knowing the meaning of a sentence in a language requires knowledge of grammatical rules and a vocabulary, whereas our understanding of images does not require any such knowledge (Schier, 1986, p.51). Images

as such lack a vocabulary, grammar and syntax. Instead, Schier argues, images “trigger” our ordinary capacities to recognise people and scenes by the way they look. When we see an image, such an understanding (an interpretation or judgement with a propositional form) is simply caused in us, not unlike the way in science fiction stories mental states are covertly implanted in people.

Now it is unclear to me how some things’ “triggering” our ordinary capacities to recognise things by their looks could account for the fact that images present us—*visually* present us—with the appearance of things. But currently I do not wish to criticise Schier’s account on its own terms. What matters is that, from the beginning on, Schier assumes that the distinctive way images represent may be approached by considering the distinctive way images can be understood. Seeking to classify representations in terms of how we understand them is not unreasonable. Several philosophers have suggested that, at least in the case of language, a theory of what a statement is about presupposes a theory of how we come to understand it (Dummett, 1973 is most famous for this programme). Schier’s project makes for an interesting attempt to develop a related theory of the way images are of something or other.

However, it is the starting assumption of such an attempt that is problematic. Schier must assume that apprehending the representational properties of images is a form of understanding. Yet, if pictures make things visible, then it is natural to think that the distinctive way we can apprehend their representational properties is simply an act of *seeing*. A viewer apprehends what an image represents in this way just if some of the things the image represents visually appear to that viewer. They get to see those things. And at least on a plausible conception of vision, acts of seeing are not themselves acts of understanding. (Of course, you may respond with “I see” when I try to explain something to you and you want to convey to me that you understood what I meant; but that is only a loose use of ‘see.’) Standing in my kit-

chen at 6:30am, I may see the raccoon in the backyard, scavenging for food. Merely by being in that state I do not yet understand anything, not, at least, if understanding something requires grasping something propositional. Of course, in seeing the raccoon I may come to see that a raccoon is scavenging in the backyard. Here I come to know or understand something about my surroundings on the basis of what I see. But, crucially, the perception itself and the knowledge I arrive at are distinct. The perception may be how I come to know that fact about the raccoon, but that does not make it itself an act of knowing (the point is argued for at length in Travis, 2004).

At the end of his book Schier concludes triumphantly that he has been able to explain what an image is without having to assume that an image makes something absent visible (Schier, 1986, p.208). Yet this can hardly come as a surprise, given that Schier's assumption that apprehending the representational properties of images is a form of understanding had already closed off that option right from the start. It has ruled out from the beginning the idea that images represent by making absent things visible. For this reason, Schier's encouragement that an account of what images are should start out from the question how we understand images is far from innocent. This proposal must be resisted by someone who thinks that images represent by making scenes visible in their absence.

Theoretical frameworks are not neutral. The injunction to focus on understanding in explaining images has already implicitly taken a stance on the far more basic question whether or not images make the scenes they represent visible—a question which will have consequences for the way we conceive of the nature of representational properties, and how we apprehend them. For this reason, then, the question whether images make the scenes they represent visible or not is fundamental, and must not be skipped.

1.5 The way ahead

The central disagreement about visual images is a dispute about the role they can play in perception. When we see an image, we are made aware of the visual appearance of some figure or scene. I have shown how it is not obvious what exactly this comes down to. What is it for an image to present us with the visual appearance of a figure or scene?

As long as we distinguish between sensory and sensible appearances, the options before us are clear. Either images can make for a sensory appearance of what they represent, that is, they can enable an event in which absent scenes visually appear to a viewer, as accepted by medieval authors. Or alternatively, images merely copy the sensible appearance of a tridimensional scene, that is, copy the visible qualities that determine the way such a scene looks—a view popularised in the Renaissance.

An important drive behind both of these views are specific theories of vision. Crucially, absent scenes can visually appear to a viewer only if vision as such does not wholly depend on what is present to the senses. A dominant tradition in the later Middle Ages conceived of vision as wholly active. During that time, as I have shown, it was frequently taken for granted that visual representation depended on an active collaboration between image and viewer.

It is not obvious that we can accept the perceptual possibilities that medieval authors saw, given the way we conceive of vision. Current thinking about perception and human psychology differs considerably from the Augustinian doctrines I discussed. Perception is standardly conceived of as a passive event in which some external object acts upon a perceiver, making them aware of the presence and qualities of that object. Further, in part because of the development of modern optics, few scholars today feel the problem of vision as a pressing concern. They typically reject the idea that vision is mediated by some sensible intermediary, be it visible species or sense-data. Yet as I will bring out in the chapters that follow, there is enough to be salvaged

from the history I retraced. I will defend a number of core elements of the late medieval ideas about images, and show that they are in fact already embedded in the way people today naturally think about how they perceive the world.

There is good reason to pursue this project. The main alternative in making sense of the fact that images can make us aware of the visual appearance of some absent scene is the copy theory of visual representation. That view today stands out as the standard view among philosophers of perception and aestheticians. In the next chapter I will show that the copy theory of representation is unacceptable as a philosophical account of the role images play in perception.

2

Copying appearances

Plato gives a notorious outlet for the idea that the painter merely copies the look of things. In the *Republic*, Socrates talks about the art of painters, and how they make a guild capable of producing all the things that other kinds of craftspeople do severally. The painter's art is a marvel. With a mere brush stroke she is able to render any item from the visual world, and place it before our eyes! Whatever visible thing, artefact, or organism you come up with, the painter is able to show it to you with her painting. Glaucon, Socrates' interlocutor, deems that if painters are so artistically versatile, they pull off a truly clever and wondrous feat. But Socrates dismisses his admiration with a satirical quip: the way the painter produces all things may be easier than you think!

SOCRATES: You could do it most quickly, if you should choose to take a mirror and carry it about everywhere. You will speedily produce the sun and all the things in the sky, and speedily the earth and yourself and the other animals and implements and plants and all the objects of which we just now spoke.

GLAUCON: Yes, the appearance of them, but not the reality and the truth.

To this, Socrates responds with delight.

SOCRATES: Excellent! And you come to the aid of the argument opportunely. For I take it that the painter too belongs to this class of producers, does he not?

GLAUCON: Of course.

SOCRATES: But you will say, I suppose, that his creations are not real and true. And yet, after a fashion, the painter too makes a couch, does he not?

GLAUCON: Yes, the appearance of one, he too.

(Republic, 596^e)

What Plato writes about painting in the tenth book of the *Republic* mainly concerns the relationship between painterly representation and the visible world. Stephen Halliwell stresses this point (Halliwell, 2002, p.137). As Halliwell observes, at least in this specific discussion about painting and mirrors Plato does not seem to be primarily concerned with expounding general metaphysical insights about truth or being. On closer inspection it is a puzzle about the cognitive value of painting, when painting is understood as copying the visual appearance of things. With Socrates' provocative mirror analogy Plato wants to present a puzzle specifically about those representations that address the sense of sight.

In Classical Greece, there seems to have been a growing conviction that the painter's skill consists in rendering or preserving the way things look from a specific point of view (for a seminal study of this artistic movement, see Bruno, 1977). By arranging coloured pigments on a bearer, the painter was thought to be able to capture all of the visible world with her art.

Through Socrates' words, Plato presses us to reconsider the idea that painters copy the appearance of whatever they encounter in the visible world. Are we willing to accept its implications? If painters indeed produce objects of which the appearance copies how some scene

or individual looks or could look from a specific viewpoint, then their skill does not require knowledge of the world. Just as someone unable to read Latin can still succeed in copying one of Vergil's poems, a painter can copy the way a human torso looks without knowing anything about anatomy or the nature of humanity. A painter who sets out to copy the mere look of things does not get us any closer to knowledge or reality. As Halliwell writes,

Socrates issues a challenge to those who value visual art, just as he later does to the lovers of poetry, to find a justification for pictorial representation that will endow it with something other than the cognitively redundant value of merely counterfeiting the "look" of the real. (Halliwell, 2002, p.139)

This means that the cognitive value of painting is unclear. Painting is therefore inherently problematic—or so Plato seems to suggest.

Commentators typically attribute to Plato himself a 'copy theory' of visual representation, and cast his criticism of painting in light of that view. However, I suggest that this interpretation is not mandatory. Plato need not be defending a copy view. Instead of a strong philosophical claim about the nature of representation, Plato could be defending the more restricted claim that painting *done in a certain way* is worthless. It is likely that Plato was well aware of the details of the craft (ancient biographies go so far that Plato studied painting himself, though this is potentially apocryphal, see Riginos, 1976, p.42). His criticism is more plausibly directed at a specific painterly *style*. Just as someone may dismiss the Impressionist movement in nineteenth century Paris as "superficial" without rejecting all of painting, Plato's objections may have a target more precise than the art of painting as such.

Read in this way, the notorious criticism of the tenth book of the *Republic* becomes an art theoretical one. Plato's use of the mirror in

the passage is consciously polemical. It merely states that *if* and insofar as painters copy appearances, their art does not manifest genuine knowledge of reality. For, even a mirror can do *that*. Surely, a mirror lacks skill or knowledge. Moreover, if painters in copying appearances do not need any real knowledge of reality, then we, as viewers, can not distil any knowledge of reality from their products either. Therefore, *if* that is what painting comes down to, then it just is not of much cognitive value (see also Janaway, 1995, p.118). Plato's discussion would still ban the Greek painters of his day from the state, given that those painters, like the impressionists in Paris, were staunchly in the grip of a controversial, deplorable artistic fashion. Yet it does not mean Plato himself thought that the painterly styles of his day exemplified the nature of representation in pictures (several scholars have suggested that this is indeed how we should read Plato, e.g. Bruno, 1977).

Even if Plato himself did not defend a copy theory, the ideal of copying appearances that is embodied in the fashions he criticizes—the project of representing something in a picture by copying its look—is nowadays still regarded as a viable philosophical account of visual representation as such (see Halliwell, 2002, ch.5 for a historical discussion). More strongly, I regard versions of it as the main competitors to the view that images represent by making things visible. In what follows I will show that, read as a theory of visual representation, the copy theory makes a muddle out of the relations between viewer, painting, and the visual world. Two observations about painting seem to me beyond question. First, that a painting, as an image, represents what it does specifically for the eye. What it represents must be visually detectible—we must be able to see the painting as an image. Second, that a sufficiently skilled painter can visually represent anything that she can see. My argument in this chapter is that the copy theory of representation can only account for these two facts at an unacceptably high cost. This, I will conclude, is sufficient reason to consider more seriously the idea that images, instead of being mere

copies of the looks of things, in fact make those things visible to the eye.

2.1 Ringers and reproductions

To assess whether the copy theory of representation is viable, we need to get a more precise understanding of what copies are. According to a notorious story by Jorge Louis Borges, the author Pierre Menard composed several passages that coincided—word for word, line by line—with excerpts from Cervantes' *Don Quixote* (Borges, 1964). Borges invites us to imagine that a page of Menard and a page of the much older Cervantes are both originals, despite their having the same words on them, share their word order, and have the same punctuation pattern and capitalisation. The story presents them as different works of art that are, at least in a superficial sense, dead ringers. Did Menard copy Cervantes?

In analytic philosophy Borges' example is familiar from late twentieth century discussions about the conditions of identity of works of art (see e.g. Goodman, 1976; Wollheim, 1980; Levinson, 2012, p.50). Are works of art necessarily concrete particulars, or do they allow for multiple instantiations? When do we have another instance of an earlier work? When do we have an entirely fresh one? It is not unlikely that Borges himself had such philosophical questions in mind when composing his amusing thought experiment. Philosophers' intuitions have diverged on what to make of it. But not only does reflection on Menard's fictional composition jog our intuitions about originality in art. It also helps us bring to light two distinct ways in which something can be a copy.

According to the story, Menard conceived of the fruits of his labour as a piece of original writing, not as a copy. Borges defends him by highlighting the differences between Menard's and Cervantes' work. Menard's story is almost infinitely richer than what Cervantes wrote,

Borges suggests. He compares a line of Cervantes,

... truth, whose mother is history, rival of time, depository of deeds, witness of the past, exemplar and adviser to the present, and the future's counselor.

with a line of Menard,

... truth, whose mother is history, rival of time, depository of deeds, witness of the past, exemplar and adviser to the present, and the future's counselor.

Borges thinks the comparison is revelatory, and writes how, written in the seventeenth century by a 'lay genius', Cervantes' enumeration does not get beyond a rhetorical praise of history. On the other hand, Menard's witty image of history as the mother of truth is astounding. Manifestly in the grip of William James' pragmatism, Menard presents historical truth not as what has happened, but as what we judge to have happened. This makes the second phrase above, in many respects, a better work of prose composition.

Despite Borges' cunning defence, someone may still insist that, no matter how we interpret the fruits of Menard's labour, the second passage just is a word for word copy of the first. Even though Menard did not set out to copy Cervantes' work, and explicitly renounced any intention to do so, in some sense he nonetheless ended up doing just that. So did Menard copy or not?

What Borges' story brings out masterfully is that we can have it both ways. This is because there are two distinct ways in which something can be a copy. A strict formal resemblance suffices for something to be a copy in one of these ways, but not in the other. On the one hand, Menard's work can count as a copy of *Don Quixote* because it is a textual *ringer*. Being a ringer of something requires no more than a qualitative resemblance to it. On the other hand, because Menard self-consciously avoids mechanically transcribing Cervantes' text, his

work could also fail to count as a copy because it is no *reproduction* of it. Assuming that Menard is sincere about how he worked, he did not take a past exemplar and mechanically generate a new incarnation. Yet that is precisely what a being a copy in this other way requires.

Ringers and reproductions, so understood, are distinct kinds of copies. It is for this reason that, paradoxically, some copies can be originals.

Let me say a bit more about this distinction. Being a copy in the sense of being a ringer requires no more than qualitative resemblance. For this reason, something can be a ringer quite accidentally. Two people may be ringers by coincidence, for example. Further, there is no such thing as a 'poor' or 'inaccurate' ringer. Naturally, someone may have come close to being a ringer for the sheikh, but this means not that their countenance is somehow imprecise, but just that it is no more than notably similar to that of the sheikh. By contrast, copies in the sense of reproductions require not qualitative resemblance, but a certain form of production. Reproductions are always the product of some mechanism designed or evolved to preserve features of an original (cf. Evans, 1982, p.125). Implementing that procedure can always be more or less successful in light of its ends. Therefore, reproductions allow for greater and lesser degrees of accuracy. Something is a reproduction of something else only if it is the output of such a mechanism, and for that to reason always preserves the features of an original to some degree.

What do these two accounts of copying mean for how to understand images? The French painter and art theorist Roger de Piles (1635-1709) claimed that the essence of painting is the imitation of visible objects by means of form and colour. He meant in part that the colours of her paint and the way she applies them are the principal means the painter has at her disposal. Especially if we limit ourselves to painting, we must acknowledge that an image consists of no more than some pigments and brushwork applied to a surface. To produce

an image is to alter or configure the way a surface looks. If images are copies, their copying is carried out in the domain of visible appearance. We therefore cannot expect images to be *dead* ringers or *exact* replicas. Rembrandt did not represent gold braid with gold, but used some well-placed strokes of umber paint. Jan Weenix did not paste real feathers on his canvas to represent the details of a peacock, but used clever brushwork to get the bird's appearance right (see also Kulvicki, 2014b, p.67). If images are visual copies, then what they copy is limited to the visible side of things.

Both copy theories claim that in producing an image, a surface needs to be altered such that its look copies the visible appearance of what it represents. The two ways of thinking about copies result in two versions of the copy theory of representation. On the one hand, a painter could be thought to produce a visual ringer. They could be preparing a marked surface that happens to visually resemble some scene. On the other hand, she could be thought to make a reproduction of some object or scene. The painter could produce a marked surface following some mechanism designed to preserve the visual features of an original, and so aim to reproduce that object or scene's visual appearance.

These two versions of the copy theory are subtly different. If a painter produces a ringer, then an unfortunate brush stroke may make her image represent some entirely different scene, given that it will now visually resemble something else. An untargeted smudge may make an emerging portrait of the sheikh into a portrait of the sheikh's brother, given that it now resembles him instead. Yet if a painter reproduces an appearance via some mechanism, then the scene represented remains the same despite such an unfortunate gesture. Such a slip could at best make the image inaccurate. The copy theory has its appeals, but I suggest that in neither of its guises it is attractive to pursue. Let me begin considering the view that an image represents an object in virtue of being a specific sort

of visual ringer—also known as the *resemblance theory* of visual representation.

2.2 Sight's motley crew

What is the relation between an image and what it represents? The answer that an image represents an object in virtue of being a specific sort of visual ringer may very well be the oldest of its kind (Newall, 2011, p.91). Recall, early atomists such as Democritus thought that parts of the material world were visible only because they emitted visible films of atoms to the eye. Lucretius likened those films to the skin of a snake—something that conveys to us the visible qualities of the animal, yet may be distinct from the snake itself (Hyman, 1989, p.3). These immediately perceptible films or *eidola* emitted by material things were taken to be images of the material world; images that represented distal objects by visually resembling them.

The resemblance theory of visual representation centrally claims that an image represents something only if it visually resembles it. The claim is explicit in Catharine Abell's recent work, when she suggests that for x to depict y , x must resemble y in the relevant visible respect—or at least it would have to do so if y existed (Abell, 2009, p.217). Abell advances this as a necessary condition on visual representation. John Hyman, in a similar spirit, maintains that “representation in the visual arts depends on resemblances between works of art and the objects they represent” (Hyman, 2012, p.149; for other versions, see Hopkins, 1998; Greenberg, 2013; Blumson, 2009).

Take the wood block print of a group of Japanese scouts by Yūsai Toshiaki (Fig. 5). The differences between this print, published in 1895, and the skirmish it represents are countless. The print is a smallish, flat surface covered in dry layers of coloured ink; yet it represents a tridimensional scene with fighting soldiers, horses, and the mountain Daheshang in the background. However, such differences



Figure 5: Yūsai Toshiaki, *Kinshū Daioshōzan sekkō funsen no zu* (Japanese scouts fighting at Daheshang Mountain near Jinzhou), 1895. British Museum

need not take away that, as Hyman observes, “there is a strict and invariable relationship between the shapes and colors on a picture’s surface and the objects that it depicts” (Hyman, 2006, p.73). Despite their differences, Toshiaki’s print and the figures it represents have key chromatic and formal aspects in common.

How to understand the chromatic and formal respects in which images resemble what they represent? Different authors have offered slightly different suggestions. Hyman, who offers what I take to be the most effective analysis, claims that an image visually resembles both the *occlusion shapes* and the *aperture colours* of the scene it represents. They can be characterised as follows:

Occlusion shape When we trace the outline of the tabletop on a pane interposed between us and the tabletop, this results in an elliptical shape. This is the object’s occlusion shape relative to this particular line of sight; it is the shape the mark on the pane would have to take in order to occlude the tabletop precisely.

Aperture colour When we look through a cardboard tube at a coloured surface, the colours we see only depend on

the light reflected from that surface, and not on illuminants or ambient lighting external to it, because the tube prevents us from seeing the wider scene. The colours we see are the surface's aperture colours.

According to Hyman, by copying the occlusion shapes and aperture colours of a scene, a painter is able to copy a stable and distinctively visible aspect of that scene. The occlusion shapes and aperture colours in the central parts of the print will resemble those of a uniform and kepi worn by an officer as represented by the print—there is an exact resemblance between these shapes and colours (Hyman, 2012, p.143).

So understood, the resemblance view conceives of an image's representational features as intrinsic to it. An image represents what it does because it has the chromatic and formal properties it has, and chromatic and formal properties are intrinsic features. This exclusive focus on colour and form is no coincidence. All versions of the resemblance theory must accept that only such simple visual resemblances matter for representation. This is for two related reasons. First, as I noted earlier (p. 69), a painter only has pigments at her disposal—her art consists in arranging colours and forms in a specific way. Second, and more fundamentally, only visual resemblances matter, because images represent things specifically to sight. The resemblance theory holds that purely visual resemblances constitute the representational properties of an image. If this were not the case, the resemblances might as well escape the eye. These reasons are related: the painter's materials take the form they do because images represent things specifically to sight.

Note that the resemblance theory does not maintain that an exact resemblance between chromatic and formal aspects is sufficient for representation. Any visible surface will resemble indefinitely many things in this way. Many philosophers have remarked that if such resemblances were sufficient for representation, it would make large swathes of the visible world representational. Such ubiquity would

render the concept of visual representation vacuous, they claim. Instead, proponents of the resemblance theory typically suggest that something only is a representational image if it came about in the right way, or if its resemblances are intended to play a representational function (Hyman, 2006, p.256n15; Abell, 2010a, p.275; Hopkins, 1998, p.71). However, once an item did come about in the right way or is used appropriately, as an image it represents something only if it visually resembles it.

‘An image represents something only if it visually resembles it’. In this claim lies the resemblance theory’s fundamental weakness. In claiming that representation depends on visual resemblance, all versions of the theory in effect circumscribe what can potentially be represented by an image. Soviet era officials proclaimed on ideological grounds that genuine painting only dealt with subject matters that are proletarian, part of everyday life, and serving the communist cause. Not unlike this, but now on purely philosophical grounds, the proponent of the resemblance theory claims that painters can only represent visually what paintings can visually resemble. As I will explain shortly, they must hold that the painter deals only with extended, coloured bodies. As may be clear on reflection, both convictions vastly restrict the actual range of possible subject matters of the painter. At least if she is skilled enough, a painter can set out to represent anything that can be seen.

That the painter’s actual subject matter is the visible world is beyond question. Although sometimes attributed to the art theorists of the Renaissance, this assumption goes back to antiquity. Indeed, it forms the background of Plato’s discussion in the *Republic*, as covered above. If something can be seen, it can be reflected in a mirror, and it can just as well be represented by means of an image. This same assumption allowed some of the medieval image enthusiasts to defend the representation of saintly figures. When a sceptic suggested that an angel could never be represented in an image, a common retort was

that some pious individual had in fact *seen* an angel in a nightly vision or in some other kind of revelation. This was taken as proof that the angel could be represented after all, given that whatever can be seen can be represented in an image (Brubaker, 2012). Much later, the principle is still embraced by authors very far removed from the medieval tradition. Gotthold Ephraim Lessing (1729-81), for example, takes it to be a commonplace that the power of painterly representation “extends over all visible nature, of which the beautiful forms but a small part” (Lessing, 1874, p.28). In the twentieth century similar a assumption drives Wittgenstein’s discussion of pictorial representation as a criterion of visual experience. If we were not able visually to represent everything we saw, then the “representation of ‘what was seen’ ” could never fulfil such a criterial role (Wittgenstein, 2009, pp.197–198).

The painter can represent whatever she encounters in the visual world. But this puts pressure on the claim that representation requires resemblance in form and colour. This is because some things are visible without having the chromatic and formal properties that could make such resemblances obtain. Events and states such as Gerald’s swimming, a sudden sneeze, or the messy state of the hotel room, are not coloured, nor do they have a particular shape. Yet I can very well observe acts of swimming and sneezing, or the state the hotel room is in. These events, states and acts are all visible. Given that the painter can represent whatever she encounters in the visible world, she can also represent Gerald’s swimming and sudden sneezes. But given that sneezes and the swimming as such lack shape and colour, resemblance in form and colour simply cannot be a necessary condition on representation. To illustrate this more clearly, let me discuss a specific kind of example, the representation of visible events.

When Toshiaki carved his woodblock, he produced a painted image of a skirmish between young Japanese army scouts. Although his aim was to document what is now known as First Sino-Japanese War

of 1894–1895, it is hard to verify whether he set out to represent a particular event as it actually happened, or merely aimed to represent some type of event that could stand for what went on near mountain Daheshang. Either way, his woodblock image stunningly represents the event of a soldier's being blown away by a forceful equine kick. Similarly, the print represents the slash of a captain's sword. It represents the captain as well, of course, and the sword. But Toshiaki's print is marvellous in that it also represents the slash (and the kick). Observing all this, the resemblance theory of representation would have to say that the print represents what it does because the image (or some part of it) visually resembles what it does. But that cannot be right. For the represented events of slashing and kicking lack the colour and form requisite for any such resemblance. A kick as such is shapeless and colourless—it's nature is temporal, though of course the bruises it causes may be blue and round.

Events, at least typically, lack colour or form (flashes have been considered exceptions, though the point is controversial; see e.g. O'Connor, 1945, p.70). The precise nature of events is disputed, but most philosophers agree that events principally differ from bodies in the way they relate to space and time. Unlike bodies, events are generally understood not to have spatial dimensions. Furthermore, the exact period they occur in time is much more central to particular events than temporal moments are to bodies (so-called four-dimensionalists may disagree with the details here; cf. Sider, 1997). However, the lack of spatial dimensions does not take away that events can be seen. As P.M.S. Hacker writes,

One can observe, watch, look at the falling of leaves or the performance of a play. But events have neither shape nor colour, just because they do not have spatial dimensions, do not fill space, and do not consist of stuff (Hacker, 1982, p.8).

Hacker's point is that, despite their difference from material bodies, we naturally think of events as potential visibilia. To me it seems undeniable that events can be seen. Moreover, the thriving field of perceptual psychology dedicated specifically to the visual perception of events lends support here (W. H. Warren & Shaw, 1985; Zacks et al., 2007). It is also part of common sense. Consider for example the formal requirement that the signing of every legally binding will needs to be witnessed by at least two adults. Surely, the signing of a will is an event. Denying that we can see events would come at a considerable legal cost.

We can see movement and change. Some even say that we can see the passage of time itself (e.g. Phillips, 2014). If an image could only represent what it resembled in chromatic and formal respects, then those visibilia that lack chromatic and formal features could not be represented. Yet they clearly can. Visible events are but one example of a broad family items that fall outside the comfortable category of bodies. What about Jinan's anger (a state of mind)? The brightness of St. Peter's Basilica (a state of illumination)? The buttering of toast (a process)? The shininess of a helmet (a property)? Examples are legion. The resemblance theory requires that the painter's subject matter is limited to those things that possess colour and form—which most naturally would come down to a restriction to visible bodies. Painters have never felt constrained in this way, and rightly so. The restriction should strike us as misplaced. The painter's subject matter is the visible world, and we naturally conceive of the visible world as a diverse realm comprising visible states, processes, events, and properties, in addition to bodies.

2.3 Ringers in the realm of vision

Traditionally, philosophers have granted the resemblance theory that it is able to capture at least a necessary condition for representation.

This is because they take on board the commonplace that resemblance is ubiquitous; everything resembles everything else to some degree (Nelson Goodman's version of this suggestion is well known, Goodman & Elgin, 1988, p.112). This is easily extrapolated to the case of visual resemblance. A ripe tomato and a sewing needle are visually very dissimilar, but still they have some visible properties in common as well (for instance, both have a smooth exterior). Does that mean that the claim that an image visually resembles what it represents is trivially true? I have made clear why we should resist that temptation. There are many visible items that images do not—could not—visually resemble.

Someone could suggest that my argument here overlooks something about the way we see the world. For example, it could be that in real life we only see events because we see the bodies that partake in it. But if that is right, then could an image not represent a visible event by representing some of its visible participants? The peasant wedding itself may lack colour and form, but the people celebrating it—their hands, their dresses, their food, their props—do not. A painter may represent the wedding by representing its participants. And the representation of wedding guests and their accessories could straightforwardly be explained in terms of visual resemblance.

I do not think this is right, for two reasons. True, sometimes we do seem to see an event by seeing the bodies taking part in it. Yet this it is not the only way in which we can see events. When cheering the riders of the Tour de France, you may see in the corner of your eye the event of some thing's whizzing past. There is no reason to suppose that seeing that event depends first and foremost on your seeing any material body that partakes in it. Here what we see is first and foremost the event. If we saw, say, the cyclist that whizzed passed in that way, we did so because we saw the event. For the debate about images, this means that events may be represented in a way that does not depend on the representation of bodies. Umberto Boccioni's Fu-

turist painting 'Dynamism of a Cyclist' (Fig. 6) is a serious attempt at doing just that. If Boccioni's image represents the cyclist itself at all, it does so only because it represents the event in which that cyclist participates, and not vice versa.



Figure 6: Umberto Boccioni, *Dinamismo di un ciclista* (Dynamism of a Cyclist), 1913. Gianni Mattioli Collection

But there is a second reason to reject the suggestion that an image can represent events by representing the objects that partake in them. Even if we set cases of independently visible events aside, the suggestion lacks force even on the resemblance theorist's own terms. It threatens to undermine the theory's main claim. Suppose that events could indeed be represented by representing the bodies that partake in them. That does not establish anything about any visual resemblance between image and event. On this suggestion, images would still represent events without resembling them. Assuming that representation without resemblance is possible simply goes against the central claim of the resemblance theory.

Behind the difficulties with the resemblance theory of visual repres-

entation are three commitments that, when accepted conjointly, are in conflict:

1. All visible things can be visually represented.
2. Some item x can only be visually represented if the chromatic and formal properties of its representation visually resemble x 's chromatic and formal properties.
3. Not all visible things have chromatic and formal properties.

The first commitment is a general assumption about the relation between the painter and the visible world. I take the second to be the core claim of the resemblance view. The third is a general assumption about the nature of the visible world. Accepting all of these commitments, we saw, casts a shadow over the resemblance theory. Is there a version of the view that could reject at least one of these claims, and still explain compellingly the way images represent? I suggest there is not. Giving up on one or more of the above three commitments will not make the resemblance view any more attractive.

Representing 'non-depictively'

I will begin by considering the first commitment, namely that all visible things can be visually represented. That the painter is able to represent in her painting everything that can be seen I take as read. But could it not be that, although the painter can represent all visible things, she cannot represent all visible things in the same way? Consider a familiar example. A painting can represent the holy spirit by representing a dove. While the dove is visually represented, the holy spirit surely is not visually represented; it is represented because of a conventional, symbolic relation between it and the figure of a dove. The representation of the dove is doing the work in this image. Similarly, the suggestion goes, an image can only ever *visually* represent bodies, but it can nonetheless represent visibilia such as events and

states in some derivative, non-visual way. This suggestion is made by Hyman (2006, p.62) and Robin Le Poidevin (Le Poidevin, 1997, p.182).

On reflection, this response is a sophisticated reincarnation of the objection I considered a few paragraphs back. Le Poidevin suggests that

Depiction is just one form of representation. Essentially, depiction is representation by means of resemblance. A picture of a sheep depicts a sheep by resembling it (in certain respects). But pictures represent more than they depict. In particular, they may represent aspects of time that they are unable to depict. [...] temporal order is non-depictively represented by spatial order [...]. (Le Poidevin, 1997, p.182)

The fact that an image represents an extended body and the fact that the image represents an event necessarily are distinct facts, so the suggestion goes. They must be distinct because the one fact explains the other (cf. Hyman, 2006, p.63ff).

There is something odd about the distinction Le Poidevin presses. Even if we concede that not all things in a picture are represented visually (or, in his terms, 'depictively'), it is not at all clear that only material bodies can be represented visually. Someone could suggest, not unreasonably, that Boccioni represented a cyclist *by* representing the event of his whizzing past. Here it seems that the event is represented visually, whereas the cyclist itself is not visually represented, but merely implied or suggested. It would be question begging to assume on a priori grounds that visible events, such as that of a cyclist whizzing past, can never be represented visually. Yet the resemblance theory would only be saved if material bodies exhaust what can be represented visually.

Further, the use of a distinction between two sorts of facts of representation is severely limited. Indeed, it works for purely conventional or symbolic cases. The fact that a painting represents a dove and the fact that it represents the holy spirit are distinct facts—this is precisely why the representation of the dove can secure the representation of the holy spirit. With the right set of conventions in place, the same visual representation could be used to represent symbolically more or less anything. Yet things are different when it comes to cases that rest far less on symbolism and convention. In an image that represents the slash of a captain's sword, the relation between the slash on the one hand and the combination of the captain and his sword on the other seems much tighter. Presumably, the defender of the resemblance theory would maintain that the image represents the slash because it represents this captain and that sword. But the 'because' here hardly captures an instrumental relation. Instead, the relation is constitutive; the representation of the slash just is the representation of the captain and his sword. If this is right, then the fact that it represents the one and the fact that it represents the other aren't distinct facts. Hence, we lack any reason to accept the distinction Le Poidevin urges us to accept. Without this distinction, this case in defense of the resemblance theory collapses (for a similar criticism, see Wollheim, 1980, §26).

Resemblance in the visual field

How about the second commitment, namely that something can only be visually represented if the chromatic and formal properties of its representation visually resemble its chromatic and formal properties? Some may suggest I have misunderstood the nature of the relation at the core of the resemblance theory. What is supposed to resemble what? I assumed that it holds that visible properties of the image must resemble visible properties of what it represents. No doubt the resemblance theory typically takes this form. However, a close relative of the

theory could be immune to my criticism. It would focus not on resemblances between visible features, but on resemblances between experiences. Instead of assuming that a picture represents something only if it visually resembles that thing, we could suppose that a picture represents something only if the experience of looking at the picture resembles the experience of looking at the thing. The mistake of typical versions of the resemblance theory would be that they construe representation as a resemblance relation between *bodies*, while not everything we can see is a body. Instead, what they could have done is think of representation in terms of a resemblance between possible *experiences*, because everything we see is seen by having such experiences.

Similarities between experiences are rife, of course. The experience of frostbite resembles the experience of a burn, and looking around in a pitch dark room may resemble the experience of having one's eyes closed. But it is unclear that these resemblances are of any help. Images, we saw, represent things specifically to sight; whatever represents or gets represented is and must itself be *visible*. And though frostbite *feels like* a burn, peering around in a dark room surely doesn't *look like* having one's eyes closed. In order for such resemblances to be visually salient, we would have to suppose that the experiences themselves possess some visual qualities of which we become aware.

Christopher Peacocke defends a version of the resemblance theory along these lines (Peacocke, 1983; Peacocke, 1987). He assumes that, indeed, visual experience itself possesses visual qualities of which we become aware. It comprises a *visual field*. With this he means not just a field of view, but some entity that itself has something like chromatic and formal properties. One's current visual field can visually resemble another possible visual field. According to Peacocke, an image represents something by affecting our visual field in a specific way. It enables us to become aware of a resemblance between formal features of our visual field when looking at the image, and formal fea-

tures of a visual field one could have if one looked at the represented object. Peacocke's proposal would meet the condition that images represent things specifically to sight, because our perception of images itself has visual qualities and for that reason could strike us as visually like perceptions of what the images represent (see also Budd, 2008b).

Of course, if one already believed in visual fields of the sort Peacocke postulates, then the idea that visual perception as such is first and foremost a mode of awareness of the visual qualities of our visual fields may seem attractive. On such a theory of vision, one could accept that the representational role that images play in perception is to affect our visual fields in a specific way. Yet if visual fields were postulated in an attempt to rescue the resemblance theory of visual representation, then it strains the imagination. It is not at all clear why we should accept the idea that experiences themselves have visual qualities. Indeed, Mark Johnston has recently argued that the very idea of such a visual field is incoherent (Johnston, 2011). If it turns out that the resemblance theory requires there to be such visual fields, then that would make the theory much more controversial than is commonly believed.

I do not need to settle this dispute here. Suffice it to say that the alternative proposal defended by Peacocke departs from the typical version of the resemblance theory, and that as a defence of the resemblance theory it comes at an unreasonably high cost.

Visual austerity

I now want to turn to the third of the three commitments that jointly challenge the resemblance theory of visual representation. This is the commitment that not all visible things have chromatic and formal properties. The commitment is formulated negatively, because on a natural conception of the visible world, it is not the case that it comprises only the bearers of chromatic and formal properties. We can

see much more than that.

When laying out his theory of visual art, Erwin Panofsky rejects precisely this common sense view of what—strictly—can be seen. He does so by a familiar sceptical exercise (see Clarke, 1965 for a helpful diagnosis). In his seminal work on iconology and meaning in the visual arts he writes that

When an acquaintance greets me on the street by removing his hat, what I see from a *formal* point of view is nothing but the change of certain details within a configuration forming part of the general pattern of colour, lines and volumes which constitutes my world of vision. When I identify, as I automatically do, this configuration as *object* (gentleman), and the change of details as an *event* (hat-removing), I have already overstepped the limits of purely formal perception and entered a first sphere of *subject matter* or *meaning* (Panofsky, 1972, p.3).

This passage can be read as a statement about the visible. Panofsky seems to urge that only general patterns of colour are ever *really* visible; only lines and volumes constitute one's world of vision (a similarly austere view of our visual world is assumed by Descartes and Lessing). Should it be surprising that Panofsky presents this view of the visible precisely when attempting to restrict what can be visually represented in the arts? Panofsky talks as though he transfers the results of his 'analysis' of everyday vision to a theory of representation in visual art. The primary or natural subject matter of a work, he thus proposes, is given merely by certain configurations of lines and colours (Panofsky, 1972, p.5). However, it seems more accurate to say that in fact Panofsky does precisely the converse. Panofsky's observations about the limitations of the visible emerge as nothing but an expression of his theoretical commitments about the art of painting.

Panofsky's gesture can figure as an example. The assumption that

our visual world as such is limited to general patterns of colours, lines and volumes makes the idea that visual representation is a matter of resemblance in colour and form immensely plausible. But it is an unreasonably austere conception of the visible. We naturally take ourselves to see so much more: processes, events, and activities that themselves lack colour and form. If the resemblance theory about visual representation can only seem well-motivated on a much more austere model of the visible, then it is not clear what attraction it really holds.

“Fatal difficulties,” Dominic Lopes writes, “have made resemblance theories historical curiosities.” (Lopes, 2005, p.26) We see how, as a diagnosis, this would be too simple. On specific understandings of our visual world, the resemblance theory may be defensible. Yet as I have shown, such commitments make the resemblance theory much more controversial than typically thought. When it makes such substantive assumptions, a resemblance theory of visual representation is far from trivial. More likely, it will face resistance from anyone attracted to a more natural conception of what can be seen.

2.4 The visible past

Photographers are frequently taken with their equipment to capture the appearances of things in a purely optical and mechanical way. If you consider photography to be a paradigm of visual representation, you may find it attractive to assume that images represent by virtue of their historical connection to things placed in front of the camera or to things that were inputs to some comparable mechanism. Recall, Plato satirically likened the art of the painter to the skill required for holding up a mirror. The mirror shows whatever you point it at. To people impressed by the very possibility of photography, Plato’s analogy may suggest that, similarly, an image represents whatever a camera (or equivalent mechanism) was directed at. In this way, the core

of visual representation is a historical tie between image and world, and in particular a historical tie between the visible features of the image and the way things in front of the camera looked. Images are visual reproductions of parts of the visible world.

This *reproduction theory* of visual representation conceives of images as copies in the sense of reproductions, as discussed above (p. 67). Images are tied to what they represent by their unique histories. Indeed, it is a familiar suggestion that images, especially photographic ones, are a specific sort of historical ‘traces’ or ‘indices’ (see e.g. Krauss, 1985). Images are traces, because they represent what caused them to be the way they are; they are bearers of what has been called ‘natural meaning’ (Grice, 1957). They are quite special traces because in its connection with the past, an image privileges one or more items in its history—the item or items whose look is responsible for how the image looks. M.G.F Martin defends this idea, and explains how we can combine a key insight about natural meaning with an insight about a specific way a surface can look. As he writes,

Natural meaning is present wherever the object bearing the meaning is the way it is in virtue of the object meant having been the way that it was: smoke, tree rings, footprints and fossils are all bearers of natural meaning in this sense. But natural meaning (as those who have sought to use it as the basis of a theory of representation are keenly aware) is indiscriminate with respect to which objects are meant. An object is meant by something bearing meaning where there is a suitable correlation between the two. For example, in a sequence where we make a copy based on the last copy—as when we photocopy a photocopy, or photocopy the photocopy of a photocopy—then the last in the sequence will naturally mean each earlier member in the sequence. Natural meaning gives us no way of privileging any particular stage in the sequence, nor any partic-

ular predecessor of a bearer of meaning. In contrast, a natural image, such as a hologram, does privilege one entity in its history: whichever solid object's appearance is responsible for the appearance the hologram now presents. (Martin, 2012, p.342)

The reproduction view conceives of an image as an object with a specific kind of look—a look that 'belongs' to some other object, in that it is reproduced from the look of that object. The reproductive relation obtains by virtue of the image's particular causal history. This presupposes that the way one thing looks can be responsible for the look of another. In the right circumstances looks can have causal effects. As O'Shaughnessy illustrates, it was the *look* of Helen's face, rather than its chemical or electrical properties, that caused a furore in ancient Greece (O'Shaughnessy, 2000, p.570). Allegedly it was that look that caused a thousand ships to leave the port. Similarly, because some object before the camera looked the way it did a photograph may look the way it does.

A reproduction theory of visual representation contrasts starkly with a resemblance theory. On the latter, an image represents because of its intrinsic features; its chromatic and formal qualities. The reproduction theory denies this. It maintains that representation lies in features extrinsic to the image: facts about its history. In one way, this may seem a virtue for the theory. If an image's history secures what it represents, then its representational properties do not solely depend on its visual features, such as its colour and shape. The move to extrinsic characteristics liberates the reproduction theory from a constraint that was fatal for the resemblance theory. For the reproduction theory there seems no obstacle to saying that photographs, or images in general represent everything that can be seen, including events, processes, states, and the like. Although not all visibilia have a shape or colour, given that they are visible there must be some way they look. The 2003 ground invasion of Iraq

looked a certain way. But that does not require the invasion itself to be coloured or to have a shape; we may suppose that the way the ground invasion looked crucially depended on the way the various extended bodies that took part in the invasion looked. But the precise way in which complex or higher-order looks are constituted does not concern the reproduction theory. Media were able to copy the look of the first stage of Operation Iraqi Freedom by exploiting a crucial historical tie between the invasion they aimed their cameras at, and their resulting photographs and televised broadcasts. On this account, visual representations of the invasion depend historically on the looks of the various events and bodies that were part of the invasion.

However, precisely because an image does not represent by way of its intrinsic features, it is unclear if the reproduction theory leaves any representational role for the visible features of the image, apart from merely continuing a certain history. Considered from this perspective, its reliance on what seems no more than a contingent historical tie threatens to become the reproduction theory's greatest vice. Let me explain this problem in more detail.

One feature naturally means another feature where there is a suitable connection between the two. Yet we have no reason to think that such suitable connections will themselves be visible properties of the meaning bearer. Can we see the colour of the pH test as covarying with the acidity of a substance? Perhaps not. The mere fact of suitable connections that secure natural meaning is no reason to think that reproductions of an object's appearance will themselves be visible as such. And, as shown before, nothing can be an image without being visible as an image.

A painting represents what it does because of how it looks. And so, to detect such representing, we must be able to discern its representational qualities by observing the visual characteristics of its surface. Of course we may find out that it represents, or what it represents, by

being told or by reading the label. But what is distinctive of an image is that we can in principle visually detect its representational features. Whatever more we say about the nature of visual representation, images must be able to make their representing *recognisable* to the eye, to use a phrase by Charles Travis (for this reason, Travis considers images ‘allorepresentations.’ Travis, 2013, p.26).

What kind of recognisability are we talking about? An analogy may bring this out. Some connoisseurs are able to tell whether a painting is a genuine Rembrandt by hunting down clues—visible features of the work that they know (or believe) to be reliable indicators that the painting was done by Rembrandt himself. This is a viable method, but one that crucially depends on an inference from clue to attribution. On the other hand, some specialists maintain that they are able to see Rembrandt’s hand directly. Without having to rely on reliable signs, they can (allegedly) see the touch of the master in the work. It is distinctive of visual representation that its representational properties can be visually manifest, at least to the sensitive eye, in just this way.

Some have regarded photography as aiming to capture and reproduce the appearance of things. Many of us are prone to think that photographs appear the way they do due to their almost natural reproductive relation to the past. Yet it is often overlooked that, strictly, something may be a photograph of a mountain scene regardless of how the photograph itself looks. As Barbara Savedoff cautions,

this tendency to think of photographs as faithful records of appearance can come into conflict with what we see in a photograph. Things look different in photographs than they do in person. (Savedoff, 2008, p.116)

We should not forget, Savedoff here warns, that when optical processes are harassed to serve visual representation, it is not at all obvious that the appearance of the resulting image will be familiar, or even recognisable (the Hungarian photographer André Kertész is famous

for the way he exploits this fact with artistic effect). Some lenses have a focal length that strongly differs from the focal length of the average human eye. Philippe Halsman's photograph of Salvador Dalí was made with a fisheye lens, and the resulting image clearly presents us with a distorted appearance of the artist and the surrounding crowd (Fig. 7). Yet, if we believe the reproduction theory, such distortions are quite irrelevant to its status as an image. The way Halsman's photograph looks is still suitably correlated with the way Dalí looked at the time, as the image looks the way it does no doubt because Dalí looked the way he did when the photograph was taken (for a detailed discussion of fisheye photography, see Kulvicki, 2010, p.28ff).



Figure 7: Philippe Halsman, *Salvador Dalí at a book signing*, 1963.

Extreme distortions are compatible with the natural meaning relation that the reproduction theory takes to be central to visual representation. But they render that theory too permissive. Imagine the following situation. An image is projected onto a screen using a digital projector that is connected to a photosensitive receptor placed in the parking lot outside. Each of the light-emitting diodes (LED) projects a small point of coloured light onto the screen, a point that corresponds to one of the light-sensitive pixels of the receptor. In this way, some LEDs contribute red dots to the projection, some of them green ones, and so forth. Together they project a clear image that visually represents the various cars parked outside the building. Now imagine that, for some reason, the cabling leading from the receptor to the LEDs gets badly jumbled up. Each individual beam of light is preserved, but all the projected dots are now scattered across the wall. We may assume that, although their organisation seems arbitrary, there nonetheless is some complex function that would allow us to piece them together again and restore the original image. Strictly, therefore, the natural meaning of original image has been preserved. The reproduction theory of visual representation seems committed to count the scattered jumble as an image of the outside parking lot: its appearance is uniquely causal-historically dependent on the appearance of those cars in a way that preserves natural meaning. The scattered jumble of dots surely looks the way it does because the cars outside look the way they do. According to the reproduction theory, we still get to see an image. Yet it equally seems evident that the image has disintegrated and disappeared.

This is of course a dramatic way of bringing out a simple, though easily overlooked problem. The problem is already captured in Savelloff's observation. In the end, images represent to the eye—it is what we get to see that matters. For something to count as an image, we must be able to see it as an image. The idea that some visible appearances reproduce other visible appearance may be coherent, but it does not establish that all such reproductions are images.

To be sure, the reproduction theory of visual representation does not imply that visual reproductions of appearances will always be unrecognisable or will never be images. Neither does it imply that we cannot come to learn how to see certain distorted reproductions as images. Our familiarity with fisheye photography makes it not hard to see Halsman's visual representation of Dalí as such. Someone may suggest that this saves the reproduction theory altogether. If it is possible to learn to see fisheye photographs as images, then it must be possible to learn to see any visual reproduction of appearances as an image, as long as the style of distortion becomes familiar enough (a suggestion made by Schier, 1986, p.40ff). But this suggestion misses the point. First, the reproduction theory can help itself to logical possibility, but it requires psychological possibility. Compare, it surely is logically possible to store more than forty-five items in short-term memory, yet repeated experiments suggest that we can in fact not do this. Similarly, we would actually have to be able to learn to see distortions of the kind I discussed above as images, and it is just not obvious that we could ever do this. Second, the real objection to the reproduction theory is conceptual. The theory omits our seeing an image as an image from its account of visual representation. According to the theory, something can count as an image regardless of whether anyone can see it in that way. And that is unacceptable. This leaves the reproduction view dangling over a dangerous edge.

2.5 Traces and symbols

In his inaugural lecture 'On Drawing an Object', delivered at University College London in 1965, Richard Wollheim critically discusses the suggestion that visual representation is a conventional relation between an image and some object or scene. According to this *semiotic* view (as he calls it), images represent what they do because they are conventional symbols or signs. Just as any other conventional symbol or sign requires no more than a conventional association with

what it means or stands for, images require no more than a conventional association with what they represent.

The semiotic view proposes that images represent based on conventional associations. A natural way to develop this view is to attribute a syntactical structure to the image; a structure that encodes the image's semantic information. So understood, images would represent in the way many philosophers think a linguistic expression is meaningful: a basic pictorial vocabulary comprising simple shapes and colours encodes a more complex representational meaning or content (for a recent development of this idea, see Greenberg, 2012). The relation between syntax and semantics is governed by a set of conventional, arbitrary rules (see Wollheim, 1996 for a critical discussion of this approach).

On the face of it, the semiotician's claim that visual images represent by an arbitrary association between their elements and what they represent seems very different from the suggestion that they represent via a process of reproducing the visual appearances of things. Yet these superficial differences obscure a fundamental point of agreement. Both take representation to be constituted by extrinsic properties of the image, not by intrinsic ones. Both the reproduction theory and the semiotic view accept that, just by knowing all intrinsic features of an image, we do not yet know anything about what the image represents.

Because of this fundamental point of agreement, Wollheim's main objection to the semiotic view—first suggested in his inaugural lecture—could take the same form as the objection that left the reproduction theory of visual representation dangling. As Wollheim explains,

we could imagine a painting of a landscape in which, say, the colours were reversed so that every object—tree, river, rocks—was depicted in the complementary of its real colour: or we could imagine, could we not?, an even more

radical reconstruction of the scene, in which it was first fragmented into various sections and these sections were then totally rearranged without respect for the look of the landscape, according to a formula? And in both cases it seems as though there is nothing in the present view that could relieve us from classifying such pictures as representations. (Wollheim, 1974, p.25)

Wollheim's point is that if an image is no more than a conventional symbol that represents according to some arbitrary set of rules or formulas, then all sorts of things could be visual representations—including things that we could not possibly see as such. But that result is problematic. And as I noted earlier, the ability to see something as an image is integral to visual representation. Wollheim too observes how, even though technically no information was lost, we have no reason to think that the jumbled up painting retains its representative character. Just as in my earlier example of a messed up digital projection, such a total rearrangement would leave us with, in his words, “not a picture that we look at, but a puzzle that we unravel” (Wollheim, 1974, p.26).

As I showed earlier (p. 93), one cannot just insist that we could become used to the function of the scattered projection. Similarly, one cannot just insist that we could become used to the rules of this new symbol system. For as things stand, there is no compelling psychological evidence to think that we could. Any account that takes representation to be constituted by properties extrinsic to the image faces pressure in light of these observations. However, the semiotic view and the reproduction view come apart, in that the latter seems to have one last move left. It can resort to accuracy.

Recall, every reproduction has a certain degree of accuracy. This is because a reproduction is always the product of some form of mechanism or procedure designed or evolved to preserve features of an original (cf. Evans, 1982, p.125). The implementation of that proced-

ure can always be more or less successful, given its ends. As Gareth Evans writes about the photographic process:

the output is of those objects with which we have to compare it to judge the accuracy of the mechanism at the time the output was produced. (Evans, 1982, p.125)

Something is a reproduction of something else only if it is the output of such a mechanism. For that reason, reproductions always preserve the features of an original with a certain degree of accuracy. These facts about it give the proponent of the reproduction view something to work with: it is always in principle possible to assess the efficacy of the mechanism, or accuracy of the output. A defender of the reproduction theory could therefore object that I have presented their view in an unfavourable light. I have equated it with the much cruder idea that visual representations are mere causal-historical traces of what they represent. A skid mark on the road is a trace of the car that passed it too speedily. By virtue of its history, the skid mark counts as a trace of that particular car, and no other. Yet a skid mark is not as such a more or less accurate trace, because it is not the product of a mechanism that was designed or evolved to preserve the features of an original.

A reproduction theorist may have the following move in mind. Instead of thinking of images as mere visual traces from the past, we should think of them as in some way preserving something from that past (the point is emphasised in Martin, 2012). In particular, the reproduction theory may stress that images preserve the sensible appearance of the objects and scenes that suitably figured in their history with a degree of accuracy. And, so the suggestion goes, the jumbled images that seemed to pose a problem for the reproduction theory can hardly be said to preserve that appearance accurately.

I agree that accuracy is a crucial dimension of the concept of a visual reproduction. I also agree that the jumbled images I discussed hardly

preserve any appearance accurately. But why is that so? What is it about them that makes it so obvious to us that they are inaccurate? To answer this, we are in effect asked to identify the standard of accuracy that is implicit in our assessment. While in theory numerous standards could be candidates for this role, it seems clear that the standard considered here must itself be a visual one. We hesitate to count a jumbled visual reproduction of a scene's appearance as accurate, because the result does not *look* right. That is, because it is somehow visually manifest that the image itself fails to present us with the visual appearance of the object or scene it represents.

The standard of accuracy implicit in the reproduction theory of representation is a standard of visual appearance. It tells us that an image accurately reproduces a visual appearance to the extent that it presents us with the visual appearance of the object or scene it represents. But that observation brings out a tricky point. For what is this other than the idea that we began with? The idea that images present us with the visual appearances of things has returned with full force. As it stands, the reproduction theory constrained by accuracy exploits that familiar starting point without elucidating it.

From here we can see how the story goes. The idea that images present us with the visual appearances of things can be elucidated in several ways. Yet adopting any of these would transform the reproduction view's baseline that images represent purely in virtue of their causal-historical tie to an original scene.

One option would be that images present us with the visual appearances of things by copying their appearances. As I have explained, copying the appearances of things can be understood in two ways. Copying can either involve creating something which resembles an original, or it can involve creating something which is causal-historically reproduced from the original. Taking the latter understanding of 'copy' would amount to elucidating the reproduction view's implicit standard with a reiteration of the reproduction

view. Such circularity seems to me in no way useful. The only option left, in that case, would be to import aspects of the resemblance theory of representation. That way, the reproduction view could maintain that a reproduction is accurate to the extent that the image visually resembles the items of which it reproduces the appearance (roughly, I think, the way Peirce conceived of photography, see Short, 2007). Yet that route now seems hardly attractive. I have already shown that the resemblance theory is too restrictive as it stands, and can only be defended on grounds that are controversial and lack independent motivation.

Another option would be that images present us with the visual appearances of things by making the things they represent in some way visible to a viewer. To my mind, this has seemed the most natural way to go all along. (I myself will defend a related view in the chapters to come, though without a commitment to reproduction.) Note that explaining the standard of accuracy in this way would significantly transform the reproduction view. On this understanding, something would only be an accurate reproduction of some scene's visual appearance to the extent that this scene is made visible to viewers of the image. An image would represent not just by virtue of its history, but also by virtue of making the past scene, the look of which it reproduces, visible to the eye. Pressingly, this would import into the reproduction view aspects of a competitor to the family of options to which the reproduction view belongs. It would be to supplement a version of a copy view of representation with elements of a principal alternative to copy views. It would let go of the main attraction of supposing that copying appearances forms the basis of visual representation.

2.6 Images without copying

What is it for an image to present us with the visual appearance of the scene it represents? The copy theory of visual representation maintains that images copy the visual appearance of what they represent. I have shown that, regardless of whether we think of such copying in terms of visual resemblance or in terms of visual reproduction, it is not at all attractive to conceive of visual representation as a kind of copying of appearances. That conception makes a muddle of the relations between viewer, painting and the visible world. Therefore, as a theory of visual representation it can only be defended on grounds that to many will remain too controversial.

I do not think Plato himself took the copy theory seriously as a philosophical understanding of visual representation; this much is manifest in the polemical tone of his discussion. Instead, as becomes clear in his brief but insightful argument, Plato provocatively shows that if the copy view were true—*if this were all the painter could do*—it would render painting a quite trivial affair.

The problem in Plato's days was that, more and more, copying appearances was all the painters in fact did. In the *Republic* Plato expresses a critical attitude towards a specific form of naturalism or 'skiagraphic painting' that had come into fashion (cf. Keuls, 1975). Plato does not think that the project of naturalistic painting is incoherent or impossible to carry out, but he wants to highlight that naturalism in painting is perhaps not the most valuable or worthwhile enterprise. As speaks from the dialogue between Socrates and Glaucon I began with, if painters in copying appearances do not need to have any knowledge of the world—as Plato seems to have thought—then we, as viewers, cannot distil any knowledge of the world from them either. But to this extent, Plato's criticism is at the level of art theory (and perhaps epistemology), and should not be confused with a more substantive claim about what visual representation is.

When we take it to embody a philosophical claim about the nature

of visual representation, the copy theory is much more problematic than it may have seemed initially. This gives us sufficient reason to consider more seriously an alternative. What then it is for an image to present us with the visual appearance of the scene it represents?

In the next chapter I will propose and defend a novel alternative to the copy theory. That alternative accepts that pictures make things visible, and therein I myself will make a controversial claim. Given that in ordinary vision we cannot genuinely see what is not there, how could it be possible for an image to make an absent scene visible? The copy theory may seem *prima facie* attractive, because it is not obvious how an answer to the previous question could be anything other than negative. Yet, I want to show how a straightforward answer to this question is available, and is rooted in our everyday notion of perception.

The answer I will defend has received little attention. In part this seems to be because, in order to get it into view, somewhat paradoxically we need to turn our backs to vision—at least for a brief while. As I will show in what follows, by reflecting on the nature of auditory perception we can uncover a concept of perceptual representation, a concept already implicit in the way we think about our perception of the world. Images, I will conclude, are perceptual representations in *that* sense.

3

The eye's contraband

In east London's Cheapside stands St. Mary-le-Bow. Built on stone arches, the church is also known as St. Mary de Arcubus. Christopher Wren designed the present building right after the Great Fire of 1666. The old church was mostly destroyed in the blaze, but Wren managed to use some of the old foundations. Although Wren's building received architectural praise for its classic visual appearance, the older, original church of Bow made history mainly because of the way it made itself heard across the city: its bell defined the true Londoner. To be a veritable Cockney, people agreed, one had to be born within earshot of Bow Bell. And the bell's audibility reached far.

In a famous satire from 1604, *The History of Richard Whittington*, the historical character of Dick Whittington (c. 1354–1423), mayor of London, decides to leave London to find his fortune elsewhere. Early in the morning, he sets out for the North. While climbing Highgate Hill, half a dozen miles from the city, he takes a break. Sat on a stone, suddenly the sounds of Bow Bell that had started to ring in the distance struck him. Listening more carefully to the distant bell, the faint sounds that reached him appeared to be voices!

Turn again, Whittington, Lord Mayor of London.

Turn again, Whittington, Lord Mayor of London.

According to the story, the message Whittington believed Bow Bell to transmit took possession of him. He returned to the city promptly, trying his luck by running for Mayor. No doubt the alleged message was a mere product of Whittington's fancy. Yet that on Highgate Hill he could hear the bell way down in the city was real, and it conveys an important truth about audition. Just as in vision, in auditory perception we can perceive concrete bodies that are manifestly at a distance from us. Sitting on a stone in Highgate, Whittington could hear a bell several miles removed from him, and from there he heard it ringing all the way back in east London's Cheapside.

Audition is a distal sense. Our capacity to hear is a capacity to perceive things at locations that can be spatially removed from us. Striking about the distal senses is that they seem not to require contact with their objects. For this reason, reflection on them potentially paves the way for philosophical puzzles about how such perception is possible. We have already encountered how reflection on visual perception gave rise to such puzzlement (see chapter 1, p. 37). Yet, interestingly, even though it is a distal sense, hearing has not given rise to the kind of puzzlement that vision has. For some reason, how hearing a bell in the distance is possible has seemed obvious to people.

That audition did not puzzle philosophers is not just because, traditionally, authors have taken vision for the highest or most noble of the senses. It is true that for this reason, vision has received most philosophical attention, allowing many of its puzzling features to become exposed. But the true explanation lies in something much more straightforward. Audition has proven less puzzling, because it is obvious to us that the sense of hearing is acted upon by sounds. When we reflect on what it is to hear, we can discover proximal objects of hearing that reach our ears. Indeed, over and over again, philosophers characterise hearing as a passive bombardment by sounds in our environment. We conceive of the sounds that reach us as pounding our ears, as waking us up, and as acting on our sensory capacity to

hear. The very structure of auditory experience, as a form of distal perception that is mediated by sounds, is manifest (a clear and early observation is made by Plato, *Tim.* 67^b1-4).

In what follows I will show that we already conceive of auditory perception as a form of distal perception that is mediated by sounds. When a sound makes itself present to our ears, it can confer audibility on a world beyond itself. In this way sounds play a mediating role in perception. They enable us, in the right circumstances, to hear and listen to things wholly distinct from those sounds. This in turn matters for an understanding of images. It shows that we already conceive of sensory modalities as potentially having a mediated structure. This, I suggest, helps us see what images are. Just as sounds, images play a mediating role in perception. This enables us, in the right circumstances, to see things wholly distinct from the images we are presented with in sight. To defend this novel approach to images and visual representation in more detail, we need to begin by considering what is involved in hearing the world. In particular, we need to ask what sounds are.

3.1 Sounds

What is a sound? Sounds are emitted when an body is struck or moved in the right way. A wooden hammer causes the copper bell to vibrate for a period of time. Vibrations of that sort cause air to move in a way that gives rise to a sound. If the bell is large enough, and if it is struck with sufficient force, that sound can travel or propel all the way across a city. When those sounds reach the ears, people without hearing impairment may be able to hear them. Unlike light, which itself only becomes visible to us in exceptional circumstances, sounds are heard in any conscious act of hearing. This makes the possibility of hearing quite unproblematic, akin to the way the possibility of touching is obvious. It is manifest how distant objects can have an

effect on us: by producing audible sounds.

The metaphysics of sounds is controversial. Most philosophers agree that sounds are public phenomena that do not depend on any particular mind or act of perception. Some think sounds simply are waves. Sound waves are produced when the distribution of molecules of a surrounding medium, typically air, is disrupted by some vibrating body in that medium. The disruption spreads out, and propels itself in many directions, somewhat like a ripple spreading out across the face of a pond (O'Shaughnessy, 2000, p.445). But even if sounds spread out in this way, this does not itself imply that they are identical to waves. The wave view is controversial. Some authors suggest instead that sounds are audible individuals distinct from those waves. For instance, Casey O'Callaghan maintains that a sound is an event that a vibrating body undergoes (O'Callaghan, 2007), while Roger Scruton argues that a sound is a 'secondary object,' a distinct kind of *audibile* that, qua individual, parallels the way colour qualities are the irreducibly visible cloaks of things (Scruton, 2010; see also O'Shaughnessy, 1972).

The dispute about the metaphysics of sounds need not be settled here. What matters is that sounds are wholly distinct from the material bodies in our environment, in a way that the visual and tactual features of material reality are not distinct from what we see or touch (cf. O'Callaghan, 2008a, p.804). All the above positions can accept this. That sounds are distinct is already obvious in experience: in reflecting on our experience of hearing a bell in the distance, one can confirm that sounds are wholly distinct from the material bodies they make heard. They are objects of audition—*audibilia*—by means of which we can hear the material world. The distinctness of sounds can also be given a more precise formulation, so let me now turn to that.

That sounds are *audibilia* means that they themselves can be heard. Indeed, sounds need to be heard in order for someone to hear anything at all. People with tinnitus merely *seem* to hear a high-pitched

tone, precisely because they do not hear any sound. We can only hear the police car in the street to the extent that some sound that reaches our ears is itself perceived. Whenever we hear concrete bodies such as a car, we always hear sounds as well. In a sense we always hear 'doubly' when we hear bodies at a distance.

That sounds are wholly distinct from the bodies they make heard means that they are neither identical to those bodies, nor have any parts in common with them, nor are properties of those bodies (see van Inwagen, 1994, p.214 for discussion of Armstrong's slightly weaker notion of wholly distinctness). That sounds are not identical to material bodies is clear. The bell is made of copper, but its sound, though possibly a coppery one, is not made of any metal at all. Neither does the sound of the bell have any part in common with the bell or with any other body. This follows from the fact that all of a body's parts are spatially extended, while sounds lack spatially extended parts. Even though a sound can be in a space, it does not take up space. A room may become noisy if it is filled with sound, it does not in this way get crammed. Lastly, that sounds are not properties of bodies equally stands. The sound of the bell can cease or be muted without the bell undergoing any change, for instance when the surrounding medium is removed by placing the bell inside an air pump. If sounds were properties of bodies, a change to the sound would be a change to the body.

Some have resisted the manifest character of hearing. They deny that sounds are wholly distinct from bodies, appearances notwithstanding. But they generally do so because they are gripped by an argument. They think that we can only perceive something if we perceive at least some part or property of it. Perceiving something wholly distinct from it, they think, is not going to help us perceive any distant body. Now, surely we can hear distant bodies, such as the ringing bells of St. Mary-le-Bow. This has led these authors to conclude that sounds must be parts of the concrete individuals we hear, or proper-

ties of them. How else could we hear them? (for a view motivated by this worry, see Pasnau, 1999). The problem with this line of reasoning is the contentious assumption it relies on: the assumption that we can only perceive a thing by perceiving some part or property of it. This is simply without foundation. I suspect that what drives this thought is a dogmatic conviction that perception is *uniform*. That is, that processes of perceiving must always behave in the same way. True, in paradigm cases of vision and touch we perceive objects because we perceive their properties or parts. You cannot touch a table without being in physical contact with at least one of its sides. Yet, why should we expect all perception to conform to this principle? Hearing distant bodies manifestly relies on hearing something that reaches us, while seeing and touching manifestly put us in contact with parts or properties of bodies (we manifestly reach *them*, so to say). Already these manifest differences are sufficient to resist fitting all forms of perception in a uniform mould (cf. Batty, 2010, p.513).

Sounds are objects of audition by means of which we come to hear a distal, material world. In the next section I will discuss this specific function of sounds in perception, and bring out its relation to perceptual activity.

3.2 Beyond the world of sound

Discussing the different modalities of perception, J.O. Urmson writes that

Clearly one may see, hear, feel, smell, and taste physical objects like motor cars and apples; it requires ingenious stage setting to make 'I hear an apple' or 'I taste a motor car' sound natural ... but one may certainly hear a motor car or taste an apple (Urmson, 1968, p.117).

I think Urmson is right. Just as we can see a magpie, or taste vinegar in our salad, we can hear things such as people in the corridor, aeroplanes overhead, and a pub fight down the street. Perceiving mundane things like these is what auditory perception is for (cf. Nudds, 2014). The audibility of material bodies or events depends in part on what happens to or in them. Some material bodies do not by themselves produce sounds in the way a waterfall or motor car can. Take an apple. For it to be heard it needs to partake in an event during which something happens to it; someone may throw the apple against your front door, for instance—Urmson’s ingenious stage setting. By being thrown against a wood surface, apples, together with a host of other things—such as the event of the apple’s being thrown, the hollow space inside the door, and perhaps the antique brass door knocker—can become audible. Urmson’s point is that, allowing ingenuity, more or less anything in our surroundings can become audible to us.

We can hear mundane objects and events. This suggests a crucial duality in audition: the sense of hearing can take both sounds and everyday bodies and events as its objects. The observation that we can hear bodies moreover rules out the restrictive view, often attributed to Berkeley’s mouthpiece Philonous in the *Three Dialogues*, that sounds are all we ever hear (Berkeley, 1954). At least in the world in which we evolved, our sense of hearing goes beyond that.

What is the relation between sounds and the other things we hear? I accept that sounds are the primary objects of awareness in audition. By this I mean that whenever we hear something, we always hear some sound. But in being heard, sounds can fulfil a crucial role in enabling us to hear other things. It is by means of sounds that we can come to hear other things. Consider the following case. When you hear an apple smash into your front door, your perception manifests a specific structure: you hear the sounds *immediately*, and you hear the fruit *mediately*. You hear the apple in virtue of a primary object

of hearing, namely the sound of its impact on your door (see Jackson, 1977, p.15ff for a seminal discussion of 'in virtue of'). The structure is one in which the sound mediates your perception of the apple. You hear the sound immediately, and you hear the apple mediately, by means of the sound. I suggest that only this structure helps explain how we can hear anything beyond sound. We cannot hear concrete bodies or physical events immediately. Sounds are the primary objects of auditory awareness, and so every auditory perception is a perception of some sound. But we can hear more than that, by means of the sounds we hear.

I want to discuss the conclusion that sounds mediate the hearing of bodies and events in more detail. How is the mediate structure of audition manifest in experience? There may not be a single answer to this question. But one feature undeniably stands out. This is how hearing sounds enables us to single out and listen to material bodies, events, and their attributes.

Imagine that at night you are woken up by a series of noises coming from your living room. Reluctant to get out of bed, you first listen more carefully to what goes on. A high pitched squeak. A thud. Did you leave a window open? Is someone down there? Suddenly you hear a set of cushioned feet sliding over the wooden floor. You realise it's the cat. The animal must have managed to open the kitchen door, again. What is it doing? You listen more carefully. Your ears are now no longer focused on the sounds themselves, as they were initially, but instead you're listening to what is going on in the living room downstairs. On the basis of the various sounds that reach you, such as a rustle of newspapers, the clattering of some pens falling on the floor, and the sound of nails on a glass surface, you are able quite reliably to track the animal through the familiar space downstairs. In all of this, you listen in a way that is utterly dependent on the sounds that reach you. Still, importantly, the objects of your perceptual activity—hence, the objects of hearing—extend well beyond those sounds.

Recently a number of authors have argued, both on scientific and common sense grounds, that in audition we can successfully single out, attend to, and perceptually track things that are not sounds (O'Callaghan & Nudds, 2009; Matthen, 2010; Cottrell & Campbell, 2014). If this is right, then it not only underscores that audition can be of vastly more things than sounds alone. It also brings out a central role for perceptual activity in helping us to study and clarify the structure of auditory sense perception.

In general, successfully attending to an item in perception (be it in the form of active touching, looking, or indeed listening) can be taken as a criterion for a person's perception of the item (see Crowther, 2009, p.23ff for a programmatic inventory of such directed sensory activities). That someone listens to a cat entails that they, at least for some of that time, actually hear that cat. No doubt the theoretical possibility of illusion or hallucination complicates how we can apply this criterion for perceiving. It may be that someone merely thinks they are listening to something, while in fact they are not—think of a sudden onset of tinnitus which someone mistakes for a defect in the stereo set. That there is room for error in applying the criterion does not make the criterion itself any less adequate.

Experimental psychologists have in recent years started to study the place of attention in auditory experience in more detail (for an explicit defence of the present point based on experimental results, see Nudds, 2014). But note that this structure as such is most familiar to many of us whenever we listen to things in everyday situations. We are happy to count the case described above as one of 'hearing a cat', at least in part because it is manifest in experience that this is what we attend to and succeeded in perceiving.

Our capacity to hear enables us to engage in perceptual activities directed at the material world (cf. Gaver, 1993). Significantly, our capacity to hear requires perception of some intermediary object—a sound—as a precondition for the exercise of such activities. It is

only by virtue of hearing a sound that you can direct your sensory attention to some object or event. Auditory perception of the material world manifests a structure foreign to what we find, say, in typical visual encounters with that world. Normal audition possesses a specific 'mediated' structure that makes auditory perception of a distal environment possible. This mediated structure, instead of posing a problem for the possibility of hearing distal things, precisely enables it. The sounds we hear enable us to listen to bodies and events wholly distinct from them. They confer audibility on a world beyond sound. Sounds function as *perceptual mediators*, as I will call them. Something functions as a perceptual mediator if perceiving it enables us to attend perceptually to some thing wholly distinct from it.

Anyone able to hear can hear sounds: all that needs to happen is that a sound affects them strongly enough; a sound can even wake us from our sleep. In some cases, hearing distal bodies by means of sounds involves some further effort. Someone just woken up by a strange noise may not yet be able to detect the lorry that it makes audible. They may not immediately be able to hear it as the sound of that object. Coming to hear the sound in that way will typically involve listening. In such a situation, coming to hear the lorry requires us actively (though not necessarily deliberately) to direct our sensory attention. Hearing the lorry can be a sensory achievement, in a way that merely hearing its sound is not. Through such sensory activities we can complete the natural function of the sense of hearing—not unlike the way lacing up a new shoe allows it first to function properly (I borrow the image from O'Shaughnessy, 1972). When humans in this way exploit the sounds that reach their ears, their hearing goes beyond the basic receptivity of the auditory sense.

Frequently the way we listen out for things is guided by implicit questions or search tasks, as we could put it. Auditory perception of our environment is a clear case in which gaining information can make you more sensitive to the various audible particulars in your

environment. Because you already know roughly what kind of sound you are hearing, you know what to listen out for. By knowing what you are listening to once you have achieved that state, you will be able to attend more determinately to what you hear, enabling you to hear more and discover more specific things. In this way, hearing things beyond the sounds that reach your ear depends to some extent on knowledge of the world, and this again brings out the prominent role of perceptual activity.

In all this, I am talking about perceiving particulars. Hearing a cat is not just a kind of epistemic (or ‘propositional’) hearing. That is, we can hear more than just *that* a sound is of a cat, or *that* the cat is in the room (the relation between perceptual activities and epistemic perception is the main concern of Crowther, 2009). No doubt such epistemic hearing does occur; you can indeed come to learn new facts through hearing. Yet my argument here identifies a form of perception that seems more fundamental. We can only hear that such-and-such is the case, or hear what something is doing, by hearing some particular thing or other. Indeed, it is natural to think that hearing that the cat broke the vase is rooted in an experience of hearing particulars that are not themselves sounds: you hear the cat, you hear its running through the room, its jumping on the table, and you hear the vase when it lands and shatters into fragments on the wooden floor.

Hearing distal bodies at least typically involves some form of perceptual activity. However, there is no reason to infer that this reflects a universal principle of distal perception. Hearing distal bodies tends to depend on perceptual activity in a way that vision does not. In vision, many philosophers think, merely being in a state of seeing suffices for vision of a distal, material world. “With sight,” Hans Jonas writes, “all I have to do is open my eyes, and the world is there, as it was all the time” (Jonas, 1954, p.512). Suppose that each sense has an associated ‘field’ such that, as soon as items happen to be within it, they are perceived. On that presupposition, Jonas’ brings out that

the items within your field of vision can be concrete, distal bodies. Suppose that in your bedroom at night someone suddenly turns on a light bulb. The bulb is something that can fall in your field of vision. It seems that, as soon as the receptivity of your eyes is functioning properly, you cannot but see that bulb. Getting that bulb into view requires no active direction of your gaze; it is hardly a sensory achievement. The sense of sight *per se* is receptive to the distal, material world.

Cases of hearing concrete bodies and seeing them differ subtly but importantly. The sole occupants of the 'sensory field' of hearing are sounds. This is why hearing more than just sound typically involves at least some active engagement with the things that lie beyond it. Ordinarily this kind of engagement is largely unreflective, but it stands out clearly when we encounter unfamiliar sounds. An unfamiliar sound at first may seem detached from the rest of our perceived surroundings, and forces us to listen more carefully. Placing the sound requires a quick auditory search, in which we tune our ears to the object or event from which the sound springs (see Crowther, 2010 for a discussion of how these perceptual activities may affect how it is to have such episodes). In this way, active listening plays its part in a broader pattern of searching and mining our surroundings for perceptual encounters with objects and features that would have gone unperceived otherwise (cf. Johnston, 2006).

As I have explicated here, auditory perception can be of more than just sounds. More specifically, sounds enable us to listen to things wholly distinct from them, and so make mundane objects and events audible. This mediating role of sounds is manifest in experience. Making something else perceptible is a defining characteristic of perceptual mediators. Perceptual mediators confer perceptibility on objects wholly distinct from them, and so enable us to perceive those things despite the fact that there are wholly distinct from them. My discussion of auditory perception has established that the notion of perceptual mediation is in principle viable, and finds

actual realisation in the sensory modality of hearing. Uncovering the perceptual mediators our environment has on offer is of central importance to the philosophy of perception. I will argue that it is also crucial for understanding visual representation. As I will show in the next section, the mediating role of sounds can give a representational dimension to perception. If that is right, then it opens the way for a powerful account of representation by means of images.

3.3 Representation in perception

Listening to material bodies like coaches and cats, or to physical events such as bar fights, does not depend on these objects of perception being actually present in our environment. Instead, when they are audible, they are audible solely via sounds. Their sounds, when they reach our ears, can make these objects and events audible to us. That auditory perception does not depend on the presence of some of its potential objects may not stand out clearly in everyday experience. At least typically, what we hear is located somewhere in our surroundings and likely to be perceived via other sensory modalities too, vision in particular (the relative dominance of vision is illustrated nicely in the ‘ventrilloquist effect’, see Bertelson, 1999). It is very common for people to see an object while they hear it. But it is an important consequence of the mediated structure of auditory perception that hearing objects strictly requires only the presence of sound.

Examples can bring this out well. The clearest of these have only become available after Édouard-Léon Scott de Martinville, and later Thomas Edison, conducted their pioneering work on technologies of sound reproduction. The late nineteenth century invention of phonograph cylinders enabled sounds produced at a specific moment in history to be captured, stored, and played back at a later time (cf. Sterne, 2003). A contemporary playback of Scott de Martinville’s rendition

of the French song “Au Clair de la Lune”, recorded on 9 April 1860, offers a clear case where you hear some thing or someone not present to the senses (some philosophers have observed this point, see e.g. Gregory, 2013, p.17).

The key point I want to bring out is this. Hearing's independence from at least some of its objects is not exceptional, but part of the very structure of mediated perception. I think this points to a further step. Because mediated perception has this structure, it is most adequate to think of cases of hearing material bodies as *representational*.

To find out what had happened in the Boeing 737 that crashed near Pittsburgh on 8 September 1994, a team of audio forensic investigators played back the cockpit voice recorder that was retrieved after the accident. The team was particularly interested in a malfunctioning of the plane's rudders. Hence, it is unlikely that they listened merely to the sounds reproduced by their playback equipment. Indeed, they found out that the rudder's eventual jamming was not due to inept operation of controls, but occurred because of mechanical obstruction. Their observation required a careful, well-informed ‘mining’ of an auditory scene that was made audible again by replaying the recording. This activity of listening crucially relies on an ability to attend to specific objects and properties that are made audible through a recording. The investigators' finding was made possible because they were able to single out specific items audible via the recorded sound: the plane's rudder, as it locked in its extreme position, and the members of the flight crew as they manipulated the controls.

The malfunctioning aeroplane parts and crew members were audible months after the craft had crashed. They were audible because of what had made them audible that day in 1994: their sounds. Listening to those instruments was not in any way an illusory or non-veridical experience. To the contrary, the judgements based on it turned out to be spot on. And although the forensic team exercised their capacity to hear things by means of those sounds with expert precision, they

did not engage in anything beyond the pale of ordinary auditory discrimination. This example brings out dramatically a rather mundane point: the success conditions for mediated listening do not include an object's presence.

Sound reproduction techniques allow us to recreate the conditions sufficient for hearing concrete bodies and physical events. Precisely because of this, also in listening to a recording we can focus our perceptual attention on something that goes beyond the heard sound itself. This is especially easy if the sounds are of a kind our ears are already attuned to. This everyday ability was exploited, both scientifically and commercially, as soon as the technology to reproduce sounds was invented. As Jonathan Sterne describes, very early on the genre of 'descriptive specialities' came into fashion—a precursor of the still familiar radio play. In those recordings of sounds, distributed for amusement, events were re-enacted for a recording device using a variety of techniques and tricks to increase a naturalistic listening experience. Descriptive specialities allowed listeners to have the experience of hearing events they perhaps only ever read about in encyclopedias, newspapers, or novels (cf. Sterne, 2003, p.245).

Why would such perception be representational, as I suggested earlier? This is because cases in which we can perceive things in their absence have typically been understood as instances of representation. The concept of representation has taken centre stage in recent work in the philosophy of perception. The representational dimension of perception reveals itself in cases where, allegedly, things sensorily appear to a perceiver that are not present to the senses, or where we have reason to think that the perceptual state a perceiver is in does not depend on such presence (Rowlands, 2015, p.6; Siegel, 2010, chap.6). Wherever a kind of perception does not require some or all of its objects to be present, we may take it to involve a form of representation of those objects. Hence we can accept the following as a sufficient condition for perceptual representation: something is

represented in someone's perception if they perceive something that is not present to the senses, or where their perceptual state does not depend on such presence.

Now of course this raises a difficult question. What is it for something to be present to the senses? The question is difficult, because it is at least natural to think that as soon as something is perceived, it is thereby present to us, at least in some way. Was their hearing the rudder not sufficient ground for thinking that, even after the plane had crashed, the rudder was present to the audio forensic investigators? (Perhaps in this way a distant star may be present to the eye, even if it is very far away.) The point is well taken. Just as there is undeniably a sense in which that rudder was not present to them, there is an intuitive sense in which, by means of those sounds, it did become present. This tension is no more than a consequence of the somewhat see-saw way philosophers have conceived of perceptual representation. It seems to have worked its way into the very catchphrase they have come to use for representation: *presence in absence*.

What matters is that some objects are able to determine in part the way the world appears to us in perceptual experience, whereas other objects are not. Perceiving something entails having a kind of sensory experience. In such a sensory experience, things appear to one, and they appear in a certain way. When you stare at a green wall in front of you, it is the greenness of the wall that in part determines the way the wall appears to you. Typically, the wall looks green to you because it is in fact green. An object or scene may be understood to be perceptually present to you if its properties play a role in determining how things appear to you. As Mark Johnston has characterised this understanding of presence, it is the way in which some item—object, quality and so on—discloses some aspect of its nature (Johnston, 2007, p.233). Whereas the green wall and even the distant star disclose some aspects of their nature in this way, the rudder the team heard was unable to do this; its sounds did that *for* it. In this way the

rudder the team was able to hear was not present to the senses, but represented to the team by its sound.

The mediated structure of auditory perception allows for sounds to be conveyed, reproduced, or more generally to recur in ways that make some things audible, or audible anew. Sounds can make these things heard, while figuring importantly independently of those things they make heard. This is why we may conceive of the sounds in our environment as potential auditory representations of a world that lies beyond them. Those heard sounds can play a determining role in the exercise of our capacity to listen to things that are not sounds (that they play this perceptual role is not infrequently overlooked, see e.g. Tolliver, 2007, p.143).

Just as with other perceptual capacities, we can naturally conceive of the capacity for representational perception as only exercised in response to something. We may think of them as reactive capacities along the lines of Aristotle's theory of perception (I discussed this in chapter 1, p. 45). Sound figures as an essential cog in the machinery of representation in auditory perception. In general, what we can hear is at least in part determined by the audible character of the sounds we perceive. Moreover, precisely because sounds are wholly distinct from the concrete bodies they allow us to perceive, those bodies themselves simply lack an audible character—hence material bodies and most physical events are not audible as such (cf. Jonas, 1954, p.514).

Of course, technologies of sound reproduction have only become available relatively late in the history of hearing. But this by no means implies that the possibilities they reveal are not integral to what it is to hear concrete bodies and physical events. The sense of hearing enables us to perceive things not present to it. This means that perceptual representation finds a natural home in this sensory modality.

It is not hard to see a natural tie between perceptual mediation and representation, so understood. That hearing is representational

in the way I discussed is due to the structure of the perceptual situation of hearing an object. Note, I do not claim that this is the only way of thinking about representation. Hidé Ishiguro has brought out well that philosophers have considered quite distinct phenomena to be in some way representational (Ishiguro, 1994). More recently, Mark Rowlands expresses his despair at the unhelpful equivocality of the concept of representation in philosophy (Rowlands, 2015). Yet, I think that reflection on auditory perception helps us clarify an interesting form of perceptual representation; a form that both Ishiguro and Rowlands consider to be central.

Let me emphasise that the understanding of representation I propose here contrasts with 'representationalism' or 'intentionalism' about perception, as recent philosophers have used those terms. By calling mediate perception 'representational', I am not suggesting that mediate perception does not at all depend on actual objects in our environment. I do not assume that such perception involves merely a 'construction' or 'mental model' or 'representational content' of our environment. Certain philosophers do think this. They think that perception as such does not depend on actual objects in our environment. As Tim Crane explains that view,

the intentionalist theory of perception denies that the essential phenomenal character of a perceptual experience is essentially determined (wholly or partly) by the real objects which are perceived. This is because it holds that perception and hallucination are states with the same phenomenal character and therefore of the same mental kind; but in the case of hallucination there is no real object being perceived. (Crane, 2011)

When I suggest that your hearing the lorry drive by is representational, I do not deny that the character of your experience of hearing the lorry is determined (wholly or partly) by actual audible particulars

in your environment. Indeed, I take it to be manifest that the sounds we hear do at least in part determine that character (cf. Nudds, 2014). Those sounds are there in our environment, I presume, and I find it at least natural to assume that our auditory perception depends on their presence. (Strictly, however, the model of representation as perceptual mediation I propose remains neutral on the larger question whether perception or sensory experience as such is ‘intentional’ in the way Crane describes.)

By reflecting on auditory perception, I have uncovered the possibility of perceptual mediation. I have now also shown how perceptual mediation opens up the possibility of representation in perception—perception of the absent. If something functions as a perceptual mediator, then it is a vehicle for perceptual representation. Now I am able to turn to visual perception, and explicate possibilities that seemed unavailable earlier. In the next section I will show how visual representation by means of images is no more than the visual occurrence of perceptual mediation.

3.4 Images

In 1792 the French Academy of Sciences employed the visual artist Pierre-Joseph Redouté to produce a series of botanical watercolours. Praised for their accuracy, Redouté’s meticulous renderings of blooms were expressly designed for scientific instruction and education (cf. Mallery & Mallery, 1986). Many of the plants and flowers he painted were hard to conserve, grew abroad, or blossomed only for a very brief period. In part by co-operating with learned communities, Redouté produced botanical images that still successfully function as substitutes for real-life study and botanical observation (see Kusakawa, 2012, chap.7–8 for a discussion of similar examples of early collaborations between modern science and art).

Take for instance the image he painted of a *Myosotis Scorpioides*, a

small flower better known as the Forget-me-not (Fig. 8). It is natural to conceive of such a drawing, or a good reproduction of it, as allowing us visually to scrutinise the flower's anatomy, and look at the form and arrangement of its distinctive blueish petals.



Figure 8: Pierre-Joseph Redouté, *Le ne m'oubliez pas* (Forget-me-not). In: *Choix des Plus Belles Fleurs*, 1827-1833. Paris

That we naturally conceive of Redouté's watercolours as fostering visual scrutiny of plants and their parts suggests we think of these im-

ages as making these plants visible to us. No doubt this understanding is to a large extent rooted in the fact that we find ourselves able to engage with those represented plants in a visual way. Presented with Redouté's work, we seem able to look at features of these plants; trace their shapes with our eyes; visually compare the relative proportions and colours of their stems and flowers. No doubt that students at the time could see that it was a mere drawing they were presented with. If they so liked, they could also specifically look at and appreciate these drawings for what they were, perhaps coming to admire Redouté's technique in creating them. Yet this does not take away that Redouté's plate seems to afford an opportunity for perceiving a Forget-me-not.

As I noted before, philosophers have found these data puzzling, because they conceive of sight as obviously dependent on the presence of its objects. In chapter 1 I argued that the otherwise attractive assumption that Redouté's prints make those flowers visible becomes philosophically suspect because of specific assumptions about vision. Although it may seem natural to think that his watercolours enable us to look at flowers, philosophers moved by a specific theory of vision have simply denied that visual representation of *that* sort is possible.

We can now see how this reaction is too quick. Although not presupposed in our ordinary vision of the material world, the possibility of perceptual mediation is genuine. This genuine possibility explains how Redouté's watercolours can succeed in making absent flowers visible to the observant student. They can make those flowers visible, because those watercolours can play a mediating role in visual perception.

Is it surprising that this explanation of how images can make absent scenes visible has typically been overlooked in recent debates about images? Not really. This is because the possibility and reality of perceptual mediation is not manifest to us when we merely reflect on what it is to see. We naturally conceive of vision as affording us

awareness of the material world around us. Rarely do we consider this perception to require awareness of some further item that intervenes between us and the objects we see. As long as we take vision to be the model of perception, it may be tempting to think of perception simply as an immediate mode of awareness of the world around us (for the observation that for too long philosophy has had a visuo-centric bias, see O'Callaghan, 2008b, p.316). However, reflection on perception in other sensory modalities, audition in particular, helps uncover the possibility of perception that is not reducible to an immediate awareness of the world. The reality of perceptual mediation becomes manifest to us when we reflect on how we hear and listen to the world around us. The very possibility of perceptual mediation is crucial, and our sensory lives would be impoverished without it.

Redouté's images can function as perceptual mediators in vision. This is why budding botanists can, in the comfort of their studies, week in week out, scrutinise the rare flowers he painted, peer at their many-coloured blooms, and admire their frail, delicate anatomy. Recall, an object functions as a perceptual mediator if perceiving it enables us perceptually to attend to some thing wholly distinct from it. The flower whose petals we admire is wholly distinct from Redouté's painted image, in that it is neither identical to the image, nor has any parts in common with it, nor is it a property of the image. There simply is no flower in front of us when we are just presented with this image. Nor is part of the image a part of a flower, as if some petals were glued to a page. Nonetheless, by virtue of seeing Redouté's watercolour, we can look at, and so perceptually to attend to, a flower. The image functions as a perceptual mediator, and as such makes the flowers that Redouté painted visible. Especially in a cold, Parisian winter, Redouté's prints made visible what otherwise would have remained invisible to many.

It can be tempting to use observed similarities between phenomena to infer categorical conclusions about those phenomena. If I am

right, then images and sounds figure in perception in a structurally similar way. Yet it would be mistaken to conclude that images function as mediators in vision because of their similarity to sounds. Arguments from analogy are at best suggestive, and inferences based on them cannot go beyond probability without fallacy (see Walton, 2014 for a recent review of variations of this form of argument). Thomas Reid discusses how noticing a strong resemblance between earth and other celestial bodies' orbital or atmospheric conditions leads people to speculate about extraterrestrial life (Reid, 1895, p.24). Space agencies still owe part of their funding to this tendency. Yet Reid rightly points out that such inductive analogies can at best establish likelihood, never categorical claims.

I do not suggest that it is probable that Redouté's images function in vision as perceptual mediators. My claim is a categorical one. It is based on a more general understanding of what is possible in perception. It is a conclusion drawn in part on conceptual grounds. Reflection on what it is to listen to material bodies provided a well-defined criterion for the mediating function I attribute to both sounds and images. An object functions as a perceptual mediator if perceiving it enables us to attend perceptually to some object or event wholly distinct from it. Given how successful Redouté's watercolours are as substitutes for real life study and observation, to my mind it is clear that this criterion applies to images.

We have good grounds to assume that the mediating function of images is perfectly general, and is not restricted to scientific contexts, or contexts in which observation serves an epistemic purpose. Indeed, being an image *just is* being a perceptual mediator in vision, regardless of the specific scientific, communicative or aesthetic practices in which we use it (I will come back to this distinction in chapter 5). Consider a very different example, Paul Cézanne's unfinished masterpiece *Les Grandes Baigneuses* (Fig. 9).

Cézanne's painting may strike you as no more than a jumble of

painted fields, especially if you merely see it from the corner of your eye. Such a superficial view on the painting is likely to be short-lived. The stark contours of legs and arms and buttocks almost hurl the gist of the scene at us: a chaotic group of nudes. So far so good. But how many? Where? What are they doing? It is as if Cézanne mixed such questions with his oils. Just as the eye distinguishes tones and hues, it is naturally set to search for their answers. Cézanne often plays on how visual individuation is an active process. The bathers are visible, but ill-defined. Discerning individual bodies requires us to retrace their boundaries—an ocular analysis. On the other hand, their surroundings require the different treatment of successive scanning. We need to built up a comprehensive view of the landscape, and here the group's dynamic placement is a natural aid. As Adrian Stokes writes of the group, “[t]hey absorb, and in absorbing rule, monumental, the environment” (1972). The human form points us to features of the landscape. In so opening ourselves to their cues, our sensibility for human gazes is of particular importance. Self-consciously following Titian's example, Cézanne exploits the play of eyes of his figures. It is a powerful device, as it draws on our sensitivity to visible social clues. We cannot but follow the gazes of the seated figures, our eyes stumbling upon an indistinct shape (a goose? a dog?) which they stroke from all sides. Yet, at the same time, the social dynamic interrupts our scanning of environment; suddenly our gaze interlocks with that of the right figure leaning against the tree. Someone is looking at us too.

Here we have a purely aesthetic example, in contrast with the partly scientific one discussed above. Our eyes move across the panel, but this movement is guided not by the panel, but by the scene it represents (a fact brought out well in eye-tracking studies of how we look at images; see Yarbus, 1967; DeAngelus & Pelz, 2009). Our eyes are drawn in by the painting, looking from figure to figure. We visually trace the contours and outlines of a group in a landscape. Our exploration can seek for a narrative, or it can be guided merely by formal



Figure 9: Paul Cézanne, *Les Grandes Baigneuses* (The Bathers), 1898–1905. Philadelphia Museum of Art

features, or an urge to identify, or all of these. What matters is that in every instance our exploration serves a crucial perceptual function. In our activity we uncover the represented scene. It is by looking at the painting that we get the figures that Cézanne set out to represent in view. In general, it is only because we are able to look at some scene wholly distinct from the image that the scene it represents becomes visible to us.

Here perceptual activity plays a central role in determining the objects of sight. The centrality of our activity of looking in our engagement with images more generally is a faint but unmistakable echo of the theories of vision I discussed in chapter 1 (p. 43). However, I do not assume vision is exclusively or even primarily active, as Kilwardby and Olivi did (for contemporary approximations of the idea, see Noë, 2004; Hutto & Myin, 2013). Few today would deny that the material world can affect our perception and thinking. We naturally conceive of sight as a form of being receptive to a material world. To see the world around us, we think, it just needs to be there. When it

comes to what one can see immediately, our perception does not constitutively depend on our sensory activity. My point is that things are different when it comes to perceptual representation—in this case, seeing things by means of images. I suggest that as soon as we perceive things by means of images, the story at least typically reverses: we can get something into view when we direct our perceptual attention. Our activity is able to carry our gaze beyond the image. Here lies the grain of truth in the Augustinian theories of vision. Sometimes things come into view only because of how we look at them. An image makes some further item visible, precisely because it offers us a way successfully to engage in perceptual activities directed at some further scene or item. This is a defining characteristic of perceptual mediators. Images can mediate vision, in that through our engagement with them they can make objects or scenes wholly distinct from them appear to sight.

That images invite and typically require this perceptual activity has been understood to be central to representation in the arts. Adrian Stokes emphasises how representational painting compels us actively to explore a represented scene in perception, ‘feeling’ our way around it. As he writes,

More clearly in naturalistic painting, the first test of its merit is the degree to which we become attached to the turn of the contours, the degree to which we are compelled to feel our way into spaces, whether populated or whether empty of shapes. This matter is at the heart of painting on a flat surface, distinguishing its appreciation from an apprehension of landscape itself which the eye constructs and contemplates without ado as a three-dimensional datum. (Stokes, 1965, p.27)

Stokes uses the tactile metaphor of feeling our way into the space of a picture to capture our visual activity of exploring a scene that is

merely represented by the image. In looking at the occupants of a represented scene, an image may require us actively to explore, uncover and reach for details not immediately given to us (the point is also emphasised in Collingwood, 1938, pp.146–7). In the passage, Stokes contrasts the way our visual activity brings a represented scene into view with the way we can see a landscape immediately when standing outside and looking at it. Of course we also attend to various details and actively scan when we look out on a natural scene. Yet when such a scene is immediately present to us, to see it does not require anything more of us than merely having our eyes open.

Often our directed perceptual activity contributes to what we perceive. Yet while in immediate perception such activity is optional, merely a useful guide, in the case of perceiving images it tends to become ineliminable. It enables us to move beyond what meets the eye, getting a represented scene into view. Cézanne's painting brings out the centrality of such activity even better than Redouté's relatively timid imagery. It makes an illustrative counterpart of the earlier example of tracking the destructive path of a cat in your living room. In the case of images, the explanation for why we can see a flower or a nude when we see an image most often is that we are able successfully to look at a flower, or at such a nude body, when we are presented with an image that represents it. What we manage to look at here would determine what we perceive.

3.5 The problem of visual opacity

The situation of hearing some distant body, such as hearing the bell of Bow Church, intrinsically comes with a repeatable element we could try to capture, recreate or imitate. As I explained above, whenever we hear a distant body or concrete event, we do so through some sound that reaches our ears. Sounds function as mediators. Technological advances have made sounds capturable, storable, and thereby repro-

ducible. When we play back a sound recording, we make audible a sound that was audible on an earlier occasion (Martin, 2012). Hearing that sound played back, we are put in a position to listen to objects or events from a situation in a possibly distant past. I showed that given this possibility, hearing things by means of sounds is best conceived of as representational.

So far I have compared audition and vision. But audition is not the only sensory modality that manifestly relies on perceptual mediators. Arguably, olfactory perception of objects does so too. We can only smell something if we smell a scent. Just as the sounds that reach our ears, it seems natural to conceive of the scents that enter our nose as enabling us to perceive material particulars. In recent work Clare Batty has argued that in olfaction we indeed exploit the presence of scents to perceive distant material particulars by means of them (Batty, 2015). If Batty is right, then also scents function as perceptual mediators. Can they be captured, recreated or imitated like sounds can? Currently we lack much of the technology to record mechanically particular scents in the way our iPhones can record sound. Still, like Grenouille in Süskind's *Das Parfum*, a skilled perfumer may succeed in recreating or inventing scents, and so perhaps enable us to smell whatever objects they wanted us to smell.

By contrast, sight does not seem to require perceptual mediators of the kind we find in audition and (perhaps) in olfactory experience. In reflecting on what it is to see, it is manifest that we do not normally perceive the distal world by means of some perceived intermediary that mediates between perceiver and distal world. No specific sort of object is necessarily seen in our visual perception of the world, something we must see before we can see anything else. In his discussion of the contrast between the visible and audible world, M.G.F. Martin observes that

we do not conceive of the visible world as offering us primary objects of visual awareness and attention distinct

from concrete objects through which we come to see the concrete. (Martin, 2012)

This sets the sense of sight apart from the auditory and olfactory senses. Seeing a flower in a botanical garden does not require seeing some visual particular distinct from the flower. We do not require something we could potentially capture, recreate or imitate. (This is not to say that seeing does not involve light reaching our eyes. Of course it does. The point is that this light is not itself a necessary object of vision, *pace* O’Shaughnessy, 1985) Not every distant visible scene has a corresponding visual representation—a visible image—in the way sounds or smells make the material world audible or smellable. Sure, one could pick the flower itself, dry it, and put it on display. That would have enabled Parisians to observe the flower throughout the cold winter. But it wouldn’t make the flower visible in its absence. On display would be just a flower, not a representation of a flower. As P.M.S Hacker puts it, “nothing in the domain of visibilia stands to seeing as sounds stand to hearing, smells to smelling...” (Hacker, 1987, p.209).

Visual mediators are exceptional, and go beyond the natural way sight presents its objects. A concrete object’s visibility, unlike its audibility or smellability, as such does not entail the existence of any perceptual mediators. This is rooted in the nature of the sense of sight; it is because of how seeing ‘works’, so to say. Should that make us suspicious of the kind of mediated visual representation I purport to have identified? If material bodies—including paintings and drawings—can be seen immediately, does this not mean that all visual perception must be immediate? That is, does it not rule out the possibility of mediation in visual perception as I identified it? An objection along these lines is developed by Brian O’Shaughnessy. According to O’Shaughnessy, vision is immediate because it is *directional*. He thinks this directional character of sight rules out that images, themselves visible bodies, could make something wholly

distinct from them visible.

In order to represent at all, images must themselves be visible objects. O'Shaughnessy argues that images, as visible objects, occlude some of our field of view whenever we see them. This, he thinks, entails that their appearance, instead of making other items visible, can only prevent those other items from being seen. Images, such as paintings or photographs, are not transparent, but opaque. If they are opaque, they can at best block other items from view, not make them visible. As O'Shaughnessy puts this:

If two objects are located at the same moment in time on the same spatial line coming from the eye, the only way visibility of both could be realised would be if the nearer of the two objects was what we call '*transparent*': if the near object is not transparent and occupies some given sector of the visual field, then it must *obscure* the remoter object.(O'Shaughnessy, 2003, p.175)

O'Shaughnessy thinks that if, by looking in a certain direction, the first thing that appears to sight is opaque, then another thing wholly distinct from it cannot appear in that same direction. He presents this as an utterly general principle of vision. It is how vision works as a distal sense. That Byng Place visually appears to me when I turn my eyes eastwards is because the window next to me is transparent, not opaque. Had it been wholly opaque, as a painting or photograph on that wall would have been, then the Byng Place coffee stand would have remained invisible from where I am.

O'Shaughnessy assumes that sight is essentially directional. Whenever we see some particular thing, we see it down an axis in space, a line commonly described as a line of sight. O'Shaughnessy assumes that any object of vision must appear to sight in a specific direction out from us in tridimensional space (O'Shaughnessy, 2003, p.175). He takes this as a precondition for vision's being a distal

sense. And so if the bathers in Cézanne's painting can genuinely appear to us by virtue of seeing their image, as I suggest, then they could only do so by appearing in exactly the same direction as the visible canvas.

Note, however, that direction is not location. Something can be visible along a line of sight without being located along that line of sight—as when your face becomes visible in the direction of the bathroom mirror, without being located in the direction of that mirror. Similarly, there is no reason to suppose that Cézanne's bathers must be located somewhere in the direction of Cézanne's canvas in order to appear along that line of sight. (The suggestion that images and mirrors present us with a distinct, 'virtual' space seems to me based on a tacit conflation of direction and location.)

If O'Shaughnessy is right, then it seems that Cézanne's painting could only render visible an otherwise invisible group of people if it were transparent (or if, for some reason, the concepts of transparency and opacity did not apply to it; O'Shaughnessy himself thinks that this is true of non-extended sense data, which he assigns a crucial role in perception, see also O'Shaughnessy, 2000). But paintings are not transparent, because they are opaque.

The problem for the view about visual representation I defend arises because the following three premises seem to be inconsistent (the third premise naturally follows from what I have argued above).

1. If, by looking in a specific direction, the first thing that appears to sight is opaque, then another thing wholly distinct from the first cannot appear to sight in that same direction.
2. Images are opaque.
3. If, by looking in a specific direction, the first thing that appears to sight is an image, then another thing wholly distinct from the first can appear to sight in that same direction.

I take it that O'Shaughnessy is right about the relation between dir-

ectionality and opacity, and so I accept the first premise. Am I then committed to denying the second premise? Must I deny that images are opaque, appearances to the contrary?

Denying that images are opaque need not be incoherent. Kendall Walton has suggested that at least some images make things visible because, appearances to the contrary, they are transparent. He claims that this holds for photographs. Walton argues for this claim on the basis of similarities between the photographic mechanism and the workings of other optical instruments, such as telescopes, periscopes and ordinary mirrors (Walton, 1984). Walton's suggestion is revisionary, and has been met with some resistance. Yet what matters most for the current discussion is that it conflicts with the idea of images as perceptual mediators. According to Walton, a photograph makes some scene distinct from it visible to the extent that the photograph itself is *not* perceptible, but transparent. Yet something functions as a perceptual mediator only to the extent that it is perceptible. The opacity of images as mediators is not an contingent feature; it is necessary. By giving up the assumption that images are opaque, we would be denying that they meet a necessary requirement for perceptual mediation. (Michael Newall, 2010, argues that images, though not in fact transparent, at least look to be transparent. Even if that were right, which I doubt, it would not avoid O'Shaughnessy's objection.)

The way to avoid O'Shaughnessy's objection without giving up the opacity of images, is to turn back to the structure of the visual modality. In vision, as in audition and olfaction, we find perceptual mediators. Mediators in audition are sounds, mediators in olfaction are smells. In vision, I have argued, images fulfil a mediating role. However, there are nonetheless sharp structural differences between these various sense modalities. As I mentioned above, sounds and smells are sensible particulars that figure in any act of hearing or smelling. Images figure only exceptionally in vision. Images can mediate perception, but visual perception of the bodies and events around us

does not depend on that possibility. This allows us to draw a sharp line between vision on the one hand, and sensory modalities such as audition and olfaction on the other. That images figure only exceptionally in vision implies that you can see one and the same object both immediately and mediately. You can see it immediately by encountering it in real life, but also mediately when you encounter an image of that object. At least for us, vision is what we could call a *heterogeneous* sensory modality.

Sensory heterogeneity A sensory modality is heterogeneous if and only if in that modality we can perceive one and the same item both immediately and mediately.

By contrast, in auditory perception only sounds can be heard immediately, whereas objects and ordinary events can only be heard mediately, namely by means of sounds. We cannot hear one and the same item both immediately and mediately. Hence, audition is a *homogeneous* sensory modality—a sensory modality that is not heterogeneous.

If sight is heterogeneous in the way I suggest, then the problem with O’Shaughnessy’s objection becomes clear. It is that the first and last premise can only both be true if they equivocate two ways something can appear to sight. All premises are true only if we read

1. If, by looking in a specific direction, the first thing that appears to sight is opaque, then another thing wholly distinct from the first cannot *immediately* appear to sight in that same direction.
2. Images are opaque.
3. If, by looking in a specific direction, the first thing that appears to sight is an image, then another thing wholly distinct from the first can *mediately* appear to sight in that same direction.

The conflict disappears. Reflecting on the way vision works as a distal sense, we may very well accept that, indeed, opaque objects prevent further items from appearing to sight immediately. Yet, as soon

as we accept that in visual perception we can draw on capacities foreign to the sense as such, we see how, when opaque objects function as mediators, they may expand what is visible along a single line of sight.

Suppose you look at a large television screen, presenting you with the televised image of a group of people scuffling. And suppose that behind the screen there are persons exactly like those shown on the screen, doing the same things. The image on the screen is a wholly opaque surface, distinct from the group of people it represents. As the opaque visible item it is, the screen obscures from sight whatever is located in the same direction. Yet given that it is an image, it can make the group of people visible in a mediated way. Although it prevents the group from appearing to sight immediately, it allows the fighting people to appear mediately, because of the way the visible image can function in our experience of the situation. It enables us to direct our gaze and look at a group of people and their struggles as they unfold.

O'Shaughnessy writes about the television image that despite the fact that it 'shows' the televised situation in one sense, and despite the overlap in cognitive utility between our seeing the image and our seeing that situation 'in the flesh', it is only when we look at the scene immediately that we are actually able to 'set our eyes upon' that scene (O'Shaughnessy, 2003, p.180). He takes this to prove that the way the television image presents us with the visual appearance of the scuffle does not amount to its visibility. If I am right, this is a mistake. As soon as we acknowledge that vision allows for a special kind of perceptual contraband—a kind of sensory appearance foreign to it, inessential to its normal functioning—we may accept that an image does confer a distinct kind of visibility on whatever it represents to the eye. To be sure, when a scene appears to us representationally in vision, we to that extent precisely do not perceive it 'in the flesh', and all we actually 'set our eyes upon' is an image. This however does not take away that the scene does appear to us. It appears to sight in a way that

belongs only naturally to those senses that require perceptual mediators to figure in any instance of their exercise.

I accept that vision's heterogeneous structure is not the first thing on most people's minds when they reflect on what we could possibly come to see. Few will conceive of the visible world as offering us both immediate perceptual awareness of the objects and events in our environment, as well as a mediated form of vision that exploits some image we are presented with. However, that something is not obvious does of course not mean it does not occur.

3.6 Representational appearances

The central thesis I have argued for in this chapter is that images are perceptual mediators for the sense of sight. This answers the question we were left with after having set aside the copy theory of visual representation. How is it possible for an image to make an absent scene visible? The answer is that images can do this, because they bring to vision the opportunity for mediated perception. When an image functions as a mediator in someone's perception of the world, it makes some object or scene wholly distinct from it appear to sight. It can do so in the object's absence. This explains how images represent.

If what I have argued is right, then it means that perception—vision in particular—does not require the presence of all of its objects. Charles Travis writes that

perception, as such, simply places our surroundings in view, affords us awareness of them [...] It confronts us with what is there...(Travis, 2013, p.65?)

We can now see that thinking of perception as merely a mode of awareness of what is there is too simple. Undoubtedly some instances of perception simply place our surroundings in view. Moreover, I

agree that it is natural to assume that vision always brings at least some of our surroundings into view. To see is always to become aware of something that is present to the eye. But as soon as we take account of the possibility of perceptual mediation, we see that perception—even visual perception—sometimes does more. When we are confronted with a sound or a smell, or indeed an image, some wholly distinct object or scene may be audible, smellable, or visible to us. This would enable us to perceive something that need not actually be present there at all.

It is perhaps telling that Travis restricts his discussion of perception to observations about vision. He does not consider other sensory modalities. In particular, Travis omits reference to those senses, such as audition and olfaction, that for their functioning as distal senses depend on mediators. Yet it is precisely this possibility of mediation that our perception of images exploits. It is tempting to think that precisely because of his focus on vision, Travis was bound from the outset to reach only a partial understanding of what visual perception, as such, can do.

To my mind, it is clear that philosophical inquiry about visual representation should start to consider the role images can play in perception. I have identified that role, and have shown how this fits with our everyday engagement with the images we know. Images function as perceptual mediators, and as such enable us to perceive the scenes they represent mediately. This leaves it open how exactly we should conceive of mediate perception, as a state of sensory awareness. What is the difference between merely hearing a rattling sound and hearing a coach by means of that sound? What is the difference between merely seeing the paint on the surface of Cézanne's *Les Grandes Baigneuses*, and seeing a group of nudes by means of that paint? Some have suggested that this difference is not genuinely sensory. They think that to perceive such a group of nudes by means of an image is at least in part constituted by some non-sensory act of

cognition. In the next chapter I will show how we need not accept that view. Instead, I will defend an understanding of mediate perception as a purely sensory state. Instead of relying on non-sensory or quasi-sensory effects of thought or cognition, coming to see a painted surface as an image consists in no more than discovering a previously unseen aspect of the visible world.

4

Coming into view

In the final exchange of Berkeley's *Three Dialogues between Hylas and Philonous* (1713), we find a dispute over our visual perception of images. More specifically, the two protagonists disagree about what can become visible to us when we perceive a portrait of Julius Caesar.

PHILONOUS: Is there any thing perceived by sense which is not immediately perceived?

HYLAS: Yes, Philonous, in some sort there is. For example, when I look on a picture or statue of Julius Caesar, I may be said after a manner to perceive him (though not immediately) by my senses.

Hylas thinks the portrait makes Caesar visible. He takes this observation about representational images to be obvious—a claim not in need of defence. In spite of that, Philonous at once criticises him. He takes Hylas' claim to rest on a vulgar assumption.

It pays off to understand better the condensed argument that unfolds between both students. They disagree about the mediating role of images. At the heart of their dispute lies a question about the nature of mediate perception. I have already shown that images are perceptual mediators for the sense of sight. They can make absent scenes visible precisely because seeing an image puts us in a position to per-

ceive such scenes mediately. But if Philonous' objection to Hylas is right, then the mediating role images fulfil in perception is more restricted than I have claimed. If Philonous is right, mediate perception cannot be a purely sensory mode of awareness, and so is at best sensory only in part. My main aim in this chapter is to show that Philonous' argument is unconvincing. I will defend Hylas' theory of visual representation, that seeing a painting as an image of Caesar is a purely sensory state that is not just a perception of the painting itself, but also a perception of Caesar.

4.1 Hylas' theory of visual representation

What exactly does Hylas claim? He claims that looking at an image of Julius Caesar painted by an acclaimed portraitist can make that Roman general in some way appear to sight. Think of Peter Paul Rubens' portrait of the emperor (Fig. 10). When you look at it and see Caesar by means of it, Hylas takes it to be plain that you will now perceive Caesar. He stresses that this would not be an immediate way of seeing him. But one would be seeing him none the less. Hylas takes this to reflect a general truth about images. Images allow us to see all sorts of things by means of them, things we would perhaps not even be able to see without those images. For an eighteenth century viewer, Julius Caesar would be a case in point.

These ideas about representational images figure in a broader discussion about the way we can perceive the material world. Hylas wants his observations about portraiture to reveal a truth about our perception of the material world in general. Let me briefly present the dialectic.

At this stage of their argument, Hylas and Philonous have already agreed that a material world is not something we can see immediately. Material bodies cannot affect our senses, only sensible ideas can. In vision, we immediately only ever perceive ideas of colour and



Figure 10: Peter Paul Rubens, *The Emperor Julius Caesar*, ca. 1626.
Stephen Mazoh, New York

figure. By contrast, matter cannot become present to the mind in this way. Both parties accept this. Yet Hylas is adamant that we can perceive matter in some other way. We can perceive material bodies by means of sensible ideas. He thinks that Philonous has overlooked a kind of perception we know from our experience with paintings and other forms of representational art. This is where his theory of visual representation comes in. Hylas claims that seeing a painting as an image of something is a purely sensory state. It is not just a perception of the image itself, but also a perception of the thing the image represents. Though, he emphasises, the latter perception is not an immediate one. Looking at a picture of Julius Caesar, you may come to perceive Caesar in a mediated way. Hylas suggests that sensible ideas of colour and form may function just like the painting: they function as images that mediate our perception of material bodies.

Hylas uses a common sense assumption about painting to support a theoretical conclusion. He thinks the possibilities of painting reveal

something about the role sensible ideas play in our perception of the world. Philonous does not accept Hylas' argument. He does not agree that ideas enable us to perceive the material world. This is because he does not accept the assumptions about painting it relies on. According to Philonous, we do not get to perceive Caesar just by seeing his image. That would be a vulgar misconception.

That Caesar's image can provide an onlooker with some sort of cognitive access to Caesar is something Philonous need not deny. What he denies is that the way an image can bring Caesar to mind can ever amount to a sensory appearance. When we look at Caesar's portrait, it is only an image of Caesar that appears to sight, not a Roman emperor. Philonous' reasoning in support of this claim uses a comparison of the sort Berkeley's *Dialogues* are famous for. He suggests that we compare what appears to Hylas—a viewer 'in the know', who is able to see the portrait as an image of Caesar—with what could appear to a viewer who has no idea of Caesar and who does not see the portrait as an image of Caesar.

PHILONOUS: Tell me, Hylas, when you behold the picture of Julius Caesar, do you see with your eyes any more than some colours and figures with a certain symmetry and composition of the whole?

HYLAS: Nothing else.

PHILONOUS: And would not a man, who had never known anything of Julius Caesar, see as much?

HYLAS: He would.

PHILONOUS: Consequently he hath his sight, and the use of it, in as perfect a degree as you.

HYLAS: I agree with you.

PHILONOUS: Whence comes it then that your thoughts are directed to the Roman Emperor, and his are not? This cannot proceed from the sensations or ideas of sense by you then

perceived; since you acknowledge you have no advantage over him in that respect. It should seem therefore to proceed from reason and memory: should it not?

HYLAS: It should.

Philonous begins with a rhetorical question. He asks whether Hylas, beholding the picture of Julius Caesar, sees with his eyes any more than some colours and figures with a specific symmetry and composition. That he asks just this may seem odd. Does Philonous not beg the question here? Has Hylas not just voiced his stance on precisely this issue? Hylas just claimed that, yes, we can see more than colours and figures with a certain symmetry and composition: we can see Julius Caesar!

Someone could suggest that Philonous' question is merely rhetorical, intended to highlight from the outset that Hylas' stance is plain wrong. This cannot be right. What would compel Hylas to accept Philonous' position based on rhetoric alone? Philonous knows full well he needs an argument. And Hylas' reply—"Nothing else"—makes it clear that he is in full agreement with Philonous on some point or other. But this point cannot be that we cannot see Caesar. If this is right, then to which point does Hylas assent instead?

The puzzle only arises if we take Philonous to ask about what Hylas sees, period. On that assumption Philonous' question would indeed already presuppose that Hylas can only see a configuration of colours and figures. But this is not what Philonous is after. We should take Philonous to ask whether Hylas has any visual perception that is not a perception of colour and figure. And that is a subtly different question. Philonous pushes Hylas to admit that what is seen in every visual perception is a configuration of colours and figures. Hylas can perfectly well accept this. As they agreed earlier, every visual perception is a perception of ideas of colour and form. If visual mediation occurs at all, then these ideas are the 'earliest' mediators that appear to sight (see O'Shaughnessy, 2000, p.553 for a recent discussion of the

idea of an earliest mediator). Because there is no visual mediator prior to ideas of colour and figure, we may consider our perception of these earliest mediators immediate (cf. Rickless, 2013, chap.1). Now if Hylas is right in thinking that his perception of Caesar is purely sensory then the emperor's coming into view must somehow be enabled by perceiving these earliest mediators we apprehend in sight. Caesar could only genuinely appear to sight if his appearance involved the appearance of sensible ideas of form and colour. Only those ideas could confer visibility on anything distinct from them that we may perceive. If it were possible to see Caesar by looking at his portrait, as Hylas claims, then it would first and foremost be due to the visibility of a cluster of sensible ideas of form and colour.

This shows that Philonous' argument so far is sound. He legitimately proposes to compare what earliest mediators both viewers of the portrait can get into view. With his initial question, Philonous just wants to make sure that he and Hylas are on the same page: if we set out to compare visual perceptions of anything whatsoever, then we are always comparing visual perceptions of ideas of colour and shape. Hylas expresses his agreement, because he understands Philonous' comparison is specifically focused on the sensible ideas both viewers can get into view. He has already accepted that every visual perception we have is a perception of ideas of colour and form. The fundamental dispute between the students is whether some of these perceptions are also perceptions of Julius Caesar.

Hylas' theory of visual representation claims that seeing a painting as an image of something is a purely sensory state. It is not just a perception of the image itself, but also a perception of the thing the image represents. If Hylas' theory is correct, then seeing the portrait as an image of Caesar must depend on an immediate perception of colours and forms of some visual mediator. This idea is crucial, and it is only implicit in Hylas' initial claim about depiction. Both students agree that the relevant visual mediators to consider here are

sensible ideas, but that is because of further assumptions they make about perception. What matters is that Hylas' theory reflects a subtle understanding of the relation between mediate perception and our perception of the mediators that make it possible.

Philonous unearths this subtle way of thinking about perceptual mediation only to undermine its credibility. He is happy to conceive of Hylas' perception as a mediate perception of Caesar. But he explicitly denies that this is a purely sensory mode of awareness, as Hylas thinks. The difference between merely seeing an image and seeing something by means of that image, he thinks, could not be a sensory difference.

What is Philonous' argument for that denial? He starts off from the observation that perception of ideas does not require any knowledge, experience, or expertise. Coming to perceive a sensible idea is just something that happens to us. Hylas had already conceded this point earlier in the dialogue, perhaps not fully grasping its ramifications. (Myles Burnyeat, 1982 brings out well how, as soon as we accept Berkeley's conception of ideas, the arguments of the *Dialogues* unfold smoothly) As Philonous claims, someone who had never known anything of Julius Caesar would nonetheless be able to see the exact same ideas of colour and figure as Hylas sees, and see them just as perfectly. Between Hylas and someone who does not see Caesar in the picture there need not be any difference in what appears to sight. And so Hylas' claim that some of the visual perceptions of colour and form he has when he looks at the portrait are at the same time perceptions of Caesar is demonstrably false. It is false, Philonous infers, because anyone able to see could come to perceive exactly the same pattern of colour and form, yet not see that pattern as an image of Caesar. Hence, the difference between merely seeing an image and seeing something by means of that image could not be a sensory one.

Philonous argues that the connection between the visible features

that appear to sense and the way they are able to bring Caesar to mind cannot be as intimate as Hylas supposes. This is because it is conceivable that they come apart. That is what he meant to bring out with his comparison. Someone could come to see all the relevant visible features of a painting, yet fail to see it as an image of something or other. And so, playing along with Hylas' own example, Philonous suggests that when an image of Caesar appears to sight, this can at best cause you to think of Caesar. Your perception of the image leads you to conceive of him in some non-sensory way. What Hylas mistook for a sensory appearance of Caesar in fact was a mere *suggestion* to the mind—some non-sensory act, proceeding from reason and memory.

Regardless of what we think of its conclusion, Philonous' argument embodies a deep insight about perceptual mediation. In the *Dialogues* this insight is somewhat hidden from view. One might think that Berkeley realised that any materialist sympathetic to a Berkelean psychology is likely to cling onto Hylas' specific views about mediate perception as soon as they get the chance. Berkeley wants to resist those views, and instead adopts the solution he presents by way of Philonous' line of argument. Perceiving one thing by means of something else ultimately is no more than a non-sensory act of suggestion. Such an act may occur in response to a perception, but it is not itself a sensory achievement. In the next section I will consider how well Berkeley's concept of suggestion captures the initial situation Hylas introduced. Can a non-sensory act of suggestion explain the way things appear to us when we see a portrait of someone as an image? I will show that when limit ourselves to ordinary visible images such as painted portraits or photographs, Berkeley's assumption that coming to see them as an image makes no difference to sensory appearances is hard to maintain.

4.2 Suggestion and sensory variation

Philonous' argument against Hylas' theory of visual representation turns on its head the ancient puzzle of conflicting appearances. The same spoonful of honey may appear bitter to the one, while it appears sweet to the other. The same honey cannot both be sweet and bitter, it seems. Is one of these appearances misleading? Are both of them?

Berkeley frequently uses puzzles of this form to show that the qualities that appear to us in perception are in fact only immaterial ideas on the minds of a perceiver (see Burnyeat, 2012). For example, Berkeley would reply that the honey conceived as a mind-independent substance is neither bitter nor sweet. The perceived bitterness belongs only to the idea of the honey the one perceiver has. The reported sweetness belongs to a different idea of the honey, perceived by someone else. Berkeley does not conclude from this that bitterness and sweetness are unreal. Instead, he emphasises that they are qualities that do not inhere in a material, mind-independent world, contrary to what the vulgar think.

In the *Dialogues*, we read how Hylas' suggestion about painting is countered by Philonous. He uses the same dialectical device of comparing how things appear to one perceiver with how things appear to someone else. Take two perceivers of the same portrait, and compare what is visible to the one with what is visible to the other. But now Philonous does not focus our attention on any conflict. Instead, he aims to highlight just the opposite. Philonous tries to bring out that the way things appears to someone who sees a portrait as an image of Caesar need not be any different from how things appear to someone who does not see the image as such. From this lack of a difference in appearance Philonous concludes that both perceivers are just seeing the same things. Philonous inverts the argument from conflicting appearances. Compare, if we know there is no difference between the way the same honey tastes to me and the honey tastes to you, then we have no reason to think there is any sensory difference, even though

we still may describe its qualities differently (you may say it is “sweet”, while I say it “tastes like sugar”).

Is Berkeley right in presenting the argument in this way? What exactly follows from a lack of a difference in appearance? Not, it seems, that what two people see are the same things. What appears to me when I look into the gallery may turn out to be a ringer of what you saw there yesterday. What I now see is a styrofoam replica, say, while you saw the expensive marble original. Surely both of us saw something different, yet the difference need not show up in how things appeared to us visually. This observation is reasonable enough. But we should remember that Berkeley right from the beginning wants to focus our attention on the ideas we have. And here, Berkeley thinks, matters are different. Ideas that are exactly like each other could not be different. Again it is an assumption about the peculiar nature of ideas that allows Berkeley to move to his conclusion.

On Berkeley’s picture of psychology, if you and I perceive the same ideas of colour and form, then there is no room for any variation in what each of us manages to see. This means that any further differences in our engagement with what we see must lie in non-sensory factors. Suppose you and I taste the same substance. Yet only I think it tastes like honey, while it merely appears sweet to you. According to Berkeley, the difference here could only lie in the fact that I come to think of honey when tasting the substance, but you do not. Similarly, Hylas may see the portrait as an image of Caesar when he looks at the colours and forms of the paint on its surface. Someone else may see exactly the same painted colours and forms, yet fail to see the portrait in the way Hylas does. Berkeley claims the difference between them could only be due to some non-sensory factors. What factors could this be? In the discussion about painting in the *Three Dialogues* we find Berkeley laying out the answer he prefers across the board: Caesar is merely *suggested* to someone who sees the portrait as an image. Suggestion is an act distinct from perception of sensible

ideas.

What exactly is suggestion? Berkeley conceives of it as a crucial mental operation. Suggestion governs the endogenous presentation of ideas in response to earlier ones (Graham, 1997, p.403ff). The perception of one sensible quality may lead the mind to expect another. This mental expectation would bring to mind ideas earlier stored in memory. (Arguably, here Berkeley prefigures the currently fashionable idea of ‘predictive coding’. See Clark, 2015 for an overview.) Berkeley uses suggestion throughout his writings. His discussion of the perception of distance or ‘outness’ is perhaps most familiar. According to Berkeley, we never immediately see distance. We only perceive distance mediately. On seeing an expanse of colour laid out before us, ideas of touch are suggested to the mind. They suggest to us the tridimensional make up of our surroundings. Our thoughts about the distance of things are never based on our seeing their distance, but at best based on what is suggested to our minds.

What holds for distance holds for material bodies and their overall shape and size as well. As Berkeley explains in his *New Theory of Vision*,

We cannot open our eyes but the ideas of distance, bodies, and tangible figures are suggested by them. So swift and sudden and unperceived is the transition from visible to tangible ideas that we can scarce forbear thinking them equally the immediate object of vision. (*New Theory of Vision*, p. 145)

Mediate perception, as Berkeley understands it, is a complex mental act. A sensory core gives rise to suggested ideas (Hatfield & Epstein, 1979 discuss the historical precedents for Berkeley’s complex conception of perception in early modern philosophy). The most important thing about the way Berkeley conceives of suggestion is that it is non-sensory. What is merely suggested to the mind does not ap-

pear to the senses, and suggestion cannot change how things appear to you. Berkeley consistently contrasts suggestion with sensory appearances throughout his work. Whenever an idea is suggested to the mind as a result of sense perception, suggestion and perception remain distinct, even though we may be unable to distinguish the two acts on introspection (cf. Rickless, 2013, p.77). In this respect Berkeley proves himself to be the keen student of Descartes he professes to be. Also the author of the *Meditations* maintained that the way the world affects us through the senses is independent of the ‘intuitions of the mind’ formed in response to the sensations or sensible ideas we have.

The crucial assumption Berkeley makes about suggestion is that it does not affect how the world before our eyes appears to us. This has implications for how he conceives of our perception of images. As is clear from the discussion about Caesar’s portrait in the *Dialogues*, Berkeley assumes that the way Caesar is suggested to the mind of a suitable viewer of the portrait is not a sensory act, and does not bring about any sensory change. Coming to see the painting before you as an image of Caesar does not bring about any change in how the painting appears to you.

Precisely on this point philosophers have been reluctant to follow Berkeley. And it seems clear to me that their reluctance is well-grounded. Several examples have convinced people that seeing a painting or drawing before you as an image of something or other does make a difference to how that painting or drawing appears to you. What has convinced people here is the way coming to see something as an image of one thing or another manifestly makes for a visual, sensory difference. As Malcolm Budd remarks,

the awareness of what a picture depicts is experiential, as can perhaps be seen most economically in the switching of awareness that can take place in the perception of ambiguous figures. (Budd, 2008a, pp.205–6)

Budd's point is that most of us are able to verify that seeing something by means of an image is truly visual, by reflecting on their experience of coming to see something as an image of something or other. Budd only mentions ambiguous figures. Yet there are in fact three kinds of examples that clearly speak against Berkeley's description.

Sometimes we suddenly are able see a familiar image differently. We realise that the image is ambiguous. These are the cases Budd has in mind. The Necker cube is a good example. This figure allows us, more or less at will, to shift back and forth between two different ways of looking at what it represents (Fig. 11). If you try, you can see the figure both as an image of a cube with point *A* in front of point *B*, and as an image of a cube with point *B* in front of point *A*. It seems clear to me that shifting back and forth between these two ways of looking at the figure makes for a sensory difference.




Figure 11: Necker cube. First published in 1832 by Louis Albert Necker

We find a second kind of example when we reflect on so-called puzzle pictures. Not everything a picture represents is always easy to discern. Especially when a represented scene is full of detail, and when some details are cleverly disguised, it may require effort to get specific objects or features into view. 'The Puzzled Fox' is such a

puzzle picture, designed to make it hard for us to see the specific animals it represents (Fig. 12). Can you see the young cow? In trying to answer the question you may look around and search the image. When you finally spot it, if you do so at all, you come to see part of the picture as an image of a young cow huddled against a tree. Also here it seems manifest that once we manage to see the picture in that way, the way it appears to us has changed.



Figure 12: Currier and Ives, *The Puzzled Fox*, 1872. New York

Some images make things even more difficult. They may initially present themselves to a viewer as no more than a seemingly random collection of marks on a surface. They make for a third kind of example that may convince us that seeing something by means of an image is genuinely sensory. Consider the by now familiar image of a dalmatian dog scavenging a street (Fig. 13). This collection of blobs is an image. Yet to many viewers it initially appears as no more than an array of random marks. As the psychologist Stephen E. Palmer writes,

People who have never seen this image before usually see

it as a seemingly random array of meaningless black spots and blobs on a white background. However, once they have seen it as a dalmatian with its head down, sniffing along a street, the picture becomes dramatically reorganised. Certain of the blobs go together because they are part of the dog, and others go together because they are part of the street or some other object. (Palmer, 2002, p.190)

The initial viewers Palmer describes find themselves in the kind of position Berkeley exploits in his thought experiment. They are looking at an image, but they are unable to see the colours and forms before them as an image of a dog. Berkeley wants us to believe that coming to see the dalmatian by means of those marks—coming to see them as an image of a dalmatian dog—does not comprise a sensory change. Yet as Palmer emphasises, this is not at all how it seems to us when we come to see the image in that way. As he observes, the way the image appears to us changes dramatically. When suddenly the dog pops out, it seems manifest that a sensory change has occurred.

Given the salience of these phenomena, it seems fair to conclude that Berkeley's description of our perception of visual images is unattractive. His focus on an *interpersonal* case was somewhat misleading all along. The kind of variation he aimed to rule out shows up most clearly when we consider *intrapersonal* shifts in perception. Reflection on our own experience of seeing the Necker cube figure, the puzzle pictures so popular with children, and the 'hidden' images exploited in psychological experiments, makes manifest to us that Berkeley's description falls short. Coming to see something as an image of something or other makes for a change in what appears to sight. That was Budd's point too. Coming to see something as an image is a sensory achievement, and not merely a change in what we think, believe, or know about what we see. To anyone who agrees with this, Berkeley's insistence to the contrary will seem unfounded.

To be sure, Berkeley may still insist that we are merely misinterpret-



Figure 13: R.C. James, *Dog Picture*, 1965. LIFE magazine

ing our experience. Where we are confident that coming to see that dalmatian dog is a sensory achievement, Berkeley could object that we base our confidence on no more than a deceptive interplay of perception and thought. Some of what Berkeley writes about suggestion brings out that he would indeed make this retort. Yet, to my mind such a claim would only be acceptable if there were some reason to think that what we so confidently claim is impossible. If coming to see something by means of an image simply could not be a sensory achievement, then we would have to take Berkeley's alternative seriously. But as I will explain in the next section, at least in the situation Hylas has in mind, there is just no reason to think that. We can defend Hylas' specific theory of visual representation by pointing to an utterly mundane fact about perception.

4.3 Perception's partial character

Right after Raskolnikov has gruesomely murdered two women in St Petersburg, he looks over his overcoat, trousers and boots to clean off any bloodstains that may betray him. At a first glance, he only sees some spots on the boots. He rubs them off with a wet rag. But Raskolnikov knows that there may well be some noticeable evidence he is overlooking. He is right. As Dostoevsky tells us later, thick droplets of blood are still clinging to the frayed edge of his trousers. But he simply does not see them. Perhaps it is the poor light. Perhaps his manic state of mind is to some extent impairing his vision. Either way, Raskolnikov's inability to see all that is there to see is mundane. It reflects a feature of sense perception in general. Sensory perception is typically *partial*. Whenever we perceive something, we inevitably perceive only part of it.

What exactly is it for perception to be partial? The main idea is that in any given perception, there may be more to the sensible qualities of an object or scene than is manifest to you. Some objects have properties that evidently are not perceptually available from certain perspectives. Surrounding flesh and skin hide the colour of someone's lungs from view, at least in normal circumstances. And so is the cat occluded by a large suitcase. But the point holds more generally. A perception from any perspective point does not fully determine the sensible qualities perceptible from that perspective. This is because different sensible qualities might be perceptible only to different perceivers (for further implications of perception's partial character, see Hilbert, 1987, p.37; Kalderon, 2007, p.537).

Raskolnikov's tragic inability to see all there is to see is our predicament. It does not matter how hard you try to get things into view, there are bound to be things, properties, details that you overlook. If this is right, then a full specification of what someone sees cannot just be in terms of what is visible, but also has to be sensitive to include the *circumstances* of perception and the *sensibility* of the perceiver.

Circumstances and sensibility can vary independent of one another. Each on its own can have subtle effects (Kalderon, 2007 exploits this insight in resolving apparent conflicts in our experience of colour).

Circumstances On the one hand, you may see on the knight's armour a specular highlight invisible to me, even though we both have front row seats. The brushwork of a painting may partly escape me because of the tungsten illumination in the gallery. These are subtle effects of equally subtle differences in circumstance. A slight movement may be enough to lose sight of a play of light. A specific illuminant may reduce the visibility of texture or relief.

Sensibility On the other hand, I may be able to see a ligature in the manuscript you fail to discern. You might be able to see the hand of a master, even though I see no more than yet another still-life painting. In both cases our circumstances of perception may be the same. What we see diverges because of a difference in sensibility. I know palaeography; you have carefully studied painterly techniques. Both make us sensitive to details easily overlooked.

Even though many are familiar with an image, its representational aspects may not be visible to everyone. Imagine looking at the painting from close up. It is likely that at some point you will no longer be able to see it as an image of Caesar, and only end up seeing some strokes of paint. Just as above, circumstances and sensibility influence our perception of images, by allowing or preventing us from getting a view on what is there. Such factors affect what someone is able to get into view. Reflect back on the image of that dalmatian dog discussed above. It is likely that by now you will naturally see the figure as an

image of a dog. But try to recall what happened when the represented dog suddenly dawned on you. What changed? Given our predicament as perceivers, the most straightforward answer is that you suddenly managed to see visible aspects of the image you did not see before. In particular, you managed to get into view those aspects that represented the dog. If that is what happened, your coming to see the image as an image of a dalmatian was a purely sensory achievement.

Mark Johnston illustrates how complex aspects of the visible world may only be visible to those with patience or in the know (Johnston, 2006). He writes about the experience of coming to see a bluff at poker. In some situations a player's bluff is a visible aspect of the way they appear. Something in someone's face may give it away. Or perhaps it is the way they press their heel against the leg of a chair. Trained poker players have learned visually to discern aspects of the other players that are likely to escape a novice. A trained player has cultivated a sensibility to the look of someone bluffing. This gives them what Bernard Williams called a 'positional advantage' over the novice (2002, p.42). It is natural to think that once you managed to see that dalmatian, similarly you have acquired a positional advantage over those viewers that see the image for the first time.

I should emphasise that the visible aspects you manage to get into view when you come to see something as an image are not always identifiable merely in spatial terms. Some aspects of visible objects are so identifiable. Think of the rear view of a mansion or the left corner of a painting. But not all aspects are. The different colour patches in the neck of a pigeon are a clear example. Only some of these may be visible to you at a time. Yet a visible colour patch in the pigeon's neck may coincide spatially with one currently invisible to you. Seeing the Necker Cube in one way brings a visible aspect into view that shares its location and direction with the other aspect you currently do not see. We may think of visible aspects more loosely as potentially complex visible features of our surroundings.

Vision's partial character saves Hylas' theory of visual representation. No matter how carefully we look, there always will be more to see. Even with a clear view of an object, some viewers may simply be unable to see details that to others stand out as obvious. Hylas claimed that he saw Caesar by means of a painting. He assumed that the difference between someone's merely seeing the paint on the surface of that portrait, and someone's seeing Caesar by means of that paint, is a purely sensory difference. So Hylas was confident that he was visually aware of some individual that was not immediately given to him in perception. Philonous resisted this. While accepting that Hylas may have perceived Caesar mediately, he denied that it was a purely sensory mode of awareness. This must be so, Philonous held, because someone who did not know anything of Julius Caesar would still be able to see the exact same colours and figures that Hylas saw. Moreover, that person would be able to see them just as accurately. This led Philonous to conclude that the difference between merely seeing the paint and seeing Caesar by means of that paint simply cannot be a sensory difference.

As soon as we acknowledge perception's partial character, we have good reason to reject Philonous' premise. A full specification of what someone sees must include more than what is visible before their eyes. It also has to make reference to the circumstances and the sensibility of the perceiver. When you and a companion stand in front of a painting, you may always be able to see things that your companion misses, no matter how carefully they look. Someone can simply remain oblivious to those aspects of a painting that make Julius Caesar visible. Reflection on one's own experience of coming to see something as an image can confirm this; it makes it clear that this is a purely sensory achievement. Therefore, we may accept that when Hylas sees Caesar by means of the painted portrait, his perception of Caesar, though not immediate, is a purely sensory act.

Mundane as Hylas' suggestion about visual representation turns out

to be, Berkeley seems oddly insensitive to all of this. It may seem striking that in Berkeley's discussion the partial character of perception—a pervasive and everyday fact—is left out. How do we explain the omission? I want to suggest this may have a double origin.

First, consider the situation Philonous asks us to imagine. Two viewers look at the same portrait. By hypothesis, the conditions of perception in this case are stable. We may presume that position, lighting, and viewing distance are just the same. The only thing we hear is that one of these viewers knows nothing about Caesar. Philonous' scenario fosters a temptation to think that both get exactly the same view of the painting before them. Yet we now see that this would be a fallacy. As the examples above show, it may be precisely a difference in sensibility that explains why the one viewer gets to see Caesar, while someone else does not. Indeed, it may precisely be that lack of knowledge that prevents the one viewer from seeing what Hylas was able to see.

The second origin of Berkeley's oversight lies deeper. Recall, Berkeley is mainly interested in the hypothesis that sensible ideas, as visual mediators, could enable material bodies to appear to sense. That is why he stages the discussion about painting in the first place. Hylas suggested that ideas may be just like paintings. Yet as soon as we consider our perception of mere ideas, Berkeley is on firmer ground. According to Berkelean psychology, sensible ideas make for a notable exception to the partial character of perception. As Berkeley defines them, to perceive an idea of sense is to perceive *all* of it (Berkeley, 1954, First dialogue; Burnyeat, 1982). In contrast to two viewers who look at the same painting, or at the same colours and shapes on its surface, two viewers who perceive the same sensible ideas would necessarily perceive the same. Given the nature of ideas, there would be no room for any difference in what both viewers get into view. To some extent, this fact about Berkelean psychology protects at least Berkeley himself from criticism.

4.4 Penetrating vision

In *Of the Standard of Taste* (1757), Hume presents two wine critics as they figure in a noted story in *Don Quixote*. To much amusement of the onlookers in a crowded tavern, the critics' refined verdicts about a wine seem to conflict. After they have tasted the same wine, the one judges it to have a taste of leather, while the other denies this and judges the wine to taste of metal. For a brief moment it seems a case of snobbery is unmasked. As it turns out, both are right. The guests discover a key with a leather strap on the bottom of the barrel.

We may take this appropriation of Cervantes' story as Hume's way of emphasising the partial character of perception. Because perception has such character, it can give rise to apparently conflicting verdicts about what we see, taste, or hear. Such conflicts are typically merely apparent, Hume might suggest, simply because different perceivers may be sensitive to different things. That point is thoroughly anti-Berkelean in spirit. The same wine can taste both leathery to one and metallic to another, not because these qualities are unreal, or are constituted by the perception itself, but simply because distinct qualities of a wine need not show up together in any single perception of it—something oenophiles will readily attest.

When two people are looking at the same painting, the situation is not significantly different. Some aspects of the painting may appear to the one. Yet the other fails to see them. We may reasonably suppose that the difference between such viewers is due to a possibly subtle difference in what aspects of the painting they perceive. As Hume's example brings out nicely, even people in similar circumstances can perceive different things. Variance in focus or sensibility may be the culprits. And so, when it is manifest to Hylas that Caesar is made visible by the portrait, someone else may not be in a position to appreciate this. This is not because Caesar fails to be suggested to the other person's mind, as Berkeley thought, but simply because they fail to perceive those aspects of the portrait that represent the emperor to

sight.

Someone may still wonder what made it so that only Hylas was able to see Caesar by means of his image in the envisaged situation. Similarly, why does the one critic pick up on the leather, while the other tastes only the metal? It is unlikely that there is a single answer to these questions. Hume thinks a kind of perceptual refinement he calls 'Taste' can explain why someone overlooks what is glaringly obvious to someone else. Staying closer to Berkeley, we may suppose that Hylas' knowledge of Caesar made the difference. Without such knowledge, Berkeley seems to think, Hylas would not have been able to have that mediate perception of Caesar.

Let us suppose that his prior knowledge of Caesar put Hylas in the favourable position he was in. In what way can such a cognitive, non-perceptual state make the difference? My answer is that it puts a viewer in a position to see aspects of the world that were there to be seen all along. Some cognitive states may be necessary to look at one's surroundings in the right way, but they are not themselves part of what it is to see some absent scene or figure by means of an image.

Some recent philosophers take a different stance. They rely on the assumption that some non-sensory, cognitive states can be *generative* of some of the qualities of which perceivers become aware. If such a cognitive effect occurs, it does not render visible some aspect of the world that was there to be seen all along. Instead, they think, it generates an experience of certain objects or qualities, *ex nihilo* (cf. Hopkins, 2012, p.653). This opens up a response to Berkeley that differs from mine. Recall, Berkeley thought that coming to see the painting before you as an image of Caesar does not bring about any change in how the painting appears to you, but instead depends on an act of suggestion. Instead of rejecting that claim because it does not fit well with the partial character of perception (as I suggest), one could also assume that acts of suggestion themselves are able to affect appearances, be it in some non-sensory way. Some philosophers hold

just such a view about what it is to see something as an image. I want to distance myself from this strategy. Not only is what these philosophers assume disputable, it is also redundant as an explanation, and incompatible with Hylas' central claim: that Caesar's appearance by means of his portrait is genuinely sensory.

The assumption that cognitive states can generate an experience of certain objects or qualities is a general one about the relation between perception and cognition. It is a claim that has become popular in recent years, mainly in light of findings in experimental psychology. Fiona MacPherson mentions an experiment by John L. Delk and Samuel Fillenbaum as a candidate for being a case where subjects' cognitive states can affect the qualities they experience (Delk & Fillenbaum, 1965). The experimenters showed participants orange cut-outs of various objects. Some of these were shapes of characteristically red things: such as a heart, a pair of lips, or an apple. They placed the shapes in front of a coloured background that the participants could alter. Then they asked participants to adjust the background to match the colour of the orange cut-out, so that the shape became least distinguishable from its background. Delk and Fillenbaum observed that for the shapes of characteristically red objects participants selected a background colour that was significantly redder than the background colour they selected for shapes of objects that were not characteristically red. MacPherson thinks that studies like these suggest that cognitive states or dispositions can 'penetrate' our perceptual state. As she maintains,

Although Delk and Fillenbaum don't themselves suggest a mechanism that explains why this happens, one might think that what is happening is that the subjects' beliefs, that certain of the cutout shapes were shapes of objects that were characteristically red, penetrated their perceptual experience of those cutout shapes thereby altering the content and phenomenal character of those experiences.

(MacPherson, 2012, p.39)

The way some philosophers have described what it is to see something as an image presupposes this general claim, that cognition can alter our perception. They assume that when Hylas sees the portrait as an image of Caesar, his perception is not a purely sensory achievement. Instead they think of it as a non-sensory effect. Hylas' knowledge of Julius Caesar penetrates his perception of the portrait, causing a non-sensory change in the way it appears.

Perhaps the most influential proponent of this idea is Richard Wollheim. As he writes, "thought, conceptual thought, can bring about changes in what we see in a surface" (Wollheim, 1998, p.224). Wollheim holds that when we see Caesar in a picture, we always also have some thought about Caesar (Wollheim, 2001). And just as MacPherson thinks we must identify cognitive effects to make sense of the colour experiments, Wollheim assumes that to make sense of mediate perception we must assume that certain non-perceptual states 'penetrate' or, as he calls it, 'permeate' our visual experience. In other words, not only does the surface of the image before us determine how things appear to us when we see it as an image, our perception in such a case is always in part determined non-sensorily, by our thoughts or knowledge of the things the image represents.

Dominic Lopes has recently taken a similar stance. He maintains that our seeing something as an image is enabled by our seeing a specific range of the image's visible qualities. He thinks that our perception of these qualities causes us to have a visual experience of what the image represents. Lopes defines the relevant qualities as those that comprise the image's 'design'. He thinks that a picture's design constitutes its representational aspect. As he explains,

A picture is a two-dimensional surface that depicts a scene in virtue of the way its surface is marked and coloured. Use 'design' to refer to those visible surface properties in

virtue of which a picture depicts what it does. Design comprises the surface configurations that you see when you see the picture surface without seeing anything in it and that are responsible for your seeing something in it. Not every intrinsic visible property of a picture surface is part of its design, however. We may be able to see that a picture is made of canvas or is very old, but if these are not features in virtue of which the picture depicts what it does, then they are not elements of its design. (Lopes, 2005, p.25)

In elucidating his concept of design, Lopes exploits a situation that, at least in essentials, corresponds to the one Berkeley uses in the *Dialogues*. Lopes uses the dalmatian image I discussed earlier as an example. He maintains that if someone looks at the image of the dalmatian without seeing it as an image, they may nonetheless get a full view of its visible qualities. In other words, they may see the image in just as perfect a degree as someone who does see it as an image of a Dalmatian. Seeing the image as an image of a dalmation does not make a difference to the qualities of the image you see. Instead, Lopes suggests, the difference would lie solely in the effects non-sensory states have on your perception of those qualities.

Wollheim and Lopes adopt a neo-Berkelean view. What they claim deviates only slightly from what Berkeley thinks about mediate perception. If seeing a surface before you as an image of something or other is an effect of cognitive penetration, then one denies that seeing the surface in that way is a purely sensory act. This implies, just as Berkeley wanted to claim, that Hylas erred when he advanced his theory of visual representation. The slight difference is in the details. Where Berkeley assumes that the cognitive act of suggestion does not affect how the image before our eyes appears, the neo-Berkelean view now makes room for this. It assumes that when we see the design of the painting it has an experiential effect on us (an effect not unlike the 'phosphene vision' brought about by electrically stimulating the

visual cortex. Brindley & Lewin, 1968; cf. Hopkins, 1998, p.653). It causes a cognitive act in us that influences non-sensorily how the painting appears. In other words, yes, something changes in the appearance, but no, this is not because we get to *see* more.

The neo-Berkelean view tells us what it is to see something as an image. Yet it does so in a way that is much more controversial than the earlier claim about partiality. It is controversial both in light of its general commitments about perception and cognition, and in light of its specific claims about painting.

In light of its general commitments about perception and cognition, the neo-Berkelean assumes that some variation in appearance is non-sensory. Yet many philosophers—including Berkeley—assume that any variation in appearance must be sensory. Sensory variation comes about when what we perceive itself undergoes some change, or when we get to see different aspects of what we perceive. If there is non-sensory variation, as the neo-Berkelean supposes, this would mean that a change in appearance could come about because we become aware of some non-sensory qualities of which we were unaware beforehand (M.G.F. Martin identifies this as a crucial point of disagreement between analytic ‘realists’ and philosophers working in the phenomenological tradition, see Martin, 2003, pp.523–24). It is unclear to me what such a non-sensory quality could be. Moreover, as Pryor (2000) has brought out well, the empirical evidence does not support the assumption that cognitive penetration in fact occurs. Even authors more sympathetic to the neo-Berkelean view admit this. Dustin Stokes writes that “there is no conclusive evidence for instances of cognitive penetration” (Stokes, 2014, p.6). And even MacPherson admits that the evidence there is does not lend support to cover cases of image seeing (MacPherson, 2012, p.36).

The hypothesis that only the influence of our cognitive states makes it possible for us to see images as such, makes our engagement with images more mysterious than it seems to be. Could what visitors get

to see in the National Gallery in point of fact be no more than a mere product of their thoughts or beliefs? John Hyman suspects the hypothesis to be no more than an expression of the ancient sentiment that images have something ‘magical’—that they cause effects in us that we cannot solely explain in terms of the visible objects they are (Hyman, 2006, p.145). As he characterises this sentiment,

Human beings are blessed with the capacity for unusual visual experiences that have not been caused by what they are experiences of, and a picture will generally cause a sighted human being to have an experience of just this sort. The task of the painter, on this view, is to produce something that will cause in a spectator a particular mental state of this unusual sort (Hyman, 1989, p.21).

This characterisation captures well the way Wollheim and Lopes conceive of our perception of images. It brings out that their view leaves too much unexplained. Even if one grants that perception as such is cognitively penetrable, it just remains unclear why we should think that seeing something as an image would not be possible without it. The neo-Berkelean hypothesis simply makes too big a leap from the available experimental findings to our everyday engagement with looking at and learning from images.

What matters most, however, is that the neo-Berkelean proposal is distinct from Hylas’ theory of visual representation. Hylas maintained that mediate perception is a purely sensory mode of awareness. Yet that is precisely what the neo-Berkelean view denies. It is at the core of the neo-Berkelean view that to perceive something by means of an image requires, apart from a sensory appearance of the painting itself, a cognitive state that penetrates or permeates our perception. It is because of this further cognitive effect that the way an image appears to us changes. This change is not sensory—it is not because we get more into view. For this reason, the neo-Berkelean follows Berke-

ley in thinking that mediate perception is sensory only in part. Although it undermines Berkeley's argument, the neo-Berkelean view still holds on to Berkeley's own conclusion: mediate perception is not a purely sensory mode of awareness.

4.5 Difference and identity

The key insight Hylas hoped to convey was that mediate perception is a purely sensory mode of awareness. Berkeley's response to this, by word of Philonous, is that this is not possible. Berkeley agrees that mediate perception indeed involves *an* act of sensory awareness. Yet he claims that the difference between merely seeing an image and seeing something by means of that image could not be sensory. All that can appear to sight when you see Caesar's portrait as an image of him is an arrangement of colours that could have appeared in just that way to someone who did not see the image in the way you do.

We can capture the disagreement in terms of the following two claims.

1. Mediate perception is a purely sensory mode of awareness
2. Variation in mediate perception entails variation in appearance

Hylas accepts both of these claims, whereas Berkeley rejects them.

I have defended Hylas' theory of visual representation. I explained how, even when we present two viewers with exactly the same portrait, in exactly the same conditions, there is still room for a purely sensory difference between what the one gets into view and what the other sees. I emphasised that this defence of Hylas' theory of visual representation differs from the view taken by several authors. In a neo-Berkelean spirit, they accept Hylas' conviction that (2) variation in mediate perception entails variation in appearance, but at the same time side with Berkeley. They side with Berkeley in denying (1) that mediate perception is a purely sensory mode of awareness. According

to them, the variation in appearance entailed by variation in mediate perception is due to 'cognitive penetration' or perception's 'permeability to thought'.

It is insightful to frame all this in terms of Hume's appropriation of Cervantes. After they have tasted the same wine, one critic claims it tastes of leather. The other denies this and only finds a taste of metal. In pinpointing their disagreement, Hylas is clearly on Hume's side. He thinks that what the critics report is just something that is there to be perceived in the wine. Berkeley seems to side with the sceptical guests, those who assume that the wine must taste just the same to all of them. He thinks the critics are merely bluffing when they say it tastes differently; they merely confuse what they think about the wine with what they perceive. The neo-Berkeleyan view maintains, contra Berkeley, that the wine *does* taste differently, at least to each critic individually. Yet, contra Hylas, the neo-Berkeleyan claims that the difference in the way the wine tastes to them is not due to what they perceive, but merely due to the way their judgements about the wine colour their perception of it. Their snobbery rests on the way their own cognition constitutes the distinct qualities they pigheadedly attribute to the wine.

The wine of course did taste both of leather and of metal, given that strapped key down the barrel. Once we grant the partial character of perception, it is quite obvious to see the attraction of Hylas' solution. Of course, Hume's example seems specifically designed to bring out how compelling this solution to the puzzle is. It is likely that neo-Berkeleyans such as Wollheim and Lopes would not wish to extend their analysis of how we see images to how we taste wine. But briefly pretending that they do offer such an analysis does prove insightful. It suggests that there is no reason to think that the difference between merely seeing an image and seeing something by means of that image could not be sensory.

Let me address a potential objection that someone may bring in at

this point. Some might be concerned that my defence of Hylas' theory falls short in one crucial respect. Although what I have claimed indeed explains that the difference between Hylas and his peer can be genuinely sensory, or that the shift that occurs when the represented dalmatian suddenly dawns on us is no more than a sensory change, all it leaves us with is the situation in which a viewer of an image comes to see more of *the image*. Indeed, it must be a change in which aspects of the image come into view if it is to count as sensory at all. Yet, even if we manage to get ourselves into that favourable position, we are still faced with a question of how this aspect of the painting we see relates to the material object it represents. In other words, if, as Hylas maintained, mediate perception of Julius Caesar is to be genuinely sensory, then a viewer of his portrait should also be able to come to see more of him, *Julius Caesar*. Yet, the objection presses, just by pointing to perception's partial character, all I have accounted for is the way some potentially hard-to-see aspect of Caesar's portrait can come into view (this is not unlike an objection John Searle raises in discussing sense datum theories of vision, Searle, 1983, pp.38, 58–61).

This criticism exploits an assumption we need not accept. It assumes that coming to see those aspects of the painting that represent Caesar and coming to see Caesar are distinct events. Perhaps they are events that occur side by side, or simultaneously, but the objection presupposes that they are not identical. Only on that assumption does it follow that coming to see specific aspects of an image of Caesar can never by itself amount to coming to see more of Caesar. To me it seems not at all obvious that we need to accept that assumption. Why could our perception of Caesar by means of the image not simply be identical to our perception of just those visible aspects of the image that represent him?

Brian O'Shaughnessy makes a suggestion about what could motivate the denial of such a straightforward identity claim. He explains how a tendency to overgeneralise our intuitions about specific forms

of perception may motivate the denial. He writes that there proves to be a persistent imagery at work in our minds, something that leads us to think of vision as acts of grasping or touching. Sometimes it may feel to us “rather as if the gaze literally reached out and touched its object”. He writes that

to the extent that such imagery or such a way of thinking is at work in one’s mind, to that extent one may come to consider the relation between mind and object to be such that perception of a material object could not conceivably be identical to perception of a distinct and second item [...]. (O’Shaughnessy, 2003, p.183)

To be sure, we cannot rightly describe the event of placing my hand on the rim of a glass as an event of placing my hand on something wholly distinct from that glass. To that extent we conceive of touching something in the way we conceive of biting something—I cannot bite one thing by biting something wholly distinct from it. Yet although this may be a truth about the tactile world, vision clearly works differently (see also my earlier discussion in chapter 3, p. 106). We have no reason to think that the constraints that apply to touch carry over to distal senses such as vision (or, for that matter, audition; Pasnau, 1999, pp.317–18 makes the mistake of expecting audition to conform to the constraints on touch).

If this is right, then the objection to Hylas’ theory of visual representation relies on a badly motivated assumption. Especially if we consider what may be required to get that specific aspect of the image into view—directing one’s gaze in a particular way to an absent scene only visible by means of the image—it seems more attractive to identify the two perceptual events. That is, we should prefer to consider the achievement of getting Caesar into view just to be the sensory achievement of seeing that aspect of the painting. (This makes my explanation ‘unitary’ as opposed to ‘divisive’, in Hopkins’ terms. Hopkins,

2010, p.170) This suggests that there is no obstacle to thinking of the achievement of getting Caesar into view as a sensory one, and to conceive of our mediate perception of Caesar as a purely sensory act of awareness.

Here we have one and the same event that falls under several descriptions. We can describe it correctly as “seeing those aspects of the image that visually represent the emperor” or “seeing the portrait as an image of Julius Caesar”. But we can also describe it rightly as “seeing Julius Caesar”, or “seeing the emperor by means of his image” (see also O’Shaughnessy, 2003, p.174). Here some descriptions will be of what is perceived immediately, and as such will make reference to some of the items or visible qualities present to the perceiver at the time. The descriptions that mention Caesar as an object of perception, however, characterise the perception as the mediate perception it is. As descriptions of a mediated form of perceiving, they are not constrained to items or visible qualities present to the perceiver. Instead, they will be of those items or qualities the perceiver manages to look at or otherwise get into view by means of the image. Such descriptions will specify perceived items or visible qualities that are merely represented by the image (here we have a limited application of Anscombe’s distinction between ‘material’ and ‘intentional’ descriptions of perception, Anscombe, 1965).

The earlier objection claimed that just by pointing to perception’s partial character, all we have accounted for is the way some potentially hard-to-see aspect of Caesar’s portrait can come into view. We now see how this is a mistake. The event of someone’s seeing those aspects of the portrait is at one and the same time the event of their seeing Caesar by means of the image. To my mind, this constitutes a completely adequate answer to the objection. It brings out how the achievement of getting Caesar into view is a genuinely sensory one, and it allows mediate perception to be constituted just by a sensory appearance.

4.6 Discovering an image

What we think, expect and know can help us uncover aspects of the world. It can also help us uncover images. Just as some bits of the world can be accessible only to those in the know, some bits of a painting can be too. Consider Brueghel's *Landscape with the Fall of Icarus*. Arthur Danto observes that one can easily overlook the central subject of the work. Icarus is only visible because we can see his legs, occupying only a small part of the painting. The rest of the work may overwhelm us, and may draw attention away from this key to what Brueghel wanted to show us (actually, the distraction may very well be part of what he wanted to show us). When someone points out, 'Look, that is Icarus!', one's perception of the work changes. As Danto writes,

To see the painting in these terms, if one had not seen it before, works to transform the entire composition, to pull it into a different shape... (Danto, 1981, p.119)

Danto's observation is put evocatively. Of course there is no actual transformation in the composition of the work. Brueghel's painting stays just the way it is, no matter how long you stare at it. What changes is what aspects of his composition we get into view. Getting Icarus in view requires us to look at the painting in a certain way.

How do we even begin to look at paintings in the right way? Richard Wollheim once mentioned that sometimes, the best way of finding out what an image represents is to look at its label. What Wollheim means is that looking at the label instructs us in how to look at the painting. It tells us what to look out for, and how to uncover details we would otherwise miss. A label can help us in the way Danto highlights, by telling us what to look for in a represented scene. A label can also help us in coming to see something as an image in the first place, as when you manage to see the dalmatian dog only after having been told that it is an image of such an animal. In both cases, our

knowledge transforms what we get to see simply because it makes us more sensitive to what was there to be seen all along.

Hylas was right: also mediate perception is genuinely sensory. In this chapter I have defended Hylas' position. I argued that the achievement of getting Caesar into view is a purely sensory one. It is an achievement that may well require perceptual activity on our part; it may well require us to look more specifically at what the image makes visible.

This conclusion is continuous with what we can accept for everyday cases. It fits the general fact that in looking at the world we are not entirely passive recipients, but also active explorers. In perception we aim to uncover ever more details of the visible world. In an apt analogy, F.H. Bradley likens our activity of attentive looking to that of a microscope, writing that it is "as if, so to speak, my will had served as a microscope, as if I were turning the screw and the detail were coming out" (Bradley, 1887, p.379). Here Bradley talks about perception in general, but his claim aptly characterises in particular our engagement with images. Cézanne's paintings confirm this. They successfully thematise just this activity. More strongly, deliberate attempts to draw out our activity of looking are found widely in the visual arts. T.J. Clarke writes about how Poussin's way of painting invites active exploration of elements,

astonishing things happen if one gives oneself over to the process of seeing again and again: aspect after aspect of the picture seems to surface, what is salient and what incidental alter bewilderingly from day to day, the larger order of the depiction breaks up, recrystallizes, fragments again, persists like an afterimage. (Clark, 2008, p.5)

Mediate perception exploits the active dimension of vision. However, we need not assume that we can only come to see Caesar by looking for Caesar in particular. Our active visual attention to a represen-

ted scene may very well be drawn out by the details of the image itself. We often direct our eyes towards something merely because we feel pulled in that direction, are made curious by something unresolved, conflicting, or unclear in what we see around us. Collingwood has aptly characterised the experience of such 'pre-attentive' awareness of something in our field of view as having an almost vibratory quality, inviting a more attentive gaze to resolve the tension (Collingwood, 1992, p.23; the point is emphasised by Olivi as well, see Adriaenssen, 2011, p.335).

This makes the microscope analogy doubly apt. Looking through a microscope, by adjusting the instrument's focus we not only can get into view what we set out to see, but also catch glimpses of details we had not anticipated. One gets a sense that there's something there, but one is not yet able to look at it specifically. In the same way, the visible detail of a painting may offer us cues to adjust what we look at by means of it. This activity of looking, adjusting, looking again, can go on indefinitely. Clarke captures it aptly:

Painting is making a world materialize. There are no rules, obviously, about how much or how little detail is needed for the materialization to occur. In any case, the process is often not about how much or little might be necessary, but about how long the process goes on being pleasurable to the person doing it. (Clark, 2008, pp.219–220?)

Some paintings allow for this type of peeling, in which more and more of their representational qualities can be uncovered. We have no reason to expect an easy formula here. Especially when imagery is produced by people with the purpose of informing, entertaining, or in an attempt at psychological reparation, there is no saying in advance how particular viewings will work out. What matters for the current discussion is that, while multi-faceted, this process can in principle be a genuinely sensory affair.

5

Structures of appearing

“*What do you think you’re doing?*” In 1962, an American advertising agent, Martin K. Speckter, designed a punctuation mark that allowed copywriters to convey surprised, rhetorical questions like this one more clearly. The visual symbol, which he ended up naming the *interrobang*, superimposes a question mark and an exclamation point (Fig. 14). In part because Speckter’s proposed typographical convention has remained a rarity, to the average reader the occurrence of an interrobang on a page or screen will stand out for its distinctive appearance. Although we’re familiar with how its component elements look—i.e. the question mark and the exclamation point—it is the way these elements are constructed to form a single glyph that is characteristic of the way the mark appears in a text (Houston, 2013, chap.2).



Figure 14: Interrobang, *Palatino* typeface

When we encounter the interrobang in a text, it appears visually. As a type of punctuation, the mark has no specific phoneme associated with it, and is to be discerned by the eye alone. Hence its characteristic look is essential to it. Its sensible appearance makes its identity visible to us. What is distinctive of that appearance?

A minimal account of sensible appearances or looks identifies the look of a thing with its manifest visible properties, such as shape and colour (such an account is neutral between a variety of ways of thinking about perception, see Martin, 2010). On such a minimal account, we may characterise the interrobang's distinctive look in terms of its more basic visible properties. Those properties jointly contribute to determining the way the symbol appears. At least typically, the interrobang will be printed in the same colour as the surrounding text. Therefore, just as with other typographical characters, the mark is identified by its visual shape. The symbol looks distinctive, because its shape is distinctive. We can define that shape in various ways, for instance in terms of other shapes (as the union of the shape of the exclamation point and the shape of the question mark) or in terms of more basic properties (as a collection of colour points in a coordinate system).

Here we have a simple exercise in characterising how something visually appears to us. It exploits the fact that how things appear to us in vision is structured in a specific way. It is natural to conceive of an item's look as a complex visible quality; a visible aspect that is constituted by more basic visible properties, such as shape and distribution of colour. This suggests that *analysis* is one viable way of characterising how something appears. First we identify the more basic qualities that constitute how an item looks when we see it. Having identified these more basic qualities, we can in turn characterise them in terms of the way they look. Of each quality, we can ask how it relates to the looks of other basic items that we could see. These relations among component qualities enable us to explain what is dis-

tinctive about the way a particular object of vision appears, and in what respects it looks similar to other visible objects. This analytic approach can help us identify how the interrobang's visible appearance is distinctive, despite the fact that in many respects it resembles the appearances of other typographical elements in a surrounding text.

Perhaps the best worked-out development of this simple idea can be found in Nelson Goodman's *The Structure of Appearance* (Goodman, 1977, a revised version of his Harvard doctoral dissertation, *A Study of Qualities*). As Carl Hempel characterises Goodman's project in that work,

The specific task Goodman sets himself is to construct the beginnings of a theoretical system which will permit an economic description of the structure of our phenomenal experience in the various sense realms. (Hempel, 1953, p.108)

Goodman develops a bold framework for analysing how things can appear to us—whether it is in vision, audition or some other sensory modality. He captures the sensible appearance of objects of perception at any given time as being built up out of discrete concrete phenomenal individuals. At least for vision, each of the smallest parts of the item we see comprises a chromatic quality, a spatial quality, and a temporal quality; something we see may look blue, here and now. Goodman's starting point is extremely general, taking a simple presentation of a uniform colour patch in our field of view at some point as his working example. His exposition is programmatic rather than exhaustive, and Goodman stays largely within the bounds of analysing very simple visible appearances in this way. Yet it is powerful as a move to identify a structure of appearance. Starting with a limited vocabulary, Goodman constructs a system that allows us to describe the spatio-temporal character of what we see, the ordered way things

appear in our field of view, similarities and differences between colours and shapes, and so forth (cf. Hempel, 1953, p.109).

Goodman analyses the looks of visible object as complexes of more elementary qualities. Only these more elementary qualities are, on his approach, characterised in comparison. But what is the nature of these simpler sensible qualities at the heart of Goodman's analytical approach? That remains a moot point. Goodman himself takes them to be qualities of our sensory experience, and proposes to capture the world around us in terms of how such qualities appear to us (Goodman, 1977, p.101ff). This pushes a strong phenomenalist approach to the relation between appearance and reality: it takes our sensory awareness of our own experience as given, and analyses our ordinary concepts of the empirical world in terms of that awareness. However, Goodman emphasises that his choice for a phenomenalist approach is purely methodological. (It allows for a constructional system that overcomes some of the technical difficulties that rival approaches face, such as the system of Rudolf Carnap (1928).) Phenomenalism is not essential to the project of analysing the structure of appearance as such. Alex Byrne and David Hilbert carry out a similar kind of analytic project when they explain the structure of the appearance of colours and tastes. About the simple qualities they are interested in they emphasise that

there is nothing mental about the qualities themselves, or the items that possess them. If colors and tastes are properties of anything at all, they are properties of familiar non-mental objects like lemons and drops of lemon juice. (Byrne & Hilbert, 2008, p.387)

Byrne and Hilbert investigate the structure of appearance along the same lines as Goodman, yet assume the elementary sensible qualities of how things appear to be firmly rooted in the sensible world around us. I take this to be an attractive default approach, that should be

followed unless good reason suggests otherwise. As here I see no such reason, I will follow Byrne and Hilbert in conceiving of elementary visible qualities as sensible qualities instantiated in our environment.

The visible appearance of an object is a complex of its more elementary visible qualities. Goodman identifies a structure to how things appear, by analysing looks of visible objects in terms of complexes of more elementary qualities. At this more basic level, he thinks, we can identify what distinguishes specific visual appearances. Although Goodman only discusses the most simple cases, he seems confident that his analytic approach can at least in principle extend to a complete characterisation of the way any object appears in any situation. I want to remain neutral about this. What I want to highlight, however, is that the way of analysis is but one way of approaching the structure of appearance. It is not the only way. At least for one specific way visual appearances are structured, the analytic method may very well be quite useless. In the next section I will show how the conclusions about images I reached in the previous chapters allow us to uncover something of the structure of appearance that is captured by neither Goodman's nor Byrne and Hilbert's version of the analytical approach.

5.1 Appearance and representation

Think of the participants in a life drawing class. They now study the model as it is sitting before them in the same room, then they look at the model as it appears to them by means of the image they are creating. This situation allows the artists to exploit a crucial fact about vision, namely that we can see one and the same item both immediately and mediately. The point is brought out vividly in Carl Mydans' photograph of a group of art students in Texas (Fig. 15). By looking at the model now in the one way, then in the other, the artists can tweak and adjust the images they are painting in such a way that they

represent realistically—roughly, such that the way the model looks in real life matches the way it looks in their pictures.



Figure 15: Carl Mydans, *Art students at the painting class*, 1939. Texas

Life drawing of this kind is possible and worthwhile because vision is sensorily heterogenous, as I defined it in chapter 3 (p. 133). There I suggested that we may think of a sensory modality as heterogenous if and only if in that modality we can perceive one and the same item both immediately and mediately. I have shown how this applies to vision. Our capacity to see our surroundings immediately can be supplemented by images that function as mediators in visual perception. This explains vision's heterogeneity. At first, Ishmael in *Moby Dick* has only seen whales by means of images, such as the dramatic painting hung in his lodgings in Nantucket. Only after he has sailed out on a whaling ship does he first see a whale immediately. As soon as we have an image of something which also exists out there in the world, we may be able to see that same thing now immediately, then mediately, or even, as in a life drawing class, see the thing in both of these ways, more or less at the same time.

Vision's heterogeneity implies a further distinction. It implies a distinction between two kinds of visible aspects that may be actually instantiated by items present in our surroundings. Let me present it by means of an example.

In the room Mydans photographed, the model sitting before the artists looks a certain way. She looks pale, slightly distracted, independent. We may conceive of how she looks from our own point of view—of her sensible appearance as seen from here—as a complex visible quality of that model. Compared to our description of the look or sensible appearance of the interrobang, a full characterisation of the look of the model would be a monumental undertaking. I doubt if anyone could ever give a satisfactory, informative characterisation of how the model looks just in terms of the visible appearance of the more elementary visible qualities she possesses. Pinpointing what is distinctive of the interrobang's appearance is relatively easy. Identifying what makes the model's looks distinctive seems to require a volume of literature.

Practical challenges in fully characterising the model's look in terms of her more elementary visible qualities does not undermine a theoretical point. At a metaphysical level, we may identify the model's look with a complex of elementary visible properties she possesses, such as the shapes and colours of her body. Were we to change the model's look, then we would have to alter those more elementary qualities. For example, we could ask her to wear darker clothes next time, or invite her to fix her hair so that it falls differently. Which set of qualities is identical to the visible aspect of the model when we see her? Of course we always see the model from a specific point of view, if we see her at all. We may normally identify the way she looks to us with a specific range of such visible qualities—namely those qualities visible to us from where we stand. (We could also try to define the notion of an 'overall' look and leave out the perspective. I will not do this, but see Martin, 2010, p.202.)

Now consider the way the model looks if you see her by means of one of the images the artists are working on. This too is a possible way of looking at her. Yet here you would not see the model immediately, but only via her representation—a surface covered in paint, still wet, unfinished as a portrait. Nonetheless, what the cases share is that here too it seems natural to identify the model's looks with a complex of more elementary visible qualities. We can identify her look with the complex of those qualities visible to us when we look at the model by means of her image. In this case, how we see the model is not best captured in terms of our relative location to the model. Instead, it is best captured by saying that we see the model by means of her image. That is how we see her when we see her mediately. The difference in how we see the model has consequences for how, precisely, we specify her look. When we see the model mediately, the way she looks must be identified with a complex of visible qualities of the painted image that enables us to see her. Why is this so? We can see it from a viewpoint of intervention. To change the way the model looks when seen by means of the image, asking the person on the stage to reshape their pose will have no direct effect. Instead, we would have to ask the artist to make changes. The artist could add more blue to the central parts of the painting, or darken the section that shows the model's hair, for example. The way the model looks when seen mediately must be identified with a complex of visible qualities of the painted image.

I have identified two routes for specifying looks, corresponding to two ways in which items can be seen. Looks can concern the complexes of visible qualities of things in the world when we see them immediately. The other concerns complexes of visible qualities of images as perceptual mediators, when we see things mediately. These two routes for specifying looks allow for comparisons. You can compare the way some item looks when you see it by means of an image with how it looks (or would look) when seen immediately. Conversely, you can compare the way some item looks when you see it immediately with how such an item looks (or would look) when seen

by means of an image. In both cases, you would be drawing comparisons between distinct visible aspects of what is and could be before your eyes. Such investigation allows you to spell out how different sensory appearances of the same thing relate to one another.

Comparisons like the ones sketched above are a key procedure of the representational arts. As such they are aesthetically significant. In a life drawing class the students carefully study the visible aspects of the model they see in order to improve how the visible qualities of their paintings make her look. In doing so, they precisely draw on the kinds of comparisons made possible by these structures of visual appearance.

The possibility of comparing the way items look when seen immediately with their looks when seen mediately is also important for the philosophy of perception. It helps establishing systematically how different perceptual encounters with the visible world relate to one another. It allows us to reveal something about the structure of visual appearance. In particular, we can group the visible aspects of items in our surroundings into two kinds. On the one hand we find those visible aspects that belong to the objects that instantiate them, and which we typically see as just present in front of our eyes. On the other hand we find visible aspects that are instantiated by objects present to us, but in a way belong to some other object or scene that need not itself be present. These visible aspects belong to an object or scene that we can come to see by means of the items that are present to us. How are these visible aspects individuated? They are individuated by the figures or scenes they represent (cf. Johnston, 2007, p.247). Such a basis of individuation is naturally attractive, because these aspects of objects in our environment are representational: they enable us to see some absent figure or scene by means of them. When we see such visible aspects in our surroundings, we see the item that instantiates them as an image of something else. (M.G.F. Martin has recently argued for a related point, but he neglects the role visual mediation

plays in how we identify and grasp the representational qualities of our surroundings. See Martin, 2012.)

The model looks pale, in the way one of Mydans' painters represents her. Her represented paleness is a visible aspect instantiated by the paint on their canvas. It exemplifies a visible aspect of the representational kind. Not only do these visible qualities determine that the painting has a representational character, they also determine how the painting represents. For example, the visible aspects that make the painted canvas a visual mediator at the same time determine that the model looks pale when one sees her by means of the painted image. For an item to represent in that way, it suffices to have those visible qualities.

I already mentioned how elementary visible qualities are commonly organised into families of colour, timbre, and taste. Now we see that also highly complex visible aspects of the world around us effectively fall into two categorially distinct kinds: those that determine the way things look when seen immediately, and those that determine the way items (images) represent some absent scene or other. These are two kinds of complexes of qualities. The two kinds of complexes of qualities can, and frequently do, coexist in our visible surroundings. This is a structural fact about our visual world.

There are methodological lessons to learn here. I have shown how visible aspects can come in two kinds, and how this is a fact about the structure of appearance. But note that this fact is unlikely to show up on an approach that emphasises the way elementary visible qualities are structured among themselves. It is an assumption of Goodman's analytic approach to visual appearances that the nature of more complex qualities can be explained with reference to the way the more elementary ones appear. Yet the two kinds of visible aspects instantiated by items present in our surroundings we have just identified and characterised showed up precisely because I followed a different method. My main strategy has not been to analyse the way something

appears in terms of more elementary qualities, but to compare how some item looks when one sees it immediately with how the same item looks (or could look) when seen by means of an image. I have taken a step back, concentrating not on how appearances are composed, but on the structural similarities and differences that show up in different types of engagements with more complex visible aspects. This suggests that the way representation figures in our visual world may only show up when we compare the visible qualities of things at a considerable remove from the minutiae.

Visual representation is firmly rooted in the visible world. This shows from the fact that the representational aspects visible on a painting are due to the painting's having the visible qualities it has. Suppose you press a seal onto a receptive surface, such as a slab of clay or wax. The seal itself may have an image of two wild animals engraved on it, and is thereby able to transfer that image onto the wax (Fig. 16). When you press it down, you in part alter the visible qualities of the wax, such that the slab itself now visually represents two wild animals. Part of the wax' surface has become representational. After removing the seal and looking at the wax you may discern the complex of visible qualities you just impressed on the wax and see it as an image of an oxen.



Figure 16: Reproduction of a late bronze age seal-impression

The visible qualities we see around us are instantiated by various

items in our environment. This means that the visual representation of two animals on a slab of wax, or of a group of bathers on Cézanne's painting, is as firmly placed in the world as the red colour of the tomato, or the blue hue of the sky. As I emphasised earlier, with Byrne and Hilbert, such qualities are in no way mysterious or mental. Now of course our conception of a quality, such as the colour red, may be inextricably tied up with the way such qualities appear to us in perception. But that does not mean such qualities are any less real—it only means that we conceive of them in a way that is tied up with how we perceive the world (Hilbert, 1987 has called such conception 'anthropocentric').

The conclusion to draw about the way representation figures in our visual world is entirely mundane. Just as the painting hangs in the gallery quite independent of my perception of it, the qualities that make it visually represent in the way it does, are out there too, regardless of whether I come around to see them or not.

5.2 The image's silence

Mark Johnston has characterised visual appearances in terms of how our visible surroundings make themselves manifest to us. When we open our eyes, we perceive things in our surroundings, as well as their distinctive visible qualities (Johnston considers such qualities objective 'modes of presentation' of the world, 2007, p.245). A tree in a meadow may be visible along with its colours and the crooked contour of its stem. Just as the tree is out there in the meadow quite independently of my perception of it, those visible qualities are out there too, regardless of whether I see them or not. The tree looks the way it does because of the visible qualities it has. For something to look the way the tree does, it suffices that it has just those visible qualities that the tree has.

In the previous section I outlined a structural fact about our visual

world. The figures and scenes we can see by means of images look the way they do, because the image mediating our perception instantiates specific visible qualities. The image represents by virtue of having these qualities. Moreover, these qualities at the same time determine how the things that the image represents look when we see them by means of it. This way, surfaces in our surroundings can be representational just because of their intrinsic qualities—they are visual representations purely by having the visual qualities they do. For example, Cézanne's *Les Grandes Baigneuses* represents several bathers that look a specific way when you see them by means of that painting. The painting represents the bathers, because the visible qualities it has enable it to fulfil a mediating function in our vision of the world. Therefore, for a painted surface to represent a group of bathers in the way *Les Grandes Baigneuses* does, it suffices to have those visible qualities.

My suggestion is controversial. It is common for philosophers to dismiss the idea that having specific visible qualities is sufficient for visual representation. Typically, authors writing about the nature of visual representation suppose that merely possessing a complex of visible qualities is never enough for a surface to count as a visual representation. They may accept that having some complex of visible qualities is necessary for representing a group of bathers as Cézanne's painting does. Yet they are adamant that *Les Grandes Baigneuses* only comes to represent a group of bathers because of some further condition that is met. Any other surface that just happened to possess those visible qualities would not thereby be representational. In other words, authors commonly assume that visual representation requires something other than the opportunities already available in the visual world. What could motivate this?

Here is a diagnosis. It seems to me that many philosophers are driven by the idea that something can only be a representation if it is produced in the right way. (I already touched on this in chapter 2,

p. 74.) This is regularly spelled out in terms of the ‘artificiality’ of images. Many simply take it for granted that every image is an artefact created or selected to represent something. Images are said essentially to be the product of human craft or intention, embedded in some communicative practice, and so forth. Early on in his book on visual representation, Robert Hopkins expresses this conviction as follows:

depiction is an artefact, necessarily the product of human contrivance. Unless our surface has been created, altered, or at least brought to prominence within a framework of people’s goals and actions, it cannot depict. Rather, it is at best indistinguishable from some depiction, i.e. from something very like it, but with the appropriate history. (Hopkins, 1998, p.71)

With ‘depiction’ Hopkins just means visual representation—the kind of representing distinctive of images. He takes it for granted that, unless a visible surface has been created in order to represent, or at least been introduced into a framework of goals and actions, it cannot count as a representation. Hopkins states this merely as a report of fact, and it is clear he thinks it does not require further argument.

Those who accept that images need to be produced in the right way diverge over exactly how an image needs to be produced to count as a visual representation. Richard Wollheim, for example, writes that an image represents x only if it “was intended by whoever made it to be a configuration in which something or other could be seen and furthermore one in which x could be seen” (Wollheim, 1977, p.717). Just as Hopkins, Wollheim ties representation to an artist’s intention. Similarly, Hanna Pitkin suggests that “pictures are representational insofar as the artist was alleging something about the visual appearance of what he depicted, about the way it looked or the way he imagined it to look” (Pitkin, 1967, pp.68–69). Other authors suggest that

representation is instead tied to non-intentional facts in the image's history. For example, Lopes claims that what is represented is determined by an image's causal relation to its object, via a series of preceding information states that have that object and its properties as their sources (see Lopes, 1996. A similar claim is made by the reproduction theory I discussed in chapter 2). *Les Grandes Baigneuses* is of a group of bathers, because of a causal-historical information chain from Cézanne's painting to a group of bathers.

The variation here need not concern us. What matters is that all these authors take it to be beyond doubt that an item does not represent anything *just* by looking the way it does. As Abell writes,

The pool of milk that forms when I accidentally drop the carton may be indistinguishable from a picture of a cat, but it does not depict a cat, because it was not produced in the right way. (Abell, 2010a, p.275)

Having the visible qualities Abell observes her milk to have may be necessary for representing a cat in some way, but Abell takes it as obviously insufficient for it. This assumption is typically taken on board without much discussion. It is as though visual representation simply *could not* be realised by what is visible in our environment alone (another vivid example is Bantinaki, 2007, p.277).

I see no reason to follow suit. The authors I mentioned all dismiss the idea that having specific visible qualities is sufficient for visual representation. Yet from what I have shown about how images represent, there is space to side-step this pressure. An image represents by making some object or scene visible in its absence. In doing this, images function as perceptual mediators. And images can function in perception as perceptual mediators because they possess a complex of visible qualities; a complex that may be identified with more basic visible elements such as form and colour. Considerations about how these qualities came about, or about how the image was produced,

do not seem to figure in this explanation of what it is for an image to represent. As soon as some visible surface in our surroundings instantiates such a complex of visible qualities, it is an image, and it visually represents something or other. Therefore, not only does visual representation necessarily exploit standard capacities of vision, it is a genuine part of our visual world.

To be sure, not just any set of visible qualities will do. A visible surface is only representational if its visible qualities in fact enable us to see something. A surface only represents if it can function in vision as a perceptual mediator. As you may try out for yourself, not just any visible surface enables you to see something by means of it. This means that having just any odd collection of visible qualities does not suffice for something to be an image. But it does allow that there are some ways a surface can look that suffice for it to make something wholly distinct from it visible, and therefore to represent.

My account of representation has a remarkable consequence. It implies that *what* an image represents is not determined by the visible qualities that make it a visual representation. What an image represents is left open, and will depend on what exactly a viewer finds themselves able to see in the image. Consider an example. A familiar observation is that the visible marks on the surface of a portrait of Amal may serve equally well as a representation of Amal's twin sibling, Bilal. Not only can we see Amal by means of the portrait, it makes Bilal visible as well. Things can go even further. Seen as a visual representation of Amal, the portrait makes visible a sun-tanned Amal, an Amal cast in shadow, and Amal wearing heavy make-up. It all depends on how you look at the image. Visible qualities of an image as such do not fix any particular way it can make an absent object or scene appear to sight. Which of the various manifestations of Amal we see, or whether we see Bilal rather than Amal, will depend on how we look at the portrait. Why is this so? The explanation is that many different things or situations can look the same, or at least look suf-

ficiently alike (cf. Travis, 2004). This suggests that all there is to say *a priori* about what an image represents is that it represents whatever viewers are able to see by means of it.

So it turns out that every image can, in principle, make a large number of different scenes visible. This is consistent with all I have established about visual representation. Images represent simply by making absent scenes visible, and we have seen how this is possible. This understanding of visual representation neither fixes nor restricts what an image makes visible (and hence represents). As long as it makes something wholly distinct from it visible, the image is a visual representation. About what exactly it makes visible, the image itself remains silent.

All this is no more than a consequence of the fact that to represent something visually just is to make it visible via perceptual mediation. But to many it will seem a remarkable claim. It contradicts a dominant way of thinking about the place of representation in the visible world. In the next two sections I will clear up two confusions that may motivate resistance to my conclusions about representation. In particular, I will discuss two kinds of norms that philosophers emphasise seem to be implicit in our engagement with images. Both of them have at times been taken to motivate the idea that merely having specific visible qualities is insufficient for visual representation.

5.3 The accretion of standards

As a pastime, Marcel Proust liked to visit the Louvre and try to confirm visually his amateur hypothesis about human typology—the idea that, in the end, there are only a finite number of types of people. Standing in front of paintings by past masters, Proust tried if he could see the figures of some of his own friends. In doing so, he seems to have managed to find at least some support for his hypothesis. By looking at a double portrait by Domenico Ghirlandaio (Fig. 17),

for example, Proust managed to see the left figure as the Marquis du Lau d'Ollemans, someone he knew from the Parisian suburb Saint-Germain.

Assuming that this was indeed what Proust saw, it seems right to suggest that he violates some kind of norm or standard of looking at pictures. As Proust knew full well, Ghirlandaio's double portrait is not a portrait of the Marquis du Lau. It portrays, as Richard Wollheim observes, "an old Italian prelate with a polyp at the end of his nose" (Wollheim, 2001, p.7). Proust, in some sense, was not supposed to see anyone other than the Italian man that Ghirlandaio portrayed. His behaviour in the Louvre was, at least to that extent, transgressive.

What an image such as Ghirlandaio's portrait visually represents depends on what objects or events it makes visible to a viewer. Yet, the fact that Proust so clearly set out to see the wrong thing may tempt us to think that an image does not represent just anything that a viewer is able to see by means of it. This is puzzling. On the one hand images are all about what we see, yet on the other it is obvious that Ghirlandaio only portrayed a high-placed fifteenth century Italian, and not one of Proust's Parisian friends. Apart from what the image Ghirlandaio painted makes visible, it seems, there must be some further norm that determines what the image actually represents.

Philosophers have crafted a default solution to this puzzle. Catharine Abell presents it clearly (Abell, 2010b). She explains how generally authors specify both the visible properties something must have in order to represent an object, and in addition define some *standard of correctness*—some norm that fixes whether something that has those visible properties in fact represents the object in question. Many philosophers suppose that such a norm reconciles the two conflicting intuitions that drive the above puzzle. On the one hand the view that what a painting represents depends on what object it makes visible, and on the other hand that a painting is not of everything that a viewer can see by means of it (Wollheim, 2001, pp.19–20).




Figure 17: Domenico Ghirlandaio, *An Old Man and his Grandson*, ca. 1490. Louvre, Paris

If this were the only way to resolve the puzzle, it would independently motivate the idea that merely having specific visible qualities is insufficient for visual representation; it would thereby threaten what I claimed above. Yet I can show that introducing a standard of correctness is not the only way to resolve the conflict between these intuitions. More importantly, it is not the best way either. To my mind, a much more natural response is to draw a distinction between on the one hand the images that make things visible, and on the other hand the ways in which these images can be embedded in specific,

historical, and indeed normative practices.

An image such as Ghirlandaio's portrait represents the figures it makes visible to a viewer. But this does not take away that the image was intended specifically to show us the Italian prelate that Ghirlandaio set out to portray. It is because Ghirlandaio's image was created as a portrait, that we are supposed to see this particular person when we look at the painting in the Louvre (the precise identity of the sitter is a matter of speculation, see Christiansen et al., 2001, p.159ff). Similarly, because we conceive of photography as a procedure that produces images of real particulars as they actually were, we correct the young child who insists they saw a photograph of Snow White. Of course we can acknowledge that images often come with rules or standards as to what we are supposed to see by means of them. But any such rules will attach to the image not because of its representational character—not because it makes something visible—but because of the way it gets embedded in, to use Hopkins' phrase, a framework of people's goals and actions. Proust transgressed the rules. But he transgressed the rules of portraiture, and not the rules of representation—for, if I am right, there are no such rules.

How images in fact are embedded in specific, normative practices is highly varied and flexible. The standards associated with portraiture are specific to that way of producing images, and do not coincide with other ways we exploit the possibilities of visual representation. Consider for instance the norms implicit in the practices of life drawing and diagnostic imagery.

Giorgio Morandi, the Italian painter and printmaker, is known for his still life works. Many of his sketches are drawn from life. *Still-Life* (Fig. 18), for example, is a life drawing of a small collection of rather mundane objects. In making the drawing, Morandi invites viewers of the image (which importantly include the artist himself, at a later time) to see the objects the way he saw them at the time of drawing, in all their plainness, scruffiness, emotionally tied up with a domestic




Figure 18: Giorgio Morandi, *Still Life*, 1962

life. His drawing is successful to the extent that we can indeed see them that way. One sees the drawing in the right way if one indeed sets out to get the object in view as it appeared to Morandi when he made his drawing.

Another case. In the first half of the twentieth century, the Swiss psychologist Hermann Rorschach gathered a by now very familiar collection of inkblots to be used as a projective test in therapeutic settings (Fig. 19). Initially developed for diagnosing patients suffering from schizophrenia, the set of ten Rorschach cards gradually gained popularity. An essential phase of administering the test involves free association on the part of the participant. Implicit in this first phase is a strong requirement to look at these blots and report what one sees straight away, without reflecting on the images. In its current form the test is successful because a wide variety of diagnostically significant things can be seen by means of the blots on the cards.

Compared to portraiture or photography, life drawings and diagnostic images come with a considerably different cluster of expect-



Figure 19: Card no. 1 in the Rorschach test, developed by Hermann Rorschach

ations and even requirements on a viewer. I do not claim that the particular expectations I attribute to Morandi are an essential part of the concept of drawing from life as such (though Phillips, 1992 defends this at length). However, I do maintain that the standards of life drawing frequently include that the artist draws what they saw as they saw it. Such a norm is foreign to portraiture, which commonly idealises. Again, I also do not suggest that the standards implied by Rorschach's inkblots are the only viable way to integrate images into psychotherapy. I do claim that the demand to name whatever one sees without reflecting on the image is unheard of in practices of architectural photography or medical illustration, say. What matters is the potential diversity of norms and standards implicit in our engagement with images that these examples illustrate (see Kulvicki, 2010 for a somewhat different take on such diversity).

Even internal to the practices of portraiture and photography themselves we find such diversity. Maria Loh has described how portraiture in earlier days conformed to different standards. Up to the early Renaissance, portraits were not expected to show how their sitter looked, but merely to provide for a way of securing their status, even when they were not actually around (Loh, 2009; the point is also observed by Hyman, 2006, p.247n10). Sir Francis Galton (who, incidentally, entertained a hypothesis about human

typology not unlike that of Proust) claimed to have harnessed the photographic image in such a way that it visualised types of people. Galton explicitly discourages his viewers to look at them in any other way (Galton, 1879).

Philosophers who think that a 'standard of correctness' is essential to visual representation in effect claim that such a standard is part of what it is to be an image. They assume a rather deep tie between visual representation and such a standard. But then it seems somewhat embarrassing to find that the artistic norms and standards that in fact attach to the images we find around us turn out to be a contingent mixed bag of historically determined expectations, prescriptions, and protocols. If such norms were part of what it is to be an image, this would suggest that visual representation at its very core is much more diverse than it in fact seems to be. To be sure, the kinds of images we find are extremely diverse, both in appearance and detail. Yet it seems clear that all represent in precisely the same way: by fulfilling a mediating function in our perception of the world.

The different ways we conceive of portraiture, photography or life drawing are in themselves philosophically relevant. Antonia Phillips has shown that they can reveal much about how we conceive of what we see, or how we think that images can offer us knowledge of the actual world (Phillips, 1992). But although each of these practices crucially depends on visual representation, the norms or standards associated with them do not cut as deep as some philosophers have assumed. In any case, the artistic or cultural norms associated with how we use images do not independently motivate the idea that visual representation demands more than merely having specific visible qualities.

5.4 Misperception

One further worry has spurred the idea that merely having specific visible qualities is insufficient for visual representation. This time it does not concern specific historical practices, but visual perception as such.

Consider a famous painting by Gerard ter Borch, in the Rijksmuseum in Amsterdam. To many eyes, at a first glance at least, Ter Borch's painting (Fig. 20) seems to show a friendly family scene. That is at least what Goethe took the painting to show. Goethe described the painting in one of his novels, and suggested this interpretation. A friendly family scene with a father kindly reprimanding his daughter for some small misdemeanour. More recent scholarship has pointed out that Goethe's interpretation cannot be right. It does not fit the genre categories of the time. The figure Goethe took to be the daughter was almost certainly intended to be a prostitute. The alleged father figure was intended by Ter Borch as a customer, negotiating the price. Instead of a friendly piece of family advice, Ter Borch set out to paint a brothel scene.

Unlike Proust, who merely transgressed an artistic norm when he looked at Ghirlandai's portrait with the aim of seeing some of the people he knew in Paris, Goethe made an error. The interpretation of the painting he published purported to get an historical fact right, yet did not get it right at all. Yet it is likely that Goethe's error was a compound one. To see this we may distinguish two kinds of error that may be involved when one is trying to interpret a painting, as Goethe did. The one kind of error is attributive, the other is perceptual.

In his description of the painting, Goethe did not capture what Ter Borch intended to represent. Take what Goethe wrote about the painting as an attribution of an intention or plan behind the painting. Goethe wanted to suggest that a friendly family scene was what Ter Borch set out to paint. At least if recent scholarship is sound, Goethe's claim here embodies an error we could call 'attributive'. What he said



Figure 20: Gerard ter Borch II, *Gallant Conversation*, also known as 'The Paternal Admonition', c. 1654. Rijksmuseum, Amsterdam

failed to capture the actual conditions under which the painting was produced.

Yet it is highly likely that Goethe's error in attribution stood not on its own. It was probably based on his perception of the painting. This at least is how Wollheim construes the case. As he writes,

When Goethe looked at the painting by ter Borch, what he thought he saw, as we know from his novel *Elective Affinities*, was a touching family scene in which the father, an elderly knight, reprimands his charming and beautiful daughter for some very minor infringement of manners which she has committed, while the mother looks on in

an embarrassed fashion. (Wollheim, 1996, p.31)

Typically, an artist will make sure to configure the visible qualities of her canvas in such a way that the potential viewers she envisages will see those and only those things she wants those viewers to see. For this reason, looking at what one can see most easily by means of a painting typically provides excellent evidence for the kind of interpretative attribution Goethe was seeking to make. Provided, that is, that one does not fall prey to any perceptual error and 'mis-sees' the image one looks at. And it is likely that this is precisely what happened to Goethe. Goethe thought that by means of Ter Borch's image he could see a touching family scene, yet to those familiar with the painting, there is good reason to resist that judgement.

Wollheim discusses Goethe's error to motivate the idea that merely looking a certain way, or having specific visible qualities, does not yet fix what a painting represents. Wollheim assumes that a further criterion must determine what the image actually represents. As one reason to accept such a further criterion, Wollheim notes that without it, the kind of error Goethe made would be unintelligible. As he writes,

to talk of mis-seeing a picture [...] presupposes that there is some norm of what we should see. If there is, what determines the norm? I believe that the norm comes from the intention of the artist. (Wollheim, 1996, p.31)

Now if Wollheim is right here, then that is serious. Earlier, I showed how some of the visible aspects instantiated by the paint on a canvas make that canvas a visual representation, and make it represent in the way it does. When we see such visible aspects in our surroundings, I concluded, we see the item that instantiates them as an image of something else. This implies that seeing something as an image presupposes that we have become visually aware of some complex of visible qualities present in our environment. Taking ourselves to see

something as an image implies that we take ourselves to see those qualities. And it is a very general fact about perception that what we take ourselves to perceive is always susceptible to error. I can be wrong in what I take myself to perceive. It may always happen that, unbeknownst to me, conditions are poor or the way things appear is misleading. In such a situation I may claim with confidence to see water in the distance, or to see a horse with only three legs, but I may always be wrong, depending on what I in fact saw. We cannot both consider visual representation to be a genuine part of our visual world and at the same time renounce the possibility of perceptual errors about what we can or cannot see by means of some visible qualities of our surroundings (Wollheim himself calls this a 'totally untenable form of relativism'. 1996, p.31). Wollheim suggests that accounting for the possibility of such errors presupposes that there is some norm of what we should see in the image. Hence, accounting for the possibility for error or misperception would offer independent reason to accept that there are standards about what an image represents. Were that to be accepted, then we would have to say that merely having specific visible qualities is insufficient for visual representation.

However, it is not at all clear to me that Wollheim is right. Why should we think that perceptual errors about representation are only possible if there is some norm as to what we should see when looking at an image? I see at least several non-trivial ways in which how we see things can lead us to false beliefs about what an image represents—ways that do not presuppose a norm or standard that fixes what an image represents independently of its visible qualities.

It may happen that someone is in an environment that does not contain any images at all. Yet due to their own condition, or bad lighting, they may nonetheless come to think they are able to see a face when they look at a stain on the wall, or that they can see a monster in a heap of clothes. Such cases are often associated with infancy or distinct pathologies. Uchiyama et al. (2012) point out that pa-

tients with dementia with Lewy bodies (DLB) may incorrectly take themselves to see a person in a curtain or perceive blobs on the wall as faces. Aristotle in *On Dreams* (460^b12) similarly notes that in the delirium of fever people sometimes think they see animals on their chamber walls. In such cases, it is often possible to clear up the confusion by taking a better look, turning on the light, rubbing ones eyes or, as Aristotle himself notes, by waiting until the delirium resides. Frequently, people will then come to acknowledge that they made a mistake, and that what they thought could be seen by means of those stains or pieces of clothing was in fact not visible at all.

Conversely, someone may be presented with an image but maintain that they merely see a surface with some marks. In the last chapter I discussed a by now famous image designed by psychologists precisely to get people in such a situation (Fig. 13, p. 154). Initially, most people are likely to think that the surface before them does not represent at all. Yet after they spot the dalmatian dog that is visible to the keen eye, they will readily acknowledge their initial error, and agree that the pattern of blobs before them made that dog visible all along.

These two examples of perceptual errors are uncontroversial cases where people make mistakes on the basis of how they see things. Yet neither of them presupposes that there is some norm of what we should see. Instead, what these errors presuppose is that being an image or not is an objective, visually verifiable feature of our surroundings, something we can potentially clear up by taking a better look.

Goethe of course did not make either of these mistakes. He could clearly see that he was presented with an image when he approached Ter Borch's painting. Indeed, there is no doubt that Goethe saw the image, and looked at it with care. The error he in spite of that made was rooted in his overlooking some aspect of the image. What he overlooked was that the man represented by the image holds a coin

between his fingers (Fig. 21). Even for viewers who are in a position to see the full painting as opposed to a reproduction, this requires close scrutiny. The coin is barely visible, and some have suggested that the painting has been tampered with precisely to obscure it. But, crucially, the coin between the man's fingers can be seen. Once you have noticed that coin, your entire take on the scene is likely to change. (Perhaps part of the painting may still make the scene Goethe talked about visible, but that is besides the point; Goethe's mistake was about the whole of the image Ter Borch painted.)



Figure 21: Gerard ter Borch II, *Gallant Conversation* (Detail)

Once that coin is pointed out, it has seemed obvious to many viewers that Ter Borch's image represents something other than Goethe thought. It is not unlikely that, had Goethe noticed that detail of the image, he would have agreed he got things wrong. If what I have said here is right, then Goethe's error was rooted in his perception of the image. This means that such error—such instances of misperception—can occur in a way that does not presuppose that any norm attaches to the image. Just as in the earlier two examples, the kind of error Goethe seems to have made is due to no more than a partial perception of the actual visible qualities of some object before our eyes. It is possible, because the circumstances in and conditions

under which we can come to see those qualities can be suboptimal.

It is a very general fact about perception that what we take ourselves to perceive is always susceptible to error. Our perception of images is no exception. It may always happen that, unbeknownst to you, the conditions under which you see the image are poor, or how the image appears is misleading (think for instance of the way anamorphic images do not reveal themselves when seen head-on). Such differences in how we get to see an image make for ample opportunity to get a view that is misleadingly partial, or does not put one in an easy position to tell what one sees. In such situations you may confidently claim to see no image, or to see an image of a kind not in fact around. As with any other visible quality of our surroundings, the typical route to find out if the image in fact represents how we think it does, is to look again, or look at it differently.

To be sure, not all possible disputes about images can be resolved by mere looking. There will be cases where you see things clearly, but where your mistake lies in some other false belief—something you thought you knew, something you thought to recognise. Looking at a painting may provide the final test of what can be seen by means of it, but this does not imply that our conclusions must always be safely drawn in ignorance of any further facts. As with other kinds of perception, what you know about the world makes a big difference, both to what you get to see, and to what you are prone or able to take yourself to see.

Having Ter Borch's painting and Goethe's error in mind, Baker and Hacker write how "It is amazing that different people can look at the same object and yet see completely different things" (Baker & Hacker, 1983, p.239). This is an exaggeration. Indeed, different people can look at the same object and yet see completely different things. Different people can also look at the same image and yet see completely different things. But that is hardly amazing. Understanding the nature of visual representation, and being sensitive to the partial character

of perception, should prevent us from being all too amazed.

We can of course delight in how repeat visits to a gallery can reward us with an increased level of detail, and a seemingly endless number of ways of looking at one and the same painting. But observing that images not only exploit the opportunities of our visual world, but also are a genuine part of it, that reward is more or less what we would expect.

5.5 Conclusion and summary

The central thesis I have argued for is that images are perceptual mediators for the sense of sight. Because of this, they bring the possibility of perceptual representation to visual perception. When an image functions as a visual mediator, it makes a scene it represents appear to someone who looks at the image in the right way. This means that the representational paintings, drawings, and photographs we are so familiar with make potentially absent things visible to us. The question *how* they do so, given their own visible qualities, is a psychological one. The question *why* they do so is otiose. Antoine Arnauld circumscribes the kind of conclusions I want to draw about images effectively. We must be careful, he writes,

not to seek reasons ad infinitum, but to stop when we get to what we know to be the nature of a thing, or what we know with certainty to be a quality of it. One must not ask why extension is divisible, for example, or why the mind is capable of thought, for it is the nature of extension to be divisible, and that of the mind to think. (Arnauld, 1990, p.50)

There is no use in asking, as philosophers, why images mediate vision. To mediate our visual perception of a great variety of other visibilia, I have shown, simply is their nature. The philosophical chal-

lenge images posed was how to fit the conception of images as visual mediators in our overall understanding of what vision is. Without such a fit, it is true, we would have reason to revise our conception of their nature. Without such a fit, we would have reason to look for alternative ways of understanding what an image is. Yet, precisely this challenge I have met.

In chapter 1 I identified a fundamental dispute about images. They seem to present us with the visual appearance of the scenes they represent. The dispute is about whether images do so by making absent scenes appear to a viewer, or by merely copying the way an object or event looks. This dispute is not frequently acknowledged, because so many find the very idea of seeing what is not present to the senses an oxymoron. I have shown how this apparent contradiction arises only when we make specific, controversial assumptions about vision. Late medieval traditions in the philosophy of perception, I brought out, adopt a different stance towards what vision is. Accordingly they left logical space for visions of the absent. Recent scholarship on visual culture in the Middle Ages brings out that visions of the absent mediated by images indeed played a crucial role in the way people conceived of their place in the world.

In recent years philosophers have still predominantly favoured the idea that images merely copy how things look. In chapter 2 I explained my dissatisfaction with that view, at least if construed as a philosophical theory about visual representation. I distinguished two kinds of copying, and showed that neither fits with what we already know about visual representation. The view that images visually resemble what they represent fails, because there are many visibilia that a picture cannot visually resemble. I showed that the view that images reproduce the appearance of what they represent also fails, because it disregards that images represent in a way that our eyes can detect. The unattractiveness of the copy theory in both its guises gives us reason to consider the novel alternative I have defended as a serious way for-

ward.

My main proposal to fit our conception of images in a larger understanding of vision was laid out and defended in chapter 3. How does the idea that images make absent objects visible fit with our current conception of perception? I proposed to look for an answer in the other sensory modalities. In audition, it is no mystery that sounds can make other things heard. Sounds are distinct, audible objects of perception. Hearing them is necessary for hearing anything at all. Sounds function as mediators in perception. Because hearing a sound does not depend on the actual presence of what it makes heard, auditory perception involves a form of representation—‘presence in absence’. Representational hearing (by means of sounds) and representational seeing (by means of images) are both species of the genus of representational perception. My arguments forge a connection between images and sounds that has not been explored in recent work on representation or perception. I have shown how it offers an attractive solution to our problem.

In chapter 4, I argued that we may regard mediate perception a genuine mode of perception. To many, this will have seemed required to meet the challenge satisfactorily. George Berkeley maintained that a mediated perception of something is not genuinely sensory. He thought it crucially involves a non-sensory act of cognition or thought. I responded to Berkeley’s argument for this conclusion, and showed how the theory of visual representation advanced by Hylas in the *Three Dialogues* allows us to conceive of mediated vision as genuinely sensory. Berkeley’s argument wrongly ignores the partial character of ordinary perception. Once we acknowledge that perception is partial, as we should, we can see that Hylas gets things right. Instead of relying on non-sensory or quasi-sensory effects of thought or cognition, coming to see a painted surface as an image consists in no more than discovering a previously unseen aspect of the visible world.

In the current chapter, I considered at a more abstract, systematic level how visual images contribute to the visible world. I showed how images take up a distinctive place in the structure of appearance. In visual encounters with the world, something can both be presented and represented to perceivers. Visual images introduce into our visual world distinctive representational aspects. This clarifies the relation between appearance and representation. Conceiving of visual representation as part of the visual world in the way I propose does neither suggest nor require an inherently normative dimension to visual representation. An image only accrues such norms or standards if someone or something uses it in some more specific practice. The diverging standards of photography, drawing from life, and therapeutic imagery show that, on these matters, the image as such remains silent.

I expect to be told that I have offered no more than a fragment of a theory of images as we in fact encounter them. Given the ways images have been taken up and have become the medium of art over the course of centuries, there remains a range of philosophical questions I have not addressed. I have only scratched the surface of how artists have exploited the possibility of representational imagery. My treatment of the value or significance of images has been cursory and unsystematic at best. But I have made no claim to a comprehensive theory of that sort. As far as I can tell, it is not at all obvious that there could be one theory that covers all the ways in which people have used images. What I have done is more modest, but nonetheless of importance to philosophical aesthetics and the philosophy of perception. I have shown how, in terms of a framework of perceptual appearance already implicit in our prephilosophical grasp of things, the fundamental connections between vision, images and what they represent can be given a clear, reasonable explanation. The idea that images make the scenes they represent visible, if I am correct, no longer needs to be regarded as incoherent, mystical, or philosophically naive. Visual representation simply belongs to our visual world, given how

our perceptual capacities have evolved in our engagement with the things and happenings around us. This view offers a serious alternative to theories of visual representation that assume that we cannot see beyond the image. To my mind, the very availability of the alternative I have defended undermines any theory premised on that assumption.

Bibliography

Abell, C. (2009) Canny resemblance. *Philosophical Review*, **118** (2): 183–223.

Abell, C. (2010a) Cinema as representational art. *British Journal of Aesthetics*, **50** (3): 273–286.

Abell, C. (2010b) The epistemic value of photographs. In: C. Abell & K. Bantinaki eds. *Philosophical perspectives on depiction*. Oxford: Oxford University Press.

Adriaenssen, H.T. (2011) Peter john olivi on perceptual representation. *Vivarium*, **49**: 324–352.

Alloa, E. (2013) Visual studies in byzantium: A pictorial turn avant la lettre. *Journal of Visual Culture*, **12** (3): 3–29.

Anscombe, G.E.M. (1965) The intentionality of sensation: a grammatical feature. In: R. J. Butler ed. *Analytical philosophy*. 2nd. Oxford: Blackwell.

Antonova, C. (2010) *Space, time, and presence in the icon*. Ashgate.

Arnauld, A. (1990) *On true and false ideas*. Manchester: Manchester University Press.

Baker, G.P. & Hacker, P.M.S. (1983) Dummett's frege or through a looking-glass darkly. *Mind*, **92**: 239–246.

Bantinaki, K. (2007) Pictorial perception as illusion. *British Journal of Aesthetics*, **47** (3): 268–279.

Batty, C. (2010) A representational account of olfactory experience. *Canadian Journal of Philosophy*, **40** (4): 511–538.

Batty, C. (2015) Olfactory objects. In: S. Biggs, M. Matthen, & D.

Stokes eds. *Perception and its modalities*. Oxford: Oxford University Press, pp.222–246.

Baxandall, M. (1972) *Painting and experience in fifteenth century italy*. Oxford: Oxford University Press.

Belting, H. (1994) *Likeness and presence: A history of the image before the era of art*. Chicago: University of Chicago Press.

Berkeley, G. (1954) *Three dialogues between hylas and philonous*. New York: The Liberal Arts Press.

Bertelson, P. (1999) Ventriloquism: A case of crossmodal perceptual grouping. In: G. Aschersleben, T. Bachmann, & J. Müsseler eds. *Cognitive contributions to the perception of spatial and temporal events*. Amsterdam: Elsevier, pp.347–362.

Biernoff, S. (2005) Carnal relations: Embodied sight in merleau-ponty, roger bacon and st francis. *Journal of Visual Culture*, **4** (1): 39–52.

Biernoff, S. (2002) *Sight and embodiment in the middle ages*. Basingstoke: Palgrave Macmillan.

Blumson, B. (2009) Defining depiction. *British Journal of Aesthetics*, **49** (2): 143–157.

Borges, J.L. (1964) *Labyrinths: Selected stories & other writings*. D. A. Yates & J. E. Irby eds. New Directions Publishing Corporation.

Bradley, F.H. (1887) Association and thought. *Mind*, **12** (47): 354–381.

Brady, T.F., Konkle, T. & Alvarez, G.A. (2011) A review of visual memory capacity: Beyond individual items and toward structured representations. *Journal of Vision*, **11** (5): 1–34.

Brady, T.F., Konkle, T., Oliva, A. & Alvarez, G.A. (2009) Detecting changes in real-world objects: The relationship between visual long-term memory and change blindness. *Communicative & Integrative Biology*, **2** (1): 1–3.

Brindley, G.S. & Lewin, W.S. (1968) The sensations produced by

electrical stimulation of the visual cortex. *Journal of Physiology*, **196**: 479–93.

Brubaker, L. (2012) *Inventing byzantine iconoclasm*. London: Bristol Classical Press.

Bruno, V.J. (1977) *Form and color in greek painting*. New York: W.W. Norton.

Budd, M. (2008a) *Aesthetic essays*. Oxford: Oxford University Press.

Budd, M. (2008b) The look of a picture. In: *Aesthetic essays*. Oxford: Oxford University Press, pp.216–238.

Burge, T. (2010) *Origins of objectivity*. Oxford: Oxford University Press.

Burnyeat, M. (2012) *Explorations in ancient and modern philosophy: Volume 1*. Cambridge: Cambridge University Press.

Burnyeat, M. (1982) Idealism and greek philosophy. *Philosophical Review*, **91** (1): 3–40.

Burr, D. (1985) Olivi, apocalyptic expectation, and visionary experience. *Traditio*, **41**: 273–288.

Byrne, A. (2009) Sensory qualities, sensible qualities, sensational qualities. In: B. McLaughlin & A. Beckermann eds. *The oxford handbook of philosophy of mind*. Oxford: Oxford University Press, pp.268–280.

Byrne, A. & Hilbert, D. (2008) Basic sensible qualities and the structure of appearance. *Philosophical Issues*, **18**: 385–405.

Carnap, R. (1928) *Der logische aufbau der welt*. Leipzig: Felix Meiner Verlag.

Caviness, M. (1998) Hildegard as designer of the illustrations to her works. In: C. Burnett & P. Dronke eds. *Hildegard of bingen: The context of her thought and works*. London: Warburg Institute, pp.29–62.

Christiansen, K., Weppelmann, S. & Rubin, P.L. (2001) *The renaiss-*

sance portrait: From donatello to bellini. New York: The Metropolitan Museum of Art.

Clark, A. (2015) Perceiving as predicting. In: D. Stokes, M. Matthen, & S. Biggs eds. *Perception and its modalities*. Oxford: Oxford University Press, pp.23–43.

Clark, T.J. (2008) *The sight of death*. New Haven: Yale University Press.

Clarke, T. (1965) Seeing surfaces and physical objects. In: M. Black ed. *Philosophy in america*. London: George Allen & Unwin Ltd, pp.98–114.

Collingwood, R.G. (1992) *The new leviathan: Or man, society, civilization and barbarism*. D. Boucher ed. Oxford: Clarendon Press.

Collingwood, R.G. (1938) *The principles of art*. Oxford: Oxford University Press.

Cook-Wilson, J. (1926) *Statement and inference*. A. Farquharson ed. Oxford: Clarendon Press.

Cottrell, D. & Campbell, M.E.J. (2014) Auditory perception of a human walker. *Perception*, **43**: 1225–1238.

Crane, T. (2011) The problem of perception. <http://plato.stanford.edu/entries/perception-problem/>.

Crowther, T. (2010) The agential profile of perceptual experience. *Proceedings of the Aristotelian Society*, **110** (2): 219–242.

Crowther, T. (2009) Watching, sight, and the temporal shape of perceptual activity. *Philosophical Review*, **118** (1): 1–27.

Danto, A.C. (1981) *Transfiguration of the commonplace*. Cambridge MA: Harvard University Press.

DeAngelus, M. & Pelz, J.B. (2009) Top-down control of eye movements: Yarbus revisited. *Visual Cognition*, **17**.

Delk, J.L. & Fillenbaum, S. (1965) Differences in perceived colour as a function of characteristic color. *The American Journal of Psychology*, **78** (2): 290–93.

DeLoache, J.S. & Burns, N.M. (1994) Early understanding of the representational function of pictures. *Cognition*, **52**: 83–110.

DeLoache, J.S., Pierroutsakos, S.L. & Uttal, D.H. (2003) The origins of pictorial competence. *Current Directions in Psychological Science*, **12** (4): 114–118.

Dennery, D.G. (2005) *Seeing and being seen in the later medieval world: Optics, theology and religious life*. Cambridge: Cambridge University Press.

Deshman, R. (1997) Another look at the disappearing christ: Corporeal and spiritual vision in early medieval images. *The Art Bulletin*, **79** (3): 518–546.

Dummett, M. (1973) *Frege: Philosophy of language*. New York: Harper & Row.

Edgerton, S.Y. (2009) *The mirror, the window, and the telescope*. Ithaca NY: Cornell University Press.

Emilsson, E.K. (2008) Plotinus on sense perception. In: S. Knuuttila & P. Kärkkäinen eds. *Theories of perception in medieval and early-modern philosophy*. Springer, pp.23–34.

Evans, G. (1982) *The varieties of reference*. J. McDowell ed. Oxford: Clarendon Press.

Freedberg, D. (1989) *The power of images: Studies in the history and theory of response*. Chicago: University of Chicago Press.

Galton, F. (1879) Composite portraits, made by combining those of many different persons into a single resultant figure. *The Journal of the Anthropological Institute of Great Britain and Ireland*, **8**: 132–144.

García, J.L.G. (2011) Empathetic images and painted dialogues: The visual and verbal rhetoric of royal private piety in renaissance Spain. In: S. Blick & L. D. Gelfand eds. *Push me, pull you: Imaginative and emotional interaction in late medieval and renaissance art*. Leiden: Brill, pp.487–528.

Gaver, W.W. (1993) What in the world do we hear?: an ecological

approach to auditory event perception. *Ecological psychology*, 5 (1): 1–29.

Gombrich, E. (1968) *Art and illusion*. 3rd Edition. London: Phaidon.

Goodman, N. (1976) *Languages of art*. 2nd edition. Indianapolis: Hackett Publishing.

Goodman, N. (1977) *The structure of appearance*. Dordrecht: Reidel.

Goodman, N. & Elgin, C.Z. (1988) *Reconceptions in philosophy and other arts and sciences*. London: Routledge.

Graham, J. (1997) Common sense and Berkeley's perception by suggestion. *International Journal of Philosophical Studies*, 5 (3): 397–423.

Greenberg, G. (2013) Beyond resemblance. *Philosophical Review*, 122 (2): 215–287.

Greenberg, G. (2012) *Pictorial semantics*. PhD thesis. Rutgers University.

Gregory, D. (2013) *Showing, sensing, and seeming: Distinctively sensory representations and their contents*. Oxford: Oxford University Press.

Grice, P. (1957) Meaning. *The Philosophical Review*, 66 (3): 377–388.

Hacker, P.M.S. (1987) *Appearance and reality: A philosophical investigation into perception and perceptual qualities*. Oxford: Blackwell.

Hacker, P.M.S. (1982) Events and objects in space and time. *Mind*, 16: 1–19.

Hahn, C. (2000) Visio dei: Changes in medieval visuality. In: R. S. Nelson ed. *Visuality before and beyond the renaissance: Seeing as others saw*. Cambridge: Cambridge University Press, pp.169–195.

Halliwell, S. (2002) *The aesthetics of mimesis: ancient texts and modern problems*. Princeton: Princeton University Press.

Hamburger, J.F. (2000) Seeing and believing: The suspicion of sight and the authentication of vision in late medieval art and devotion. In: A. Nova & K. Krüger eds. *Imagination und wirklichkeit: Zum verhältnis von mentalen und realen bildern in der kunst der frühen neuzeit*. Mainz, pp.47–69.

Hamburger, J.F. (1989) The visual and the visionary: The image in late medieval monastic devotions. *Viator*, **20**: 161–82.

Hatfield, G.C. & Epstein, W. (1979) The sensory core and the medieval foundations of early modern perceptual theory. *Isis*, **70** (3): 363–384.

Hempel, C.G. (1953) Reflections on nelson goodman's: The structure of appearance. *The Philosophical Review*, **62** (1): 108–116.

Hilbert, D. (1987) *Color and color perception: A study in anthropocentric realism*. Stanford: CSLI.

Honée, E. (1994) Image and imagination in the medieval culture of prayer: a historical perspective. In: H. van Os ed. *The art of devotion 1300-1500*. Princeton: Princeton University Press, pp.157–174.

Hopkins, R. (2010) Inflected pictorial experience its treatment and significance. In: C. Abell & K. Bantinaki eds. *Philosophical perspectives on depiction*. Oxford: Oxford University Press, pp.151–180.

Hopkins, R. (1998) *Picture, image and experience*. Cambridge: Cambridge University Press.

Hopkins, R. (2012) Seeing-in and seeming to see. *Analysis*, **72** (4): 650–659.

Hopkins, R. (2003) What makes representational painting truly visual? *Proceedings of the Aristotelian Society, Supplementary Volume*, **77** (1): 149–167.

Houston, K. (2013) *Shady characters: The secret life of punctuation, symbols, and other typographical marks*. New York: W.W. Norton.

Hutto, D.D. & Myin, E. (2013) *Radicalizing enactivism: Basic minds without content*. Cambridge, MA: MIT Press.

Hyman, J. (2012) Depiction. *Royal Institute of Philosophy Supplement*, **71**: 129–150.

Hyman, J. (1989) *The imitation of nature*. Oxford: Basil Blackwell.

Hyman, J. (2006) *The objective eye: Color, form and reality in the theory of art*. Chicago: University of Chicago Press.

Ishiguro, H. (1994) On representations. *European Journal of Philosophy*, **2** (2): 109–124.

Jackson, F. (1977) *Perception: A representative theory*. Cambridge: Cambridge University Press.

James, L. (2004) Senses and sensibility in byzantium. *Art History*, **27** (4): 522–537.

Janaway, C. (1995) *Images of excellence: Plato's critique of the arts*. Oxford: Oxford University Press.

Johnston, M. (2007) Objective mind and the objectivity of our minds. *Philosophy and Phenomenological Research*, **75** (2): 233–268.

Johnston, M. (2006) The function of sensory awareness. In: T. S. Gendler & J. Hawthorne eds. *Perceptual experience*. Oxford: Oxford University Press, pp.260–290.

Johnston, M. (2011) There are no visual fields (and no minds either). *Analytic Philosophy*, **52** (4): 231–242.

Jonas, H. (1954) The nobility of sight. *Philosophy and Phenomenological Research*, **14** (4): 507–519.

Kalderon, M.E. (2007) Color pluralism. *Philosophical Review*, **116** (4): 563–601.

Kalderon, M.E. (2015) *Form without matter: Empedocles and aristotle on color perception*. Oxford: Oxford University Press.

Kant, I. (1900) *Dreams of a spirit-seer*. London: Swan; Sonnenstein.

Keuls, E. (1975) Skiagraphia once again. *American Journal of Archaeology*, **79** (1): 1–16.

Kilwardby, R. (1993) *On time and imagination*. Oxford: Oxford University Press.

Krauss, R. (1985) *The originality of the avant-garde and other modernist myths*. Cambridge, MA: MIT Press.

Kulvicki, J.V. (2014a) Analog representation and the parts principle. *Review of Philosophy and Psychology*.

Kulvicki, J.V. (2014b) *Images*. Abingdon: Routledge.

Kulvicki, J.V. (2010) Pictorial diversity. In: C. Abell & K. Bantinaki eds. *Philosophical perspectives on depiction*. Oxford: Oxford University Press, pp.25–51.

Kusukawa, S. (2012) *Picturing the book of nature: Image, text, and argument in sixteenth-century human anatomy and medical botany*. Chicago: University of Chicago Press.

Lagerlund, H. & Thom, P. (2013) Introduction: The life and philosophical works of robert kilwardby. In: H. Lagerlund & P. Thom eds. *A companion to the philosophy of robert kilwardby*. Leiden: Brill, pp.1–15.

Le Poidevin, R. (1997) Time, and the static image. *Philosophy*, 72 (280): 175–188.

Leitch, S. (2001) Seeing objects in private devotion. In: L. Seidel ed. *Pious journeys: Christian devotional art and practice in the later middle ages and renaissance*. Chicago: The David; Alfred Smart Museum of Art, pp.43–60.

Lessing, G.E. (1874) *Laocoon*. R. Phillimore ed. London: Macmillan.

Levinson, J. (2012) Indication, abstraction, and individuation. In: C. M. Uidhir ed. *Art and abstract objects*. Oxford: Oxford University Press, pp.49–61.

Lindberg, D. (1976) *Theories of vision from al kindi to kepler*. Chicago: University of Chicago Press.

Loh, M.H. (2009) Renaissance faciality. *Oxford Art Journal*, 32 (3): 341–363.

Lopes, D. (2005) *Sight and sensibility: Evaluating pictures*. Oxford:

Clarendon Press.

Lopes, D. (1996) *Understanding pictures*. Oxford: Oxford University Press.

MacPherson, F. (2012) Cognitive penetration of colour experience: Rethinking the issue in light of an indirect mechanism. *Philosophy and Phenomenological Research*, **84**: 24–62.

Mallary, P. & Mallary, F. (1986) *A redouté treasury*. New York: The Vendome Press.

Marshall, P. (1981) Two scholastic discussions of the perception of depth by shading. *Journal of the Warburg and Courtauld Institutes*, **44**: 170–175.

Martin, M. (2012) Sounds and images. *British Society of Aesthetics*, **52** (4): 331–351.

Martin, M. (2010) What's in a look. In: B. Nanay ed. *Perceiving the world*. Oxford: Oxford University Press, pp.160–225.

Martin, M.G.F. (2003) Sensible appearances. In: T. Baldwin ed. *The cambridge history of philosophy*. Cambridge: Cambridge University Press, pp.519–530.

Matthen, M. (2010) On the diversity of auditory objects. *Review of Philosophy Psychology*, **1**: 63–89.

McGinn, B. (2006) Theologians as trinitarian iconographers. In: J. F. Hamburger & A.-M. Bouché eds. *The mind's eye: Art and theological argument in the middle ages*. Princeton: Princeton University Press, pp.186–207.

Meiss, M. (1951) *Painting in florence and siena after the black death: The arts, religion and society in the mid-fourteenth century*. Princeton: Princeton University Press.

Menn, S. (1998) *Descartes and augustine*. Cambridge: Cambridge University Press.

Miles, M. (1983) Vision: The eye of the body and the eye of the mind in saint augustine's 'de trinitate' and 'confessions'. *The Journal of*

Religion, **63** (2): 125–142.

Miles, M.R. (1979) *Augustine on the body*. Missoula: Scholars Press.

Miles, M.R. (1985) *Image as insight: Visual understanding in western christianity and secular culture*. Boston: Beacon Press.

Newall, M. (2010) Pictorial resemblance. *Journal of Aesthetics and Art Criticism*, **68** (2): 91–103.

Newall, M. (2011) *What is a picture? Depiction, abstraction, realism*. Hampshire: Palgrave Macmillan.

Newman, B. (2005) The visionary texts and visual worlds of religious women. In: J. F. Hamburger & S. Marti eds. *Crown and veil: female monasticism from the fifth to the fifteenth century*. New York: Columbia University Press, pp.151–171.

Noë, A. (2004) *Action in perception*. Cambridge, MA: MIT Press.

Noë, A. (2012) *Varieties of presence*. Cambridge MA: Harvard University Press.

Nudds, M. (2014) Auditory appearances. *Ratio*, **27**: 462–482.

O'Callaghan, C. (2008a) Object perception: Vision and audition. *Philosophy Compass*, **3/4**: 803–829.

O'Callaghan, C. (2008b) Seeing what you hear: Cross-modal illusions and perception. *Philosophical Issues*, **18** (1): 316–338.

O'Callaghan, C. (2007) *Sounds: A philosophical theory*. Oxford: Oxford University Press.

O'Callaghan, C. & Nudds, M. eds. (2009) *Sounds and perception*. Oxford: Oxford University Press.

O'Connor, D.J. (1945) On resemblance. *Proceedings of the Aristotelian Society*, **46**: 47–76.

O'Shaughnessy, B. (2000) *Consciousness and the world*. Oxford: Clarendon Press.

O'Shaughnessy, B. (1972) Processes. *Proceedings of the Aristotelian Society*, **72**: 215–240.

O'Shaughnessy, B. (1985) Seeing the light. *Proceedings of the Aris-*

totalian Society, **85**: 193–218.

O'Shaughnessy, B. (2003) Sense data. In: B. Smith ed. *John searle*. Cambridge: Cambridge University Press.

Palmer, S.E. (2002) Perceptual organisation in vision. In: H. Pashler ed. *Steven's handbook of experimental psychology: Volume 1, sensation and perception*. Wiley, pp.177–234.

Panofsky, E. (1972) *Studies in iconology: Humanistic themes in the art of the renaissance*. ICON.

Pasnau, R. (1999) What is sound? *Philosophical Quarterly*, **49** (196): 309–324.

Peacocke, C. (1987) Depiction. *The Philosophical Review*, **96** (3): 383–410.

Peacocke, C. (1983) *Sense and content: Experience, thought, and their relations*. Oxford: Clarendon Press.

Phillips, A. (1992) Drawing from life. In: J. Hopkins & A. Savile eds. *Psychoanalysis, mind and art: Perspectives on richard wollheim*. Oxford: Blackwell Publishing, pp.315–336.

Phillips, I. (2014) Experience of and in time. *Philosophy Compass*, **9** (2): 133–144.

Pitkin, H. (1967) *The concept of representation*. Berkeley: University of California Press.

Price, H.H. (1964) Appearing and appearances. *American Philosophical Quarterly*, **1** (1): 3–19.

Pryor, J. (2000) The skeptic and the dogmatist. *Noûs*, **34** (4): 517–549.

Reid, T. (1895) *Essays on the intellectual powers of man*. 8th edition. W. Hamilton ed. Edinburgh: James Thin.

Rickless, S.C. (2013) *Berkeley's argument for idealism*. Oxford: Oxford University Press.

Riginos, A.S. (1976) *Platonica: The anecdotes concerning the life and writings of Plato*. Leiden: Brill.

- Rowlands, M. (2015) Arguing about representation. *Synthese*: 1–18.
- Savedoff, B. (2008) Documentary authority and the art of photography. In: S. Walden ed. *Photography and philosophy: Essays on the pencil of nature*. Malden MA: Blackwell Publishing, pp.111–137.
- Schier, F. (1986) *Deeper into pictures: An essay on pictorial representation*. Cambridge: Cambridge University Press.
- Scribner, R. (2001) *Religion and culture in germany (1400-1800)*. L. Roper ed. Brill.
- Scruton, R. (2010) Hearing sounds. In: D. W. Zimmerman ed. *Oxford studies in metaphysics* 5. Oxford: Oxford University Press, pp.271–278.
- Searle, J. (1983) *Intentionality: An essay in the philosophy of mind*. Cambridge: Cambridge University Press.
- Sellars, W. (1962) Philosophy and the scientific image of man. In: R. Colodny ed. *Frontiers of science and philosophy*. Pittsburgh, PA: University of Pittsburgh Press, pp.35–78.
- Short, T.L. (2007) *Peirce's theory of signs*. Cambridge: Cambridge University Press.
- Sider, T. (1997) Four dimensionalism. *Philosophical Review*, **106**: 197–231.
- Siegel, S. (2010) Do experiences have contents? In: B. Nanay ed. *Perceiving the world*. Oxford: Oxford University Press, pp.333–368.
- Silva, J.F. (2014a) Augustine on active perception. In: *Active perception in the history of philosophy*. Springer, pp.79–98.
- Silva, J.F. (2014b) Medieval theories of active perception: An overview. In: J. F. Silva & M. Yrjönsuuri eds. *Active perception in the history of philosophy*. Springer, pp.117–146.
- Silva, J.F. & Toivanen, J. (2010) The active nature of the soul in sense perception: Robert kilwardby and peter olivi. *Vivarium*, **48**: 245–278.
- Simcock, G. & DeLoache, J.S. (2006) Get the picture? The effects

of iconicity on toddlers' reenactment from picture books. *Developmental Psychology*, **1352-1357**.

Sterne, J. (2003) *The audible past: Cultural origins of sound reproduction*. Durham: Duke University Press.

Stokes, A. (1972) *The image in form: Selected writings of adrian stokes*. R. Wollheim ed. Harmondsworth: Penguin.

Stokes, A. (1965) *The invitation in art*. London: Tavistock.

Stokes, D. (2014) Cognitive penetration and the perception of art. *Dialectica*, **68** (1).

Tachau, K.H. (2006) Seeing as action and passion in the thirteenth and fourteenth centuries. In: J. F. Hamburger & A.-M. Bouché eds. *The mind's eye: Art and theological argument in the middle ages*. Princeton: Princeton University Press, pp.336–359.

Tachau, K.H. (1988) *Vision and certitude in the age of ockham*. Leiden: Brill.

Toivanen, J. (2013) *Perception and the internal senses: Peter of john olivi on the cognitive functions of the sensitive soul*. Leiden: Brill.

Tolliver, J.T. (2007) Sensing, perceiving, and thinking: On the method of phenomenal contrast. *The Southern Journal of Philosophy*, **45**: 143–151.

Travis, C. (2013) *Perception: Essays after frege*. Oxford: Oxford University Press.

Travis, C. (2004) The silence of the senses. *Mind*, **113** (449): 57–94.

Uchiyama, M., Nishio, Y., Yokoi, K., Hirayama, K., Imamura, T., Shimomura, T. & Mori, E. (2012) Pareidolias: complex visual illusions in dementia with lewy bodies. *Brain*.

Urmson, J.O. (1968) The objects of the five senses. *Proceedings of the British Academy*, **54**: 117–131.

van Inwagen, P. (1994) Composition as identity. *Philosophical Perspectives*, **8**: 207–220.

W. H. Warren, J. & Shaw, R.E. eds. (1985) *Persistence and change*:

Proceedings of the first international conference on event perception. Hillsdale, NJ: Lawrence Erlbaum.

Walton, D.N. (2014) Argumentation schemes for argument from analogy. In: H. J. Ribeiro ed. *Systematic approaches to argument by analogy.* Springer.

Walton, K.L. (1990) *Mimesis as make-believe: On the foundations of the representational arts.* Cambridge MA: Harvard University Press.

Walton, K.L. (1984) Transparent pictures: On the nature of photographic realism. *Critical Inquiry*, **11** (2): 246–277.

Wardy, R. (1990) *The chain of change: A study of aristotle's physics vII.* Cambridge: Cambridge University Press.

Wiesing, L. (2013) *Sehen lassen: Die praxis des zeigens.* Frankfurt am Main: Suhrkamp.

Williams, B. (2002) *Truth and truthfulness.* Princeton: Princeton University Press.

Wittgenstein, L. (2009) *Philosophical investigations.* P. Michael, S. Hacker, & J. Schulte eds. Oxford: Blackwell Publishing.

Wollheim, R. (1980) *Art and its objects: With six supplementary essays.* 2nd edition. Cambridge: Cambridge University Press.

Wollheim, R. (2003) In defence of seeing-in. In: H. Hecht, R. Schwartz, & M. Atherton eds. *Looking into pictures: An interdisciplinary approach to pictorial space.* Cambridge, MA: MIT Press, pp.3–15.

Wollheim, R. (2001) On aesthetics: A review and some revisions. *Literature and Aesthetics*, **11**: 7–29.

Wollheim, R. (1974) *On art and the mind.* Cambridge MA: Harvard University Press.

Wollheim, R. (1998) On pictorial representation. *Journal of Aesthetics and Art Criticism*, **56** (3): 217–226.

Wollheim, R. (1996) On the assimilation of pictorial art to language. In: J. Thompson ed. *Towards a theory of the image.* Maastricht: Jan

van Eyck, pp.22–41.

Wollheim, R. (1977) Representation: The philosophical contribution to psychology. *Critical Inquiry*, **3** (4): 709–723.

Yarbus, A.L. (1967) *Eye movements and vision*. New York: Plenum Press.

Zacks, J.M., Speer, N.K., Swallow, K.M., Braver, T.S. & Reynolds, J.R. (2007) Event perception: A mind-brain perspective. *Psychological Bulletin*, **133** (2): 273–293.