Introduction: a trailer to the book

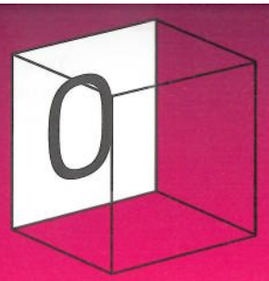

The eyes of a really scary jumping spider (*Metaphidippus*), normally resident in California. Generally, the bigger you are, the better you can see (remember Little Red Riding Hood—'what great big eyes you have, Grandma…'), but this furry little creature has exceptionally good visual acuity given its relatively small body size.

'I see nobody on the road,' said Alice. 'I only wish I had such eyes,' the King remarked in a fretful tone. 'To be able to see Nobody! And at that distance too! Why, it's as much as I can do to see real people, by this light!'

Lewis Carroll,
Through the Looking Glass

An apology

We would like to apologize in advance for the content of this book. We set out not to write the definitive textbook on vision, but to write an enjoyable one. We have occasionally (well, quite often really) sacrificed what some people hold to be 'the truth' for the sake of a coherent story, and simplified some ideas to the extent that our colleagues in vision science will be outraged to see their own area of the subject being presented so frivolously, or even ignored altogether. Therefore, what you have here is, at best, the bare bones of vision. We hope that you will enjoy our book and that you will feel inspired to go and read a book that is more advanced but probably more boring than ours (we list some of these at the end of each chapter).

The problem

Vision appears very easy. As we glance out of the window we are met with a wealth of information. We can see that there are red berries on the bush, and that they are now ripe for the birds to eat; we can see there is a slight drizzle in the air; we can see a man and we immediately know that it is someone we have never seen before, yet we can tell that, at this moment, he doesn't look very happy. All this from a single glance. We do this so easily and so quickly that it does not strike us just what an amazing accomplishment each glance actually is. But let's think about this for just a moment.

Light can be thought of as lots of little balls (photons, to give them their proper name) that arrive from the sun, or a light bulb, or some other source. These balls bounce off the objects around us, the trees, cars, humans, and so on, and some of them just happen to bounce in such a way that their path takes them from the object and through a tiny transparent window in our eye. Inside our eye, these balls are collected and vision begins. By way of analogy, we could imagine a little man with a net catching these light balls and, from the contents of his net, doing all the things that vision does—seeing colours, knowing that a car is coming towards him, recognizing that his friend has a new haircut, reaching out and picking up his coffee cup ... and a million other things. Put like this, vision appears bewilderingly difficult. This book aims to present some of what we know about how we take the contents of this net and turn it into 'seeing'.

Take a look at Figure 0.1. What do you see? The answer is probably 'not much, just a lot of spots'. However, if you continue to look at the figure for some time your

Figure 0.1 What can you see in this picture?

impression will probably change. After a while most people have a sudden revelation that it is not just a set of spots, it's the picture of a spotty dog sniffing at the ground (for those of you still struggling, turn the page and there is help for you in Figure 0.2). From now on, whenever you see this figure, you will not see just spots, it will always be the spotty dog. Somehow, your brain has interpreted this figure and made an object from the spots. This process is at the heart of what vision does—it is always trying to make sense of its input and telling us about the things we need to know about. Vision is active, not passive. Our eye is often likened to a camera—both have an aperture at the front, some type of lens to focus the light, and then something that absorbs the light at the back. This may not be a bad analogy for the eye itself, but it is a bad analogy for vision. Why? Because the camera doesn't have to do very much; it doesn't have to interpret the world and act accordingly—it doesn't have to produce dogs from spots. The camera never lies, because the camera doesn't have to tell us what it sees—but our visual system does. Our visual system is not there to faithfully record the image outside, it is there to give us the necessary information for us to behave appropriately.

The last phrase we used, '*necessary information for us to behave appropriately*', is actually rather important. It tells us that we see the world in a particular way, not because that is the way the world is, but because that's the way we are. Let's try and clarify this puzzling statement with an example. Figure 0.3 shows what is known as an Ishihara plate. It's just a bunch of spots. No dog this time, but most of us can see the number 74 hidden in the dots. However, there are some people who can see no such number. Show it to your pet dog or cat—you will find that they cannot see it either, though your goldfish will see the 74 with ease! Why? Surely the number 74 is either there or it isn't. The answer lies in your physiology. The colour of the spots has been subtly designed so that the difference between the 74 and its background can only be detected by animals with a certain type of colour vision. As we shall learn in Chapter 5,

⇒ See Chapter 5

Figure 0.2 Now can you see what the picture is? Look back at Figure 0.1 and you'll see the Dalmatian clearly, we hope.

human colour vision is different from that of nearly all other mammals in having great ability to distinguish between the colours we call red and green. To most other mammals (cats, dogs, horses, bulls, etc.) these colours do not exist—they see them all as various shades of browny yellow (we are being rather bold here in claiming to know what it's like to be a dog, but we hope you get the gist). Likewise, the small proportion of humans whom we term 'red–green colour blind' cannot see this number 74 at all; they are invariably dressed in bizarre combinations of colours and will happily put a

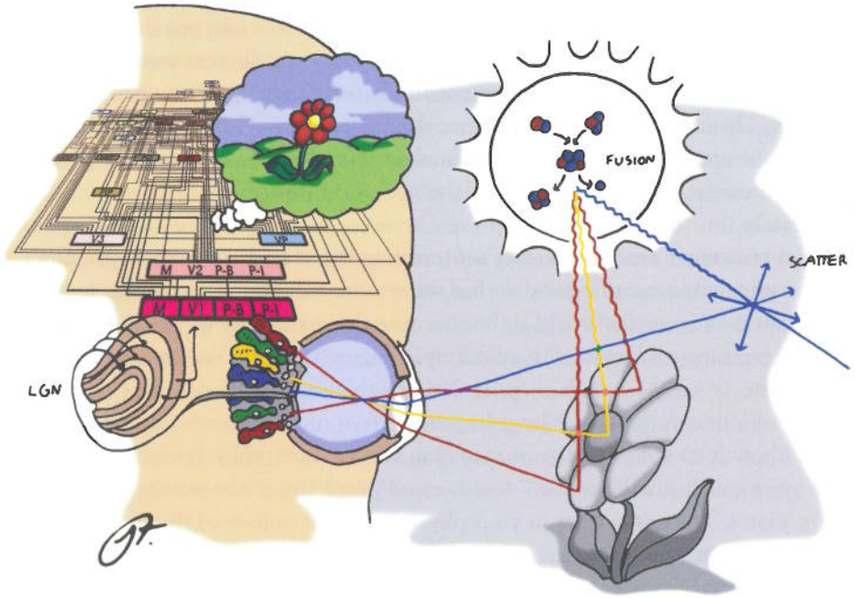

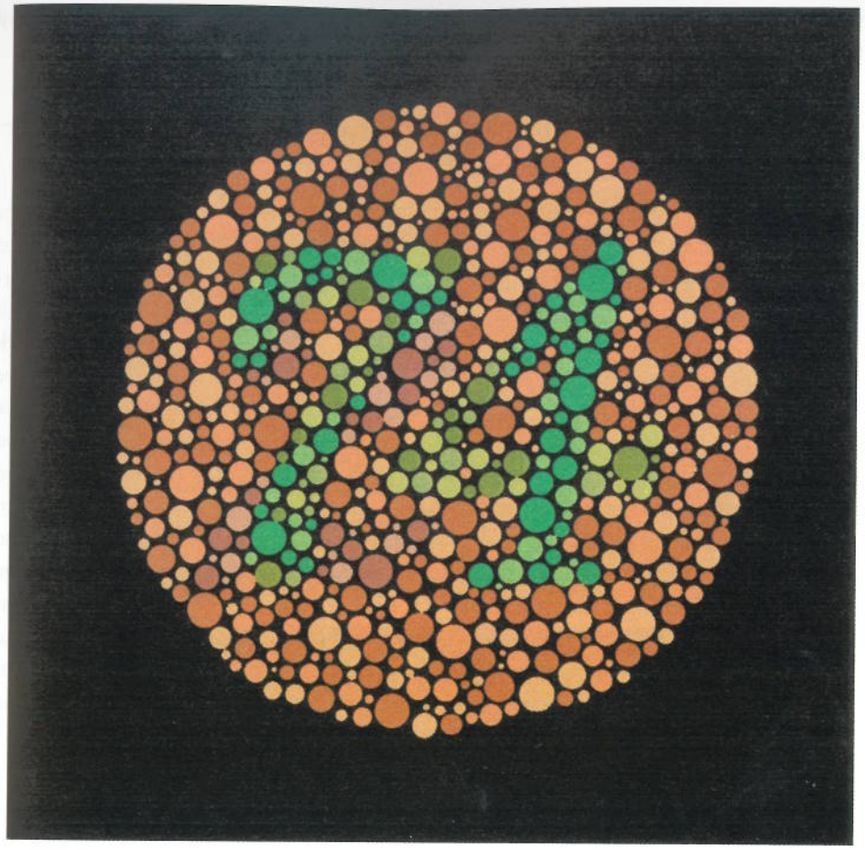

Figure 0.3 A plate from the Ishihara test for colour blindness. Can you see the number?

red sock on one foot and a green sock on the other as both look the same to them, providing great amusement for the rest of us.

At night, we lose our ability to see colours completely. If you lead a sad and uninteresting life, you may have this book as bedtime reading; if so, turn off the light and try looking at Figure 0.3 in your dimly lit bedroom. You will no longer be able to see the number 74. Indeed, a simple glance around the room will tell you that there is no colour anywhere. How can this be? Surely a green object will still be green even if there is not much light around? Not so. 'Green-ness' is an invention of your brain, and if there is not enough light around it cannot create this particular magic we call colour.

So why do we have an ability to see colour that our pet dog lacks? Is he adequately compensated by being able to lick parts of his own anatomy that we can't? The statement '*necessary information for us to behave appropriately*' seems to imply that there is something we do that needs this information that dogs don't do. As we shall see in Chapter 5, this may well have to do with picking ripe fruit, something of great importance

See Chapter 5

to our ancestors, but an activity for which most dogs lack any enthusiasm. Quite why the goldfish should want to go picking ripe fruit we shall leave for you to consider.

Vision in action

So, vision isn't there merely to form a pretty picture of the world—it's there in order for us to be able to make sense of what is out there and to interact with it, and to actively seek information about the world. If you carefully watch a person as they perform any visual task, from merely watching TV to driving a car at high speed, you will see that their eyes do not merely stand still and absorb information; instead, they flick rapidly from place to place, in almost constant motion (Yarbus, 1967). If we measure eye movements carefully, we find that the eyes move about 3–4 times per second as we perform these tasks. Figure 0.4 shows the pattern of eye movements that a person made while looking at a simple stationary picture. It shows that the person looked from place to place as they inspected this figure. It also shows that the movements of the eye were not random, but each movement caused the central part of their vision to fall upon an 'interesting' part of the picture, such as the eyes of the statue. The reason for this is that vision is not uniform across our field of view. Our central vision gives us really good detail of the point we are looking at. However, as we move further and further away from this point our vision gets worse and worse (see Figure 0.5). Thus, at any one moment we only have really sharp vision for one small part of the world and

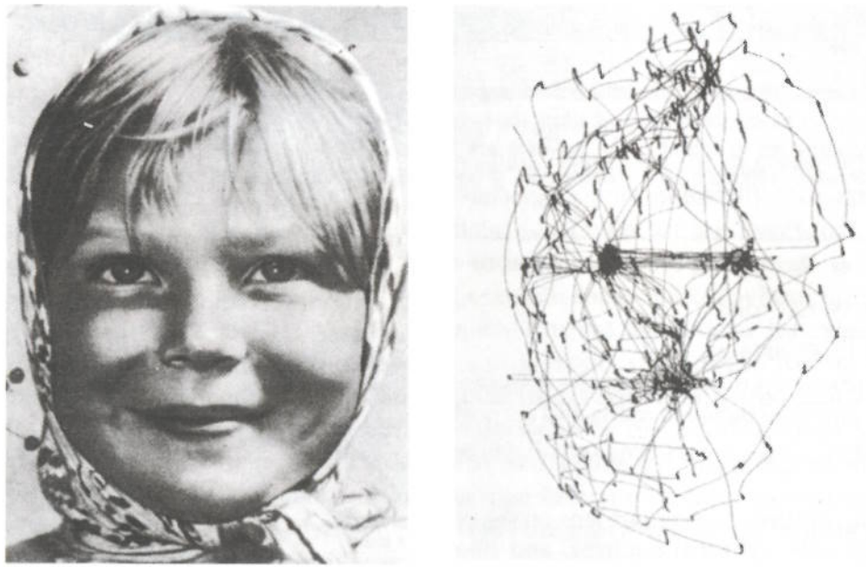

Figure 0.4 Pattern of eye movements whilst looking at a picture.

Figure 0.5 Acuity decreases with eccentricity. This picture has been tampered with so that it becomes more blurred away from the point of fixation (the central figure). This mimics the way our own vision loses the ability to see fine detail away from the centre of vision.

in order to understand what is going on we need to move our eyes around. This raises the question of how we do this—how do we decide where to look next? Given that these eye movements are highly structured to what is in the picture (Figure 0.4), it must be that we use the 'poor' information to guide these movements to the next point of interest. Clearly, if we are to understand how we see, we must also understand this active process of grabbing information from the scene in front of us.

Illusions

Figure 0.6 shows a remarkable illusion that even we still find hard to accept. It looks like a picture of two tables, one long and thin and the other short and fat. However, the table-tops are exactly the same shape (**Shepard, 1990**). Yes, that's what we said—they are the *same shape*. To convince yourself of this, trace out the shape of one and then overlay it on the other. You may be tempted to conclude that your visual system is pretty useless, as it has failed to spot that the two shapes are the same. However, that would be a mistake, because what your brain has actually done is very clever. It has turned the two-dimensional (2-D) picture on the page into the most plausible three-dimensional (3-D) objects. That is, we do not see these tables as merely lines on a piece of paper. To us they become objects, we recognize them as tables—things we put stuff on. In doing so (and we have no choice in this matter), we scale and distort these two identical quadrilaterals into one that is long and thin and one that is short and fat. As we shall see in Chapter 7, our willingness to produce three dimensions from the poorest of information can account for many so-called 'illusions'.

⇒ See Chapter 7

So, our visual system can, and frequently does, get it 'wrong'. It does not represent the physical world correctly because it makes assumptions about the world and uses these to give us the information we need. In the case of the tables in Figure 0.6, what is identical about these tables is their projection on the back of your eye. But we are not interested in what is on the back of our eye—we are interested in what is out there in the world. As we do not need to know about what's on the back of our eye, we do not need conscious access to it (it would be rather confusing if we did). Instead, our brain gives us a table of a certain shape, and if we need a long skinny table for our buffet lunch, this is far more useful.

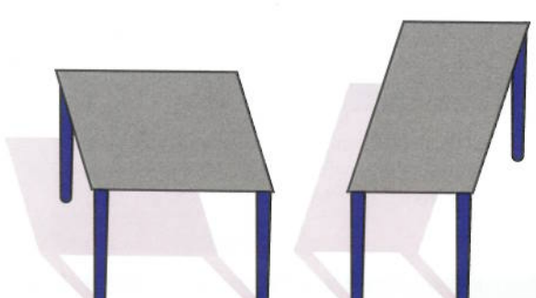

Figure 0.6 These two table tops are identical, but do they look it?

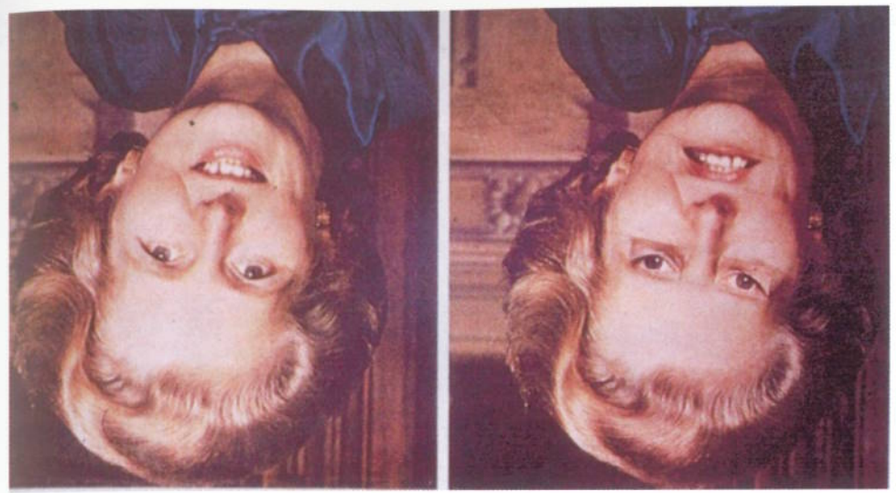

Figure 0.7 The Thatcher illusion. Turn the page upside down to see the 'real' Lady Thatcher.

Therefore, illusions provide us with a means of examining the workings of the visual system. For instance, if you look briefly at Figure 0.7 you will see two pictures of a once well-known face. The pictures look pretty much alike, but if you look carefully you will see that the one on the right has been tampered with. The eyes and the mouth have actually been cut out, inverted, and stuck back down. However, this change is quite hard to spot. Now, turn the book upside-down so that the picture itself is the right way up. Now the changes become obvious and the tampered face takes on a grotesque appearance (**Thompson, 1980**) (some might argue that attempting to make Lady Thatcher look grotesque is rather redundant—we couldn't possibly comment). This tells us that there must be something different about the face when it is upside down because we were unable to read the facial expression correctly. We are obviously used to seeing and interpreting faces which are the right way up, and perhaps the neural machinery we use to do this is only designed (or has only learned) to do this for the pictures that are the right way up (more in Chapter 10).

⇒ See Chapter 10

Box 0.1

When is an illusion an illusion?

You may notice in this chapter that we are a little coy about what we mean by the term visual illusion, so much so that we sometimes put the 'illusion' bit in quotation marks just to confuse you. At one extreme, we could opine that everything is an illusion, but that would be both trite and tiresome. Perhaps it is easier to consider an example. Consider the Adelson figure (Figure 0.1.1). The two marked parts of the figure (A and B) have identical luminance—that is, the same amount of black ink on the page. If you don't believe us, take a pair of scissors and cut out these two squares

Box continues . . .

and lay them side by side. You will find that they are identical. However, in this 'illusion' they appear to have very different lightnesses.

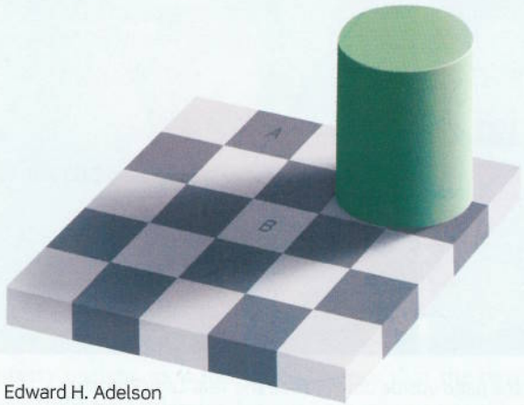

Edward H. Adelson

Figure 0.1.1 The Adelson illusion. The two squares marked A and B have exactly the same luminance but their perceived lightness is very different. We perceive A as a dark square and B as a light square in shadow. Does this show how easily fooled your visual system is or how clever it is?

Now, we could claim that this represents a visual illusion because we perceive physically identical bits of paper as being of different lightness. But, in the scene portrayed in the picture, the two chequerboard squares are really different—one is a black square and the other is a white square—so our visual systems have given us the correct answer. Why should we regard this as an illusion? Now imagine that, rather than looking at a *picture* of a cylinder casting a shadow on a chequerboard, we were looking at a *real* cylinder casting a shadow on a chequerboard. Again, we would perceive the lighter square as being lighter than the darker square, but would we want to call this an illusion just because the luminances of the squares are the same by virtue of the shadow? This would be crazy; we would be calling something a visual illusion when we perceived it correctly. Because of this, illusions are fairly rare in natural scenes—though not unknown. Generally, we see the world the way it is. Once we get in a completely artificial environment like a laboratory it is possible to create illusions with ease. A case in point would be the Hermann grid (Figure 0.1.2). This really is an illusion because the 'ghostly grey dots' at the white intersections do not exist in the picture, nor do they exist in any object that the picture is seeking to represent. At the end of Chapter 4 we shall make this point again. There, we will examine so-called visual illusions that depend upon the visual system interpreting 2-D pictures as 3-D scenes and, as a result, 'misperceiving' the size of objects. Arguably, these shouldn't be regarded as illusions at all because the visual system is accurately representing what the picture is of.

See Chapter 4

Even by our definition there are some occasions when we experience illusions in the real world, but they tend to be when things are a bit weird—when driving in fog we underestimate our speed; when swimming underwater we tend to underestimate distance and hence overestimate the size of underwater objects. Have you ever noticed that the full Moon on the horizon looks much bigger than when it's high in the sky (Figure 0.1.3). The classic explanation is that the horizon Moon looks further away and therefore must be bigger. For more details on this explanation, and why it's probably wrong, see Chapter 7.

See Chapter 7

Box continues . . .

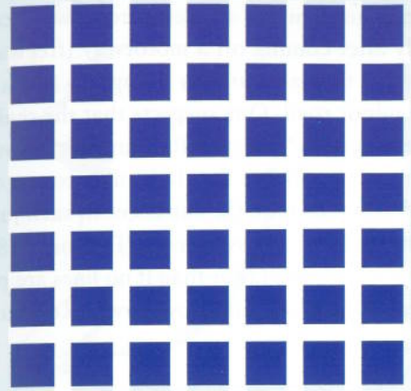

Figure 0.1.2 The Hermann grid. Do you see the ghostly grey dots at the intersections of the white lines? Do they really exist? What happens when you look directly at one of the grey dots? It should disappear.

Figure 0.1.3 The Moon illusion.

There are other sorts of visual illusion that can be used to investigate our visual senses, illusions in which we deliberately induce a change in perception by doing something out of the ordinary. For instance, we can use a technique called **adaptation** to temporarily alter our visual systems. You have probably had this experience yourself. You have been driving quickly down a major road for some, time and then you have to switch to a minor road where you are required to drive more slowly; you often get the feeling that everything is happening too slowly. The 30 mph (50 km/h) that you are now doing feels as though it is only half that. This is an example of adaptation. Something in your brain had 'got used to' the fast speed and produced this illusion. With more careful experiments in a laboratory we can use such phenomena to investigate how it is that we can sense speed in the first place, and why we get it so wrong in

this situation. This, in turn, brings us to one of the main reasons for studying vision. If we can understand why these things occur, we might be able to do something to help. In Britain (and probably many other countries too), motor accidents frequently occur as drivers negotiate a junction immediately after coming off a motorway (freeway). The nature of many of these accidents is that the car is simply going too fast as it attempts to negotiate the transition to the minor road. One suspects that the reason for this is the illusion we have just described. Drivers 'feel' as if they are going slowly when in fact they are moving much more quickly. So, we should be able to use our knowledge of speed perception to help. Indeed, in Britain, the entrances to such junctions are often now marked by a series of lines painted across the road (**Denton, 1980**). We know that such marks make objects appear to move faster than if no lines are present. Hence, drivers think that they are travelling faster and slow down. Road tests show that this works—and saves lives.

Damage to the system

> In any well made machine one is ignorant of the working of most of the parts—the better they work the less we are conscious of them ... It is only a fault that draws attention to the existence of a mechanism at all.
>
> Kenneth Craik,
> *The Nature of Explanation* (1943)

This quote from Kenneth Craik (who, incidentally, died in a road accident) helps summarize what we have tried to learn so far and what we are about to say. It emphasizes that our visual system is indeed a 'well made machine'—so well made that it works without us being aware of its processes. However, in some individuals there appears to be some fault in the system—unfortunate for them, but fortunate for scientists. When we find such an individual we can get a glimpse of what it is like not to have some process that we take for granted. Let us illustrate this with an example. In Chapter 9, we discuss cases of visual neglect—a condition in which the sufferer seems to just ignore half of the world. Figure 0.8 gives an example of such a person doing a simple (well, simple to us) task. When asked to copy the picture of a cat she did pretty well, apart from missing out one half completely; however, she believed she had drawn the whole animal. When asked to draw a clock, neglect patients draw a circle containing all the numbers (they have not lost their intelligence—they know what a clock should look like) but they put all the numbers on the right-hand side of the dial. When they read, they only read the right-hand side of the page. When presented with a plate of food they only eat the things on the right-hand side of the plate; simply turning the plate allows them to finish the meal! It seems that they are continually pulled to the right and things on the left fail to grab their attention.

See Chapter 9

From this information, we can immediately see that there must be something in our brain that normally allows objects to grab our attention and be processed more

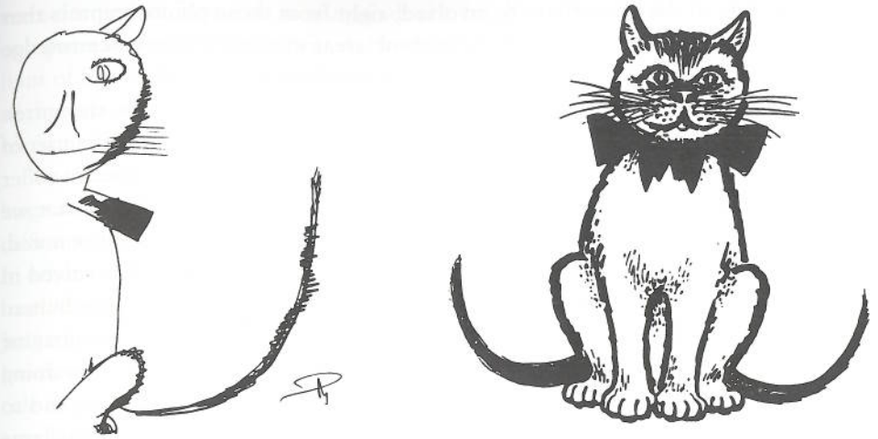

Figure 0.8 People with visual neglect often fail to copy one half of a picture.

thoroughly. Careful consideration of just what such patients can and cannot do, along with studies of just what has been damaged, allows us great insights into what the brain is doing and how it does this. Of course, we then hope to use this information to help such patients.

The list of possible visual problems caused by brain damage is huge, probably because no two cases of brain damage will ever be identical, but we shall introduce you to cases where colour or motion appear to have been affected, and to cases where, for example, a man can look at the face of his wife and have no idea who she is, although a sight of her shoes is enough to reveal her identity.

The brain

Earlier, we drew an analogy between a camera and the eye, and said that this is fine for the eye but not for vision. So where does 'vision' take place? This is a pretty hard question to answer. First, there is conscious vision, and there are theories of how and where this might arise that we shall discuss in Chapter 11. Second, there is all the visual processing that occurs without any obvious consciousness; for instance, as we walk around a room we somehow must have calculated just where the chair was and then have sent signals to our legs so that we could avoid it, but we were never conscious of doing so. Third, we know from being students ourselves many years ago that there are certain things inside the eye called **photoreceptors** that are busy changing light into electrical signals. These things have certainly got a lot to do with vision (without them we are blind), but these cells cannot be regarded as 'seeing'. So, to get back to our question of '*where does vision take place*', the boldest answer we're willing to give is '*in the brain*'. Therefore, the study of vision seems to necessitate that we understand

See Chapter 11

the various bits of the brain that are involved, right from those photoreceptors that initially change light into neural signals, to those areas that might give us knowledge of the mood of our friends, or allow us to steer a car along a busy road.

Figure 0.9 is a picture of the brain of a primate. As in the human brain, the cortex (the outside layer) is crumpled into many folds so, to make it easier, the lower part of the picture shows the cortex as if the poor monkey had been run over by a steam-roller (Felleman and Van Essen, 1991). To help us further, the parts of the brain that we know have something to do with vision are shown in colour. Two points can be noted. First, someone had to do a lot of colouring—over 50% of the cortex is involved in visual processing. Recently, we have begun to be able to chart what areas of the human brain might be involved in visual processing (Tootell et al., 1996). You might imagine that it would be far less—after all, we humans presumably need lots of brain for doing the things we think we do so well: reasoning, playing chess, doing hard sums, and so on. However, it appears that the amount of cortex we devote to vision is just as large as in this monkey—over half of your cortex is devoted to merely seeing. If we do some very rough sums and generously give hearing 10% of the cortex, all your other senses another 10%, and probably we should allow 10–20% for moving all the bits of your body, you can see that there is hardly anything left over for doing the hard bits like chess, crosswords, and perception assignments. But that is exactly the point. Vision is easy because we have lots of brain devoted to it; chess is hard because we don't have lots of brain devoted to it. If the situation were reversed and over 50% of your cortex

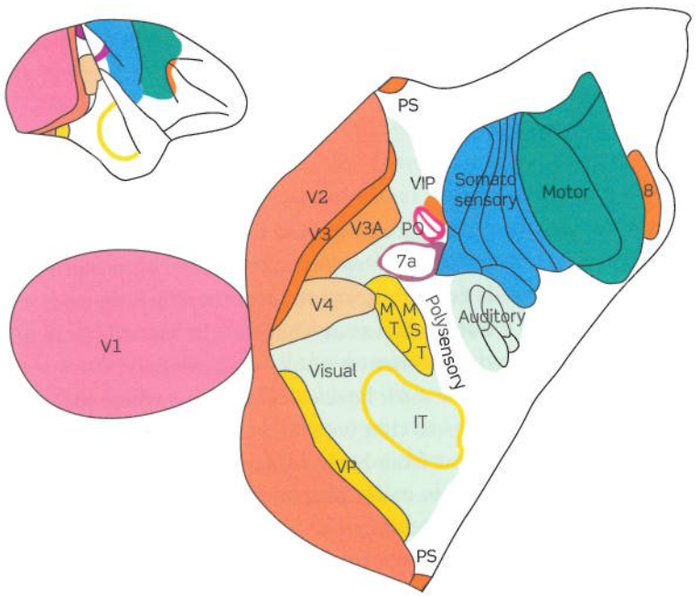

were devoted to chess, you would, no doubt, be able to take on the most sophisticated computer in the world and beat it. However, computers can now beat the most brilliant of humans when it comes to chess, but no computer in the world can put up the slightest of challenges when it comes to vision. Indeed, we would happily back a humble housefly against the most sophisticated computer in a 'vision Olympics'.

We mentioned that there were two things to note about Figure 0.9. The first was the large amount of colouring-in required; the second is that different colours are used in different places to signify that there is something different about these different areas. In other words, the area of brain devoted to vision is not one homogeneous mass, but it has lots of areas (over 30 at the last count), each of which is doing something different. Just why we have so many areas, and what each of them is doing, is a major challenge for vision science. At the time of writing, we believe that each might have a specialized role to play in analyzing some part of the scene. As an example, we could consider the area labelled MT (a small area the size of half a pea and coloured yellow for our convenience). Examination of this area, as we shall see in Chapter 6, has suggested that it must play a large role in our perception of motion. All the brain cells inside this area seem to care about the movements of the outside world—they do not fire much in response to a static scene, but they burst into action if an object moves in a particular way. If this area is involved in our sensing the motion of objects, we might predict that if we were to destroy it then we wouldn't see any motion. It may be hard to imagine, but this does indeed appear to happen. Unfortunate individuals who have damage in this region of the brain report that their world is frozen and they are unable to tell us about the movements of objects. There are also individuals who seem to have lost other very specific abilities, such as the ability to see colour, or depth, or faces, or emotional expressions, and many more besides. But we are getting ahead of ourselves. We shall return to these cases in the various chapters that consider these specific abilities.

See Chapter 6

Box 0.2

Different senses

Figure 0.2.1 shows a table of differences between the senses. These differences are due to the different properties of the physical 'messengers' and the design of the sense organs to capture these most efficiently. You will see that different senses are good for different things. Sound can travel around corners—light cannot. So we can use sound to tell us that there is something hidden behind a bush. Light is particularly good at telling us about the detail of something which is far away—it travels fast and therefore there is normally little delay between when it leaves the object and arrives at our eye, unless we're looking at some distant galaxy, in which case the delay could be anything up to a few billion years. But it requires an **imaging system** to make sure that different rays hit different receptor cells—otherwise you can't distinguish different rays, and you would lose information about detail. That may not be too tragic—you can move around even if your vision is very blurry, and simple animals like limpets manage without any sharp image. But if

Box continues . . .

you want to spot something small or a long way away, you need good eyes; humans and hawks have extremely well developed eyes that allow them to see things far away. So, how good your eye needs to be at resolving fine detail depends on (a) how far away you need to see stuff, and (b) whether you're sophisticated enough to have a brain that can handle all the information the eye sends it. In general, the larger the creature, the faster it can move, and the bigger its brain can be. Therefore, large creatures tend to have well-developed eyes and brains, and good ability to resolve fine detail. Remember Little Red Riding Hood: 'What big eyes you have, Grandma', she says to the wolf, who replies 'All the better to see you with'. If only she had asked 'What colour is my riding hood, Grandma?', she might have discovered the truth—wolves are red–green colour blind. Nocturnal animals, such as cats and owls, need to catch what little light there is if they are to use vision successfully; therefore they have big eyes, but their ability to see fine detail isn't very good. Hence the expression 'blind as a bat'. Many of these animals rely on other senses such as hearing, ultrasound, or smell.

	Vision	Hearing	Touch	Smell	Taste
Physical messenger—what carries the information?	Light	Sound	Surface shape	Chemicals	Chemicals
Distance—how far away does the sense tell us about stuff?	Close + far	Close + far	Close	Close + far	Close
Spatial detail—how well can we resolve fine detail in the scene?	High detail	Poor detail	High detail	Poor detail	Poor detail
Time detail—how well can we resolve whether things are changing fast in scene?	High detail	High detail	Medium detail	Poor detail	Poor detail
Does the sense start and stop rapidly?	Yes	Yes	Yes	No—smell lingers	No—have aftertaste

Figure 0.2.1 Some properties of our major senses.

So we have a simple notion that each of these areas is devoted to some aspect of our visual sense. This simple idea serves as a good starting point for considering the various aspects of vision, but it is probably far too simple. After all, if the shape of a London bus is processed in one part of the brain, its red colour in another part, and its speed in yet another, how do we put all the bits back together again to create the percept of the big red bus speeding towards us? Fortunately, the various visual areas of the brain are actually richly interconnected and they must work in concert to produce

'vision', but the details of how this miracle is achieved remain one of the great challenges to neuroscience.

The study of vision

The study of vision is a multidisciplinary affair. In this book you will come across experiments and ideas drawn from anatomy, ophthalmology, biology, psychology, physiology, physics, neuroscience, neuropsychology, and even computer science and philosophy. Indeed, we hope that the readers of this book will also come from a wide variety of such disciplines. This has advantages and disadvantages. In order to really understand vision we believe that all these approaches are required—each brings with it a different level of explanation. However, mastery of any one of these disciplines is difficult enough without having to know them all. The study of vision has many practical implications that we have already hinted at. One can see the obvious medical benefits that study of disorders of vision can have, and deeper understandings of problems like glaucoma and amblyopia are of benefit to many millions around the world. Vision is also a microcosm of the rest of the brain and the findings here serve to inform many other problems, such as how we can recover from brain damage and how to build prosthetics to take the place of a damaged system. We have also mentioned that the findings can be used to make our world a better place. The design of aircraft cockpits and cars, of road-markings and fire-escape exit signs, of advertisements and warnings can all benefit immensely from a knowledge of what people can see, where they look, and how it all works. In our world of TV and virtual reality, being able to present a person with the information that is needed through their visual sense is vital. Our knowledge of what is used (and what is not used) for a task or a particular percept allows us to compress images so that the minimum amount of information can be sent without the quality of the image being compromised.

Therefore, the study of vision has much to offer from an applied viewpoint. However, it can also be studied simply because it is there, in front of us, almost every wakeful moment—and so seems to dominate our conscious lives.

FURTHER READING

There are many, many books on visual perception that will give you far more detail on the subject than this book does. One of the best is Palmer (1999), which provides an account of vision from a psychological perspective. Hubel (1988) provides a more physiological perspective. Many find the book by Tom Cornsweet (1970) very easy to read and informative. For coverage of the 'higher' aspects of vision, the books by Martha Farah (2000) and by Milner and Goodale (2006) are excellent. Finally, recent research has begun to look more at how vision is used, as well as how it is done, and the book by John Findlay and Iain Gilchrist (2003) provides a perspective from this angle.

For papers on specific issues:

Illusions Some of the classic old illusions are well described by Ittelson and Kilpatrick (1951) and Gregory (1968). For a more up-to-date approach with more reference to the underlying neural structures see Eagleman (2001).

Multiple visual areas For an overview of these areas in both humans and non-human primates see Tootell et al. (1996).

POSSIBLE ESSAY TITLES, TUTORIALS, AND QUESTIONS OF INTEREST

1. Do you see what I see?
2. When is an illusion an illusion?
3. Consider two very different animals of your choice. How and why do their visual systems differ?
4. Why do we have many different areas of the brain devoted to vision and not just one large one?

REFERENCES

Cornsweet, T.N. (1970). *Visual perception.* London: Academic Press.

Denton, G.G. (1980). The influence of visual pattern on perceived speed. *Perception*, **9** 393–402.

Eagleman, D.M. (2001). Visual illusions and neurobiology. *Nature Reviews Neuroscience* **2**, 920–6.

Farah, M.J. (2000). *The cognitive neuroscience of vision.* Malden, MA: Blackwell.

Felleman, D.J. and Van Essen, D.C. (1991). Distributed hierarchical processing in primate cerebral cortex. *Cerebral Cortex* **1**, 1–47.

Findlay, J.M. and Gilchrist, I.D. (2003). *Active vision : the psychology of looking and seeing.* Oxford: Oxford University Press.

Gregory, R.L. (1968). Visual illusions. *Scientific American* **219**, 66–76.

Hubel, D.H. (1988). *Eye, brain and vision.* New York, NY: Scientific American Press.

Ittelson, W.H. and Kilpatrick, F.P. (1951). Experiments in perception. *Scientific American* **185**, 50–5.

Milner, A.D. and Goodale, M.A. (2006). *The visual brain in action* (2nd edn). Oxford: Oxford University Press.

Palmer, S.E. (1999). *Vision science: photons to phenomenology.* Cambridge, MA: MIT Press.

Shepard, R.N. (1990) *Mind sights: original visual illusion, ambiguities and other anomalies.* New York, NY: W.H. Freeman.

Tootell, R.B.H., Dale, A.M., Sereno, M.I., and Malach, R. (1996). New images from human visual cortex. *Trends in Neuroscience* **19**, 481–9.

Yarbus, A.L. (1967). *Eye movements and vision.* New York, NY: Plenum Press.

The first steps in seeing

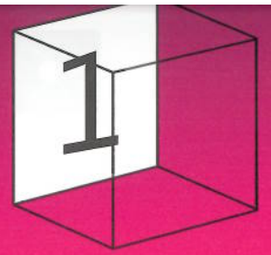

Which animals can you identify from their eye alone (answers below)? Eyes appear to have evolved on many occasions and come in a variety of forms (the most obvious distinction being between the single eyes such as our own and the compound eyes of such creatures as insects). However, you might also note that even within a type of eye there are many variations. For example, in single eyes the shape of the pupil varies greatly. Our human pupils are round, but those of the cat are more oval to allow it to close to a smaller size than ours. In other animals (e.g. geckos) the pupil is a slit that when closed becomes just three small holes in a line and thus can shut out far more light than our own.

Answers (from left to right, top to bottom): (a) robber fly, (b) jumping spider, (c) ogre-faced spider, (d) Gigantocypris (deep-sea crustacean), (e) horsefly, (f) scallop, (g) Histioteuthis (squid), (h) Dilophus (fly), (i) Tridacna (giant clam).

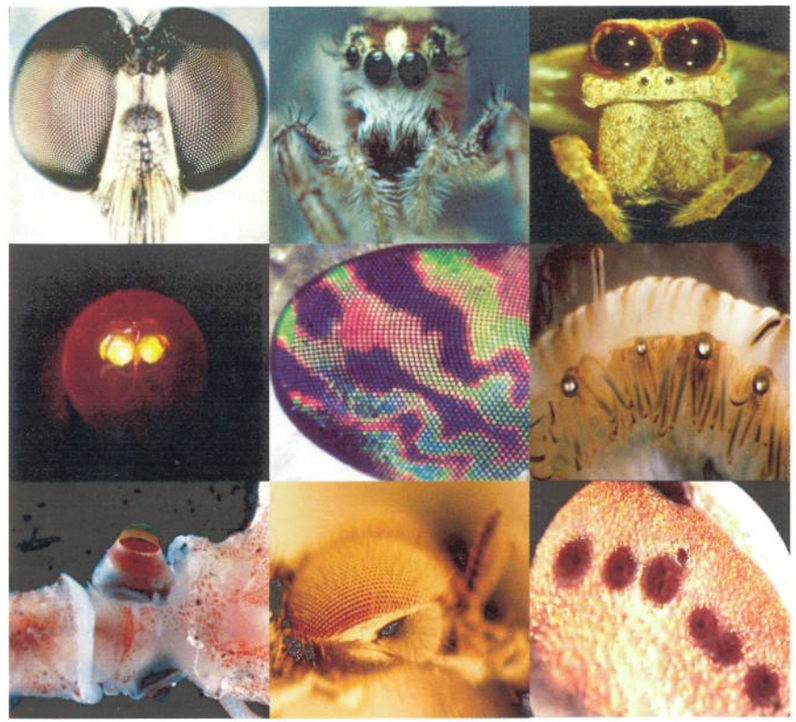

CHAPTER OVERVIEW

Even without having read any of this book, it seems likely that you have realized that the eyes are rather important for seeing. Close them and vision disappears. But what actually goes on inside these squishy little spheres? Early scientists cut a hole in the back of a bull's eye and were astonished to find that the eye produced a small upside-down image of the world. They concluded that they had solved 'vision'. Unfortunately, they had not (which means lots more reading for you!). After all, who is looking at this upside-down image? If we have a little man in our head looking at the image, how does he see? Does it matter that the image on the back of our eye is upside down? In this chapter, we explain what the eye actually does and what happens when things go wrong with it. It turns out that the back of the eye, called the **retina**, is a far outpost of the brain; it is an intricate web of nerve cells that turns the light entering the eye into a set of electrical signals. These messages begin a remarkable journey that turns light from the world into the phenomenon we experience as vision.

The eye

There is good evidence that the eyes are important for vision. Close your eyes and you can see rather little. Poke your eyes out with a pointed stick and you will see even less. So if we want to trace the visual pathway, the eyes seem a good place to start. Figure 1.1 shows a front view of a human eye and Figure 1.2 shows a cross-section through it. The first thing to notice is the **cornea**, the transparent window through which light enters the eye. The cornea is curved and acts as a lens; in fact, it is the main lens of the eye. (Three-quarters of the eye's focusing power comes from the cornea, and only a quarter from the bit of the eye we call the 'lens'.) The purpose of these pieces of optical equipment called lenses is to focus light on to the retina (the cells at the back of the eye—we'll tell you about these later).

Lenses have two important characteristics: they have at least one curved surface, and they are made of stuff through which light travels more slowly than through air. This means that light gets bent when it hits the surface (just as a stick appears bent

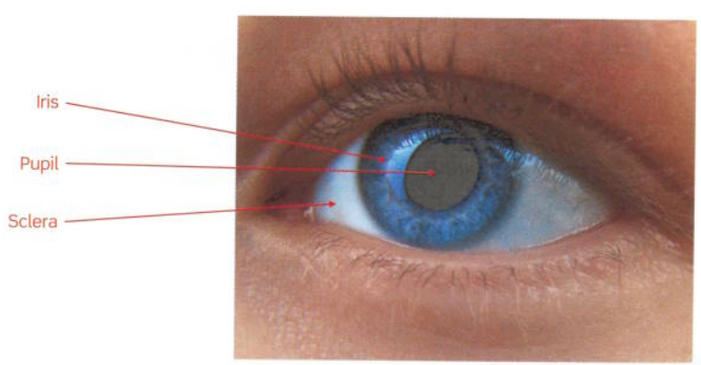

Figure 1.1 Front view of the human eye.

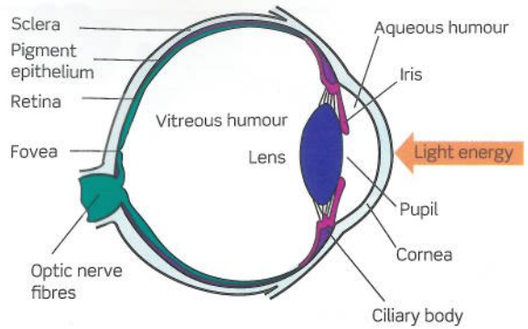

Figure 1.2 Horizontal cross-section through the human eye.

when it is half in water). Figure 1.3 shows why. Imagine you are driving a car along a tarmac road beside a sandy beach (position A). All four tyres have an even grip on the road and the car travels in a straight line. When the right-hand front wheel hits the sand (position B) the car will turn towards the right as the left-hand wheel will still be gripping the tarmac while the wheel in the sand will lose traction and move forward more slowly. Once the car is completely in the sand (position C) it will once again travel in a straight line, much more slowly, but its direction of motion will have been bent by the tarmac–sand edge. Exactly the same thing happens when light hits the cornea. Light travels at different speeds through different materials and it goes faster through air than through your cornea. How much bending takes place depends on the nature of the materials on each side of the boundary. When you open your eyes under water, everything looks very blurry. Light travels at similar speeds through water and cornea, so little deflection of the light takes place and you are unable to focus the world on to your retina (Figure 1.4a). Wearing goggles re-establishes an air–cornea boundary and, sure enough, you can see clearly again (Figure 1.4b). Fish, which don't

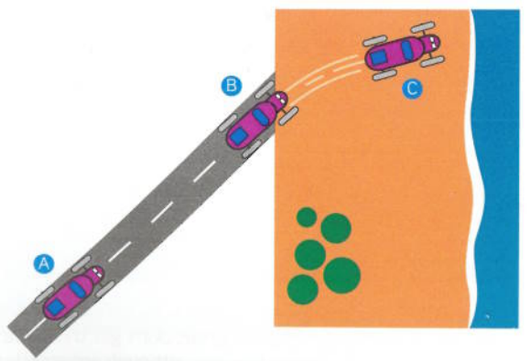

Figure 1.3 Why light is bent when it passes from one material into another (see text for details).

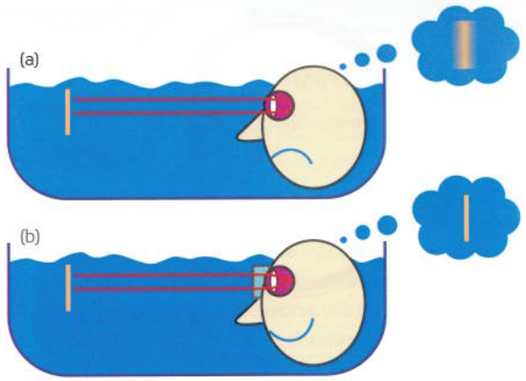

Figure 1.4 Why things look blurry when you open your eyes under water. Light travels at a similar speed in water and through the cornea, so the cornea fails to act as good lens. The solution is to wear goggles, which re-establishes an air–cornea boundary.

usually wear goggles, tend to have much more bulging eyes to compensate for the small degree of deflection between water and the lens (Figure 1.5).

Behind the cornea is the anterior chamber filled with **aqueous humour**, a watery liquid that flows from the ciliary body, through the delightfully named zonules of Zinn, through the pupil, and into the anterior chamber. From here, it passes through the trabecular meshwork and leaves the eye down (we're not joking, honest) the canal of Schlemm (Figure 1.6)—anyone trying to think of a name for a rock band need look no further. The canal of Schlemm actually runs in a circle just beneath the border of your iris (the coloured bit) and the white of your eye (called the sclera). The pressure of the aqueous humour is very important. In someone with glaucoma, a disease in which the pressure becomes too high, retinal damage and blindness can result (Figure 1.7). When

Figure 1.5 Why don't things look blurry to fish? The answer is that they have really bulging corneas which can bend light more than ours. This fish is a moray eel. Don't get this close to one and certainly don't annoy one. They have been known to attack humans, although apparently they can be quite friendly once they are used to you.

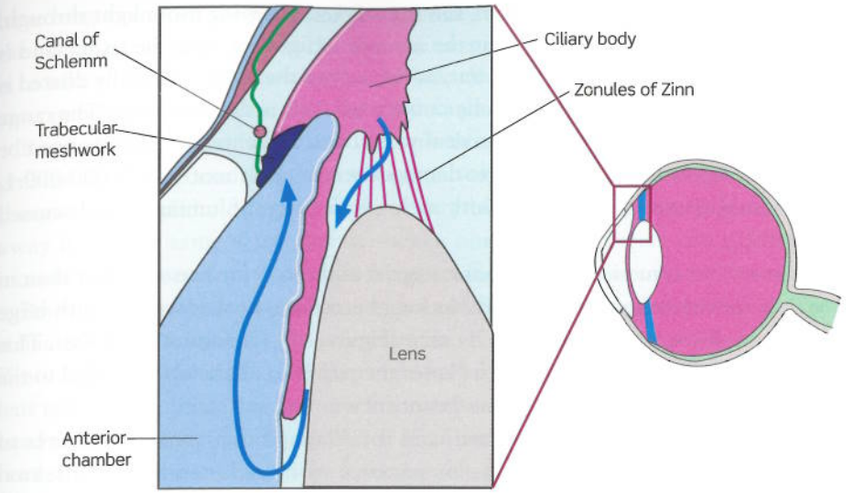

Figure 1.6 In the most common form of glaucoma the aqueous humour drains through the trabecular meshwork but gets clogged in the canal of Schlemm. This is called 'open-angle' glaucoma because the angle is clear. In 'closed-angle' glaucoma, a much rarer condition, the blockage seems to be in the trabecular meshwork.

you go the optician to have an eye examination you will probably have the pressure of your aqueous humour checked by having a puff of air blown against the cornea. This air puff is calibrated to distort the surface of the cornea by a set amount. If the cornea doesn't move enough, it suggests that the pressure is too high. In the same way, when you blow up a balloon you can check the pressure by gently squeezing the balloon to feel how much 'give' there is.

Next, we come to the **iris**, a structure that provides an adjustable aperture and is the coloured part of the eye. When light levels are high the iris constricts and the **pupil**—the hole in the middle of the iris—becomes smaller, limiting the amount of light

Figure 1.7 How the world might look to someone with glaucoma. Much of the peripheral vision may be lost without the patient realizing it.

passing through; when the light is dim, the iris relaxes, allowing more light through. This seems like a neat way of regulating the amount of light reaching the retina, and is used in cameras for the same purpose, but, as the area of the pupil when fully dilated is only 16 times bigger than when it is fully constricted, this isn't a lot of use. The range of luminance that the visual system can deal with, from the dimmest lights that can be detected up to intensities that will start to damage the retina, is about 10 000 000 000:1. How the visual system actually deals with such a wide range of luminance is discussed in Chapter 2

⇒ See Chapter 2

It turns out that sometimes our pupils constrict and dilate for reasons other than in response to the incoming light level. It has long been thought that women with large pupils are considered more attractive by men (Figure 1.8) (Demos et al., 2008). This explains why deadly nightshade, which causes the pupils to dilate when applied to the eyes, is called belladonna (which means 'beautiful woman' in Italian).

The pupils dilate when one is excited, and it is claimed that showing pictures of naked women to heterosexual men—or pictures of naked men to heterosexual women—produces pupil dilation (Hess, 1965). Of course, you get the responses you would expect from homosexual men and women as well. There is, apparently, one curious sex difference—pictures of naked sharks induce pupil dilation in men and pupil constriction in women.

Pupil dilation can, of course, reveal your excitement when you might wish to conceal it (Lubow and Fein, 1996)—hence the gambler's eye-shade to prevent his opponents gaining a clue that the poker face intends to conceal. It is claimed that jade dealers

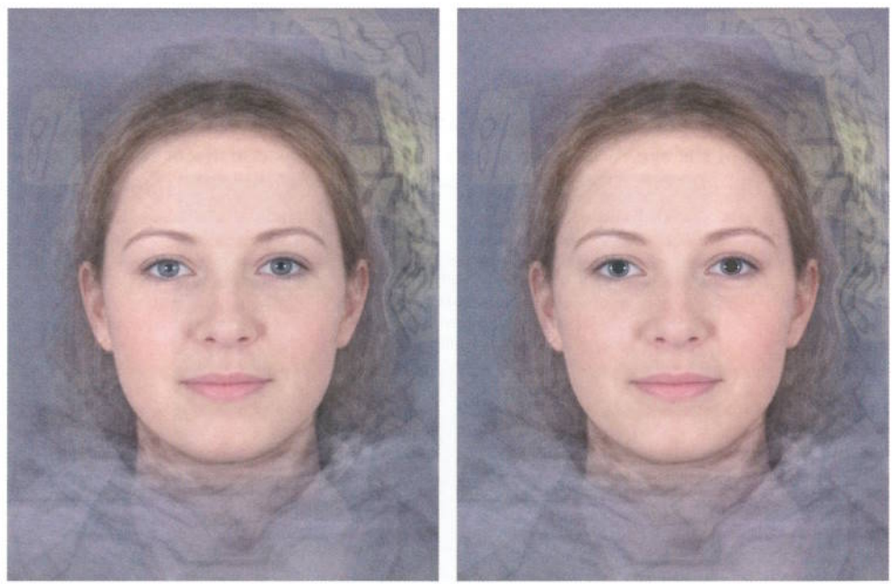

Figure 1.8 Do large pupils make a woman look more attractive? Some drugs, for example deadly nightshade (belladonna), dilate the pupils.

have long examined the pupils of prospective customers—pupil dilation will push up the asking price, in the knowledge that the customer is interested. (The solution to this is to take off your sunglasses just as the jade dealer presents you with a piece of jade; your pupils will constrict in the bright light and a bargain is assured.)

Keen photographers will know that pictures taken with a small aperture have a greater depth of focus than pictures taken with a large aperture. (One of the neat things about a pinhole camera is that everything is in focus, regardless of how far away it is.) The same is true for us—when our pupils constrict, our depth of focus increases. This is quite important as we get older, as you'll discover on the next page.

Beyond the iris we reach the **lens**. Although the lens has less power than the cornea, it has one great advantage—it is adjustable. The lens is held between the zonules of Zinn. Contraction of the circular ring of ciliary muscles relaxes the zonules so the lens gets fatter and we have more refractive power. This is what is required when we need to focus on a close object. When we want to focus on things further away, we need the lens to be stretched into a skinnier shape and this is achieved by relaxing the ciliary muscles (Figure 1.9). For this system to work correctly you need to get a few things right. First, you need to have your cornea of the appropriate curviness to focus light on the retina. Another way of looking at this is to say that you need to have your eyeball the right length for your optics. The point is that the power of the lens system must be appropriate for the size of your eyeball. Figure 1.10a shows just such an eye; the light rays from distant objects that reach the eye are near parallel and need little bending to bring them to a focus on the retina, so the lens gets pulled into a skinny shape by the ciliary muscles. Close objects send diverging rays to the eye, which need to be bent more to bring them into focus, so the ciliary muscles relax, the lens goes fat, and the light is bent more. If your eye works like this then your vision is said to be

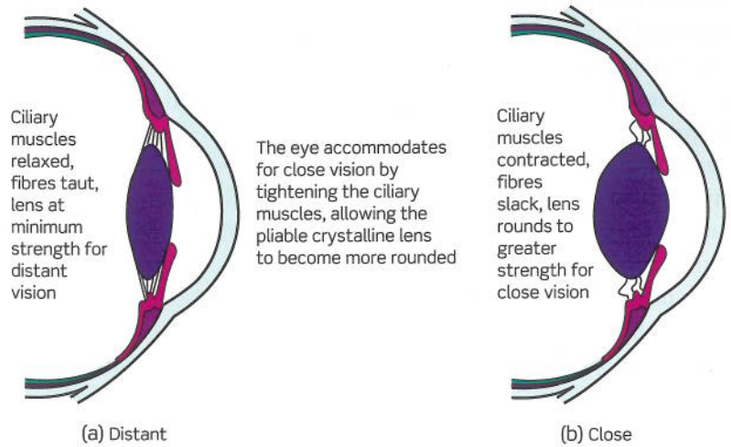

Figure 1.9 Accommodation allows the lens of the eye to become more rounded when looking at close objects by contracting the circular ciliary muscles and relaxing the zonules of Zinn: (a) distant vision; (b) close vision.

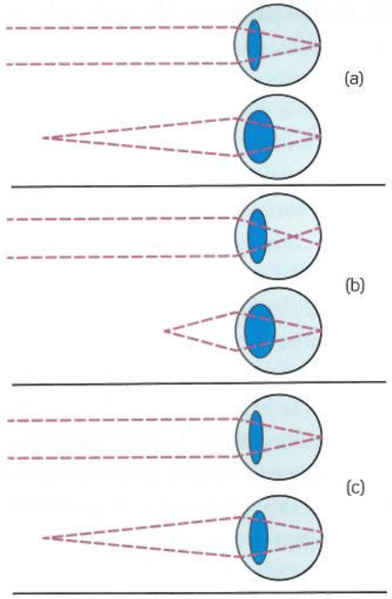

Figure 1.10 Matching the lens system to the size of your eyeball is critical: (a) emmetropic vision—both near and far objects can be focused on the retina; (b) short sight—the strong lens focuses close objects on the retina but distant objects are brought to a focus in front of the retina; (c) long sight—the weak lens is adequate to focus distant objects but is not strong enough for close objects, which are brought to a focus behind the retina.

emmetropic and you're very lucky. However, if your optics are too strong for the length of your eye (or your eye is too long for your optics—it comes to the same thing) then you are **myopic**. This is a fancy way of saying that you're short-sighted (Figure 1.10b). Distant things are blurred because they get bent too much by the optics but very close objects, which only come into focus after a lot of bending of the light, will look sharp. The opposite condition, in which the lens is too weak (or the eye too short), is called **hypermetropia**, or long-sightedness (Figure 1.10c). Distant objects that need little bending to be focused can be dealt with OK, but close objects can only be focused behind the retina.

You don't need to be a genius to see that both myopia and hypermetropia can be helped by putting lenses in front of the eyes. A hypermetrope has a weak lens and needs extra power so requires addition of a converging lens to the existing optical system. A myope needs a diverging lens to reduce the power of the optical system. These lenses can be put into frames that balance on the bridge of the nose and are held in place by struts that hook around the ears. Crude, but effective.

As we get older, the lens loses its natural elasticity (and so does much else in the body, believe us). This means that when we relax our ciliary muscles the lens does not move as much as it should and our **near-point**, the closest point we can focus on, moves away from us. This is called **presbyopia**, and it's time for reading glasses.

In Figure 1.11a we see the eye of a normal young person able to change the shape of the lens to accommodate distant and close objects. In Figure 1.11b we see what happens with age; the lens now shows little, or no, accommodation and, if distant objects are still reasonably well focused on the eye, near objects will look blurry as the ability

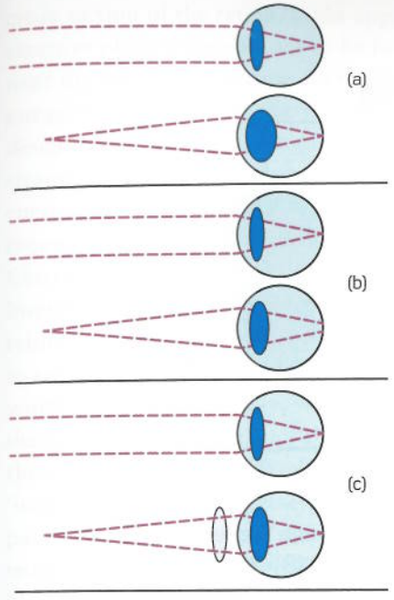

Figure 1.11 Presbyopia, the effects of ageing on our lens: (a) the emmetropic vision of Figure 1.10a; (b) inability to focus on close objects as the lens becomes inelastic. The solution (c) is to wear reading glasses.

of the lens to change shape has been lost. In Figure 1.11c we see the solution to the problem, a lens to help seeing close objects—reading glasses. The extent of the loss of accommodation is quite frightening—by the age of 30 you have lost around half of your focusing power and by 55 there is nothing much left at all (Figure 1.12).

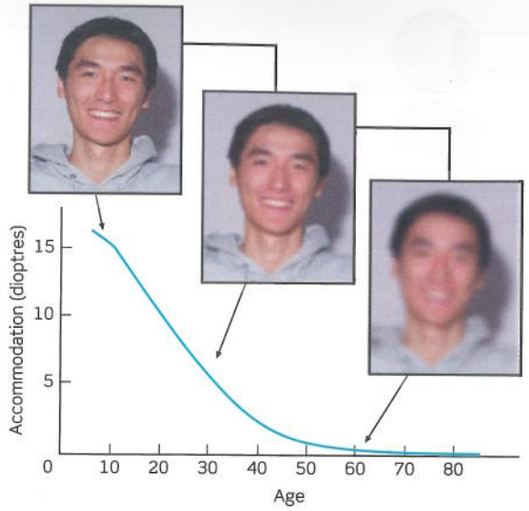

Figure 1.12 The accommodative power of the lens changes with age. The elasticity of the lens declines from birth onwards. Even fit young students are already well down the perilous slope towards reading glasses.

The main cavity of the eye, behind the lens, is filled with a gelatinous substance called the **vitreous humour**. This keeps the eyeball in shape and the retina pinned to the back of the eye. If the vitreous humour shrinks (and it does in old people), the retina can become detached and may need to be spot-welded back with a laser.

Eventually we reach the **retina**, the light-sensitive layer at the back of the eye. This is the important bit and the start of real visual processing. Figure 1.13 shows a cross-section of the retina. Light approaches from below in this view, and the light-sensitive photo-pigments are to be found in the outer segments of the **photoreceptors** near the top of the picture. So the light must travel through a good deal of neural gunge before reaching the receptors. This seems an odd arrangement—as if we designed a camera and then put the film in the wrong way round. One reason for this strange arrangement may be that the process of transduction—the turning of light energy into electrochemical energy within the nervous system—is an operation that requires a good deal of energy, and that energy must be supplied by the blood supply. Clearly, blood can be delivered more easily to the back of the eye than to the surface. Interestingly, in the octopus the receptors are at the front, rather than the back of the retina. Perhaps this shows that our own arrangement is an accident but, as we manage to see quite well despite all the gunge in front of our receptors, it was an accident not worth correcting. Actually, we still have a certain number of blood vessels that lie on the surface of the retina and, although we are normally unaware of them, we can make them strikingly visible (see Box 1.1). Notice that another penalty for having an 'inverted' retina is the presence of a blind spot in each eye where the neural bundle passes through the rest of the retina (see Box 1.2). If a creature only had one eye this would be bad news, as predators could learn to creep up on it from the blind direction. Having two eyes with blind spots pointing different ways solves this problem. Perhaps this is why, as a rule, animals have two eyes, exceptions being copopods (which are small plankton-like things and not a breakfast cereal) and the cyclops (the most famous of these, Polyphemus, ended up with a somewhat larger blind spot that he bargained for.)

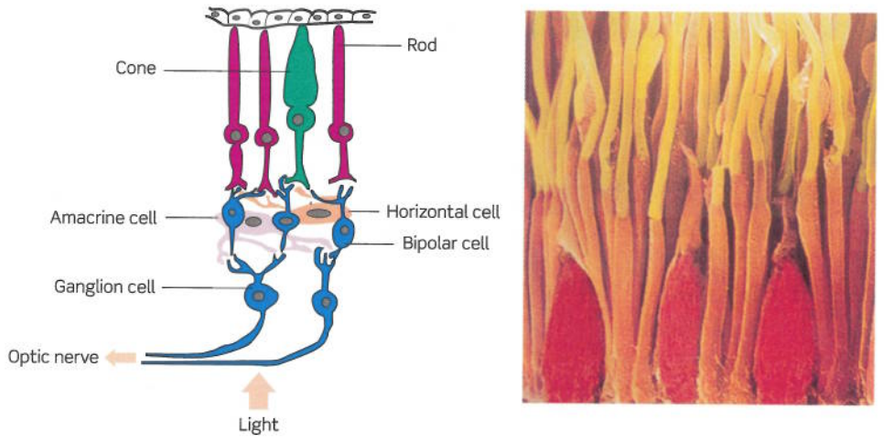

Figure 1.13 A cross-section through the retina. Note that light would come from the bottom of the figure and would have to pass through the retina before reaching the photoreceptors. The photo, a scanning electron micrograph of a primate retina, shows what the rods and cones really look like.

Box 1.1

Seeing your own blood vessels

If our retinas are covered with blood vessels that lie between our receptors and the outside world, why can't we see them all the time? There are two key reasons. First, the blood vessels don't move and the visual system is notoriously poor at detecting anything that is stationary (or 'stabilized') on the retina. Second, the light reaching the retina is sufficiently diffuse for shadows not to form on the receptors (see Figure 1.1.1). We can overcome these problems quite easily and then we can see the shadows of our own blood vessels. Follow these instructions carefully and make sure you carry out step 2 before step 3.

1. Take a piece of stiff paper or card.
2. Make a small hole in the card with a drawing pin.
3. Close one eye and raise the card to the other.
4. Jiggle the card slowly around close to the surface of the cornea while looking at a well-illuminated blank surface.

Box continues . . .

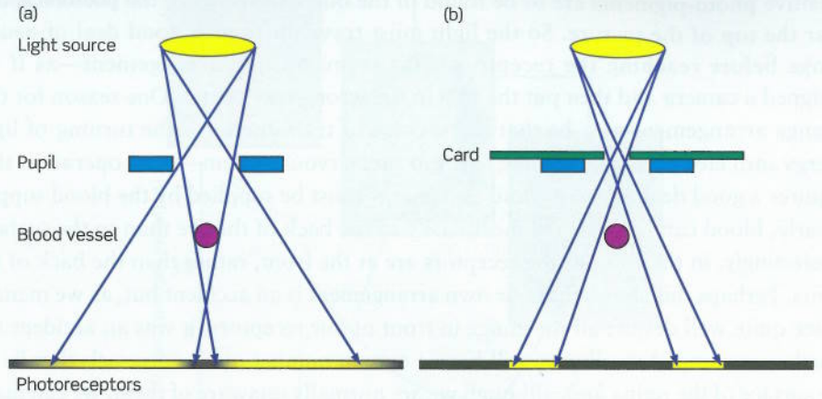

Figure 1.1.1 (a) As the blood vessels are small in comparison with the size of the pupil, and are quite distant from the photoreceptors, they cast little shadow. (b) The introduction of a small aperture—the pinhole in the card—will result in a shadow of the blood vessel falling on the photoreceptors.

You should see a pattern of shadows that resemble the veins on a leaf (see Figure 1.1.2); these are the shadows of your retinal blood vessels. Note that there is one area where there are no shadows—the point where your vision is focused. This is because there are no blood vessels covering the retina at the fovea, the area of our vision where we have our highest acuity.

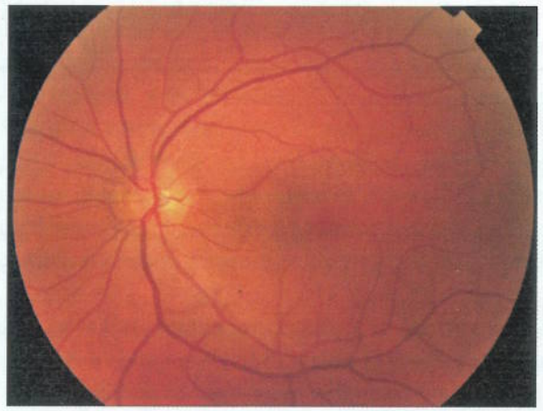

Figure 1.1.2 The pattern of blood vessels over the retina as seen through an ophthalmoscope. Note that the fovea (centre right) is free of blood vessels and the optic nerve head or 'blind spot' (centre left) is clearly visible where the blood vessels converge and leave the eye.

Box 1.2

The blind spot

Each of our eyes has a blind spot, the region where all the stuff gets in and out of the retina and therefore has no photoreceptors—and therefore no vision.

Figure 1.2.1 should help you to see your blind spot (or not see, if you see (or not see?) what we mean). Close your left eye and look at the upper green star with your right eye. One of the numbers should completely disappear! If you move your head back and forth (no sniggering) you should find that a number nearer to the star disappears when you are closer to the page. King Charles II of England (the one who hung out with Nell Gwynne, the girl with the oranges) used to play a little parlour game where he lined up the heads of people in his court with his blind spot, and thus saw them apparently beheaded. He was clearly interested in such matters since his father, Charles I, had been beheaded in a rather more permanent way.

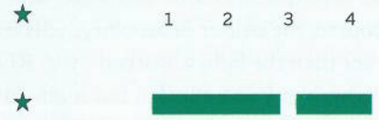

Figure 1.2.1 Seeing the blind spot.

Now move your head until the number 3 disappears and then look at the lower green star. You should find that the gap in the line disappears—but what is more it is replaced with a green line! So what is going on? It seems that some 'filling-in' process must occur. The blind spot takes on the characteristics (such as average colour and brightness) of the area around it and can even extrapolate the existence of lines. Two hypotheses may explain what is going on. The easiest explanation is that as we can't see in this spot we automatically see what is around it by default. Alternatively, there may be some active process that extrapolates what is seen around this spot and expands it to cover the blind spot.

The same lack of awareness of a blind spot occurs if, for some reason, some part of a retina stops working (as can happen as a result of retinal detachment). Although there is now an extra 'blind spot' (known as a scotoma), the person is often unaware that there is no vision in that part of the visual field. In fact, this can be rather bad news. Patients with glaucoma can lose most of their peripheral visual field, but are often unaware of this until the loss is pretty much total. How come? Nobody is quite sure but it does seem that brain processes try so hard to 'fill in' missing information that the person with glaucoma is unaware that their visual world is disappearing.

See pp. 22–23

Once the light reaches the outer segments of the receptors, neural processing can begin. The receptors are connected to **bipolar cells** and these in turn synapse with **retinal ganglion cells**. The ganglion cells are important, for it is their axons that carry information from the eye on its journey towards the visual cortex. Connecting across the retina laterally are two further cell types—**horizontal cells**, at the point where the receptors synapse with the bipolar cells, and **amacrine cells**, where bipolar cells synapse with

the ganglion cells. For our purposes, we need to know little about bipolar, horizontal, or amacrine cells, but we do need to know a little more about receptors and ganglion cells.

The photoreceptors

There are ten things you should know about the photoreceptors.

1 There are two types of photoreceptor in the human eye: **rods** and **cones**. The outer segments of the rods are generally rod-shaped and the outer segments of the cones are generally cone-shaped (Figure 1.13).

2 All our rods are basically the same; they all contain the same photopigment (the stuff that absorbs the light) in their outer segments. This is called **rhodopsin**. Rhodopsin is purple in colour and is known (not unreasonably) as 'visual purple'. Of course, the colour of an object tells us about the light reflected from it rather than the light absorbed by it. Rhodopsin reflects much red and blue light (hence the purple), but it absorbs green light preferentially.

3 Cones come in three main sorts, often (wrongly) called red, green, and blue. So-called 'red' cones contain a photopigment that is most sensitive to long wavelengths of light, 'green' cones are most sensitive to middle wavelengths of light, and 'blue' cones are most sensitive to shorter wavelengths. So we should call the cones 'long-wave', 'middle-wave', and 'short-wave', but this only complicates things so we'll call them red, green, and blue. It is because we have three cone types that we have colour vision. This is covered more fully in Chapter 5.

→ See Chapter 5

4 Rods respond very well to extremely dim light and are therefore very useful in dim conditions, i.e. at night. As the light level increases, so does activity in the rods. Because rods are so sensitive, they will be responding as much as they can when the light is still quite dim, so they are completely useless in full daylight. Cones, on the other hand (well, not really on the other *hand*), are much less sensitive, and hence are not used under dim conditions but are the ones responsible for most of our daytime vision. We have all had the experience of going into a dark cinema from a bright day and found ourselves groping in the dark before sitting in a stranger's lap. Some minutes later we can look around and see quite clearly. When we enter the dark cinema there is not enough light to get our cones going, and we have to rely on our rods. Unfortunately, our rods have just been outside in bright sunlight and are completely bleached. It takes some minutes in the dark for the rods to recover their sensitivity. When only rods are active we call our vision **scotopic**, when it is so bright that the rods cannot function we call our vision **photopic**, and the region in between when both receptor types play a role is called **mesopic**. The issue of light and dark adaptation is covered further in Box 1.3.

Box 1.3

Dark adaptation

So far, we have spun a tale about how the retinal ganglion cells can happily signal edges and are immune to the changes in overall light level. However, there is a problem in that the dimmest object that we can see, an object illuminated by dim starlight in the night sky, is about 10^{10}, or ten thousand million, times dimmer than the brightest intensity we are exposed to—the unobscured disk of the sun. Unfortunately, a typical neuron is capable of responding only over a less than 1000 to 1 range of intensities of light. So how can the visual system work over such a prodigiously large range of light intensities? We have already seen in this chapter that the eye's pupil can't be of much help because in its fully constricted state the area of the aperture is only about 16 times smaller than when the pupil is fully dilated. However, we have also learned that rods are much more sensitive to low levels of light than cones. Remember that we talked about going from a sunny day into a cinema and sitting in a stranger's lap? This was because in the dark cinema only our rods are sensitive enough to see strangers' laps, but they were yet to recover from being bleached by the bright sunshine. We can investigate this phenomenon in the lab with more precision but perhaps less excitement than sitting on strangers.

Suppose that a person looks at a very bright light for a minute or so and then is plunged into total darkness. Their task is to adjust the brightness of a small light to the point where they can just see it. The results are shown in Figure 1.3.1. In layman's terms, you can see that we need time 'to get used to the dark'; that is, initially we can only see bright lights and over time we can gradually see dimmer and dimmer lights. However, it's not quite as simple as this. The actual intensity we need to see our small spot of light is given by the purple line on the graph. As you can see, the brightness needed falls dramatically for a few minutes and then appears to level out. Just when it looks as if the threshold has settled down, after about 10 minutes of sitting in the dark, it falls rapidly once more. Indeed, in this particular example the person isn't at their most sensitive until after at least 30 minutes of darkness.

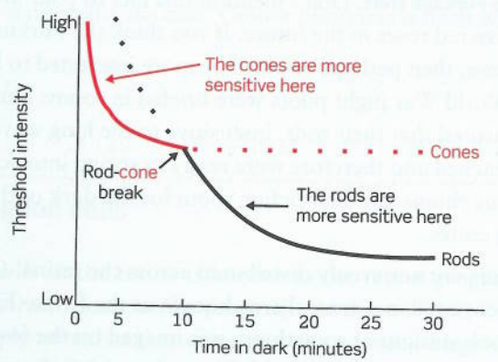

Figure 1.3.1 Dark adaptation curve. Cones recover rapidly if they are bleached, gaining their maximum sensitivity after about 10 minutes in the dark. Rods recover at a much slower rate, only gaining full sensitivity after 30–40 minutes.

Box continues . . .

We can explain the strange shape of this graph in terms of the behaviour of rods and cones. The initial bright light bleaches the highly sensitive rod photopigment as well as our cones. When we first try to detect the dim spot of light we have to rely on our cones because these recover from the bright light much more quickly than the rods. The initial drop in brightness needed to see the target spot reflects the recovery of the cones from bleaching. (The thresholds for the cones alone are shown by the green line.) These have recovered fully after about 5 minutes, so there is no great change in sensitivity for a while. However, the very sensitive rods have really been clobbered by the bright light and their recovery is much slower than that of the cones, but gradually they do recover. After around 10 minutes the rods overtake the cones to become more sensitive, and further improvement occurs until the rods recover fully after 30–40 minutes. We can confirm this interpretation of the **dark adaptation curve** (as it's known) from clever experiments where we use lights that only the cones or rods can see (and be bleached by), or by using individuals who through quirks of fate have either only rods or only cones in their retinas.

The opposite of dark adaptation is, you've guessed it, light adaptation. This is what happens when you come out of the cinema into the daylight, or when you first open your eyes on a bright sunny morning. In the darkness both rods and cones can recharge their batteries, so that as soon as we open our eyes the cones spring into life and the rods begin to be bleached.

5 Rods are most sensitive to green light, whereas the cone system overall is most sensitive to yellow light. Next time you are given a red rose in moonlight (and this will be sooner for some of you than for others) note the relative brightness of the red of the flower and the green of the foliage. You should notice that the red of the petals looks very dark and the green of the leaves looks relatively light. Next morning compare the flower and the foliage again as the rose lies in the morning sunshine filtering on to your pillow; now the red flower will look lighter than the green foliage. This is known as the **Purkinje shift**. Don't mention this fact to your loved one if you want to receive red roses in the future. If you think the Purkinje shift is just slushy nonsense, then perhaps you will be more interested to know that in the Second World War night pilots were briefed in rooms illuminated by red light. This ensured that their rods, insensitive to the long wavelengths, were not being bleached and therefore were ready to spring into action the minute Biggles and his chums left the briefing room for the dark of the airstrip and their waiting crates.

6 Photoreceptors are not evenly distributed across the retina. Cones are heavily concentrated in a central area known as the fovea (Figure 1.14). When you look straight at something, it is imaged on the fovea. Rods are completely absent from the central fovea and are most densely packed some 12–15° into the periphery (Figure 1.14). This explains why, when you look at extremely faint stars, they disappear when on the fovea but can be seen in

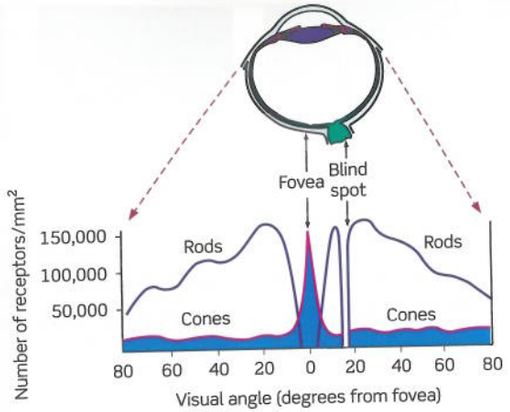

Figure 1.14 The distribution of rods and cones across the retina. Note that, although cones are most densely packed in the fovea (around 150 000/mm²), even in the periphery there are still several thousand per square millimetre.

7 There are no blue cones at all in the central fovea—blue cones are rather strange things (and are only understood by rather strange people) and somewhat unlike either red or green cones (Roorda and Williams, 1999). Interestingly, people with normal colour vision can have very different numbers of red and green cones (Figure 1.15).

8 Ganglion cell axons and blood vessels leave the eye at a point known as the **blind spot** or **optic disc**. This region, situated about 12–15° into the nasal retina, is devoid of all receptors and consequently we are blind in this area (see Box 1.2).

9 Congenitally colour-blind people have a problem with either their red or green cones. One of these cone types may be completely missing, or it may be weaker than is normally the case. Colour blindness is dealt with in detail in Chapter 5.

→ See Chapter 5

10 Er . . . that's it.

The retinal ganglion cells

The last layer of cells in the retina is known as the **retinal ganglion cells** (see Figure 1.13). We shall deal with some of the properties of these cells at length in Chapter 2. Here, we shall simply note that they come in two varieties—large and small. We shall call the large ones **M cells** (M for magnocellular) and the small ones **P cells** (P for parvocellular). These cells differ not only in their size but also in their properties. Careful study has revealed that P cells distinguish between signals coming from the red cones and those coming from the green cones. So, for instance, some P cells will be excited

→ See Chapter 2

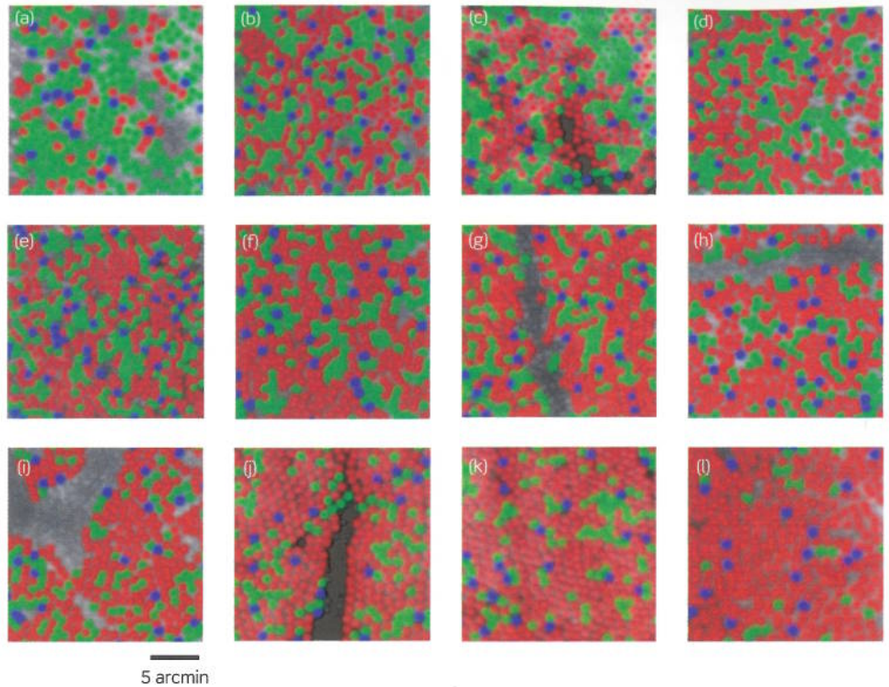

Figure 1.15 Images of the cone mosaics of 10 subjects with normal colour vision obtained with the combined methods of adaptive optics imaging and retinal densitometry. The images are false-coloured so that blue, green, and red are used to represent the blue, green, and red cones, respectively. (The true colours of these cones are yellow, purple, and bluish-purple.) The mosaics illustrate the enormous variability in red/green cone ratio. The ratios are as follows: A, 0.37; B, 1.11; C, 1.14; D, 1.24; E, 1.77; F, 1.88; G, 2.32; H, 2.36; I, 2.46; J, 3.67; K, 3.90; L, 16.54. The proportion of blue cones is relatively constant across eyes, ranging from 3.9% to 6.6% of the total population. Images were taken either 1° or 1.25° from the foveal centre. For 2 of the 10 subjects, 2 different retinal locations are shown. Panels D and E show images from the nasal and temporal retinas, respectively, for one subject; panels J and K show images from the nasal and temporal retinas, respectively, for another subject.

by the red cones and inhibited by the green cones. The M cells don't seem to care about this distinction, and mix them up. Therefore, only the P cells appear to carry the information about colour. On the other hand, the M cells appear much more suited to carrying information about the dynamic aspects of the world, such as movements and flicker. Studying these cells is quite difficult, as they are all intermingled in the retina. But one thing we do know is that both types of cell send their signals to the next visual structure, the **lateral geniculate nucleus (LGN)**, but do so in such a way that all the M cells project to one bit of the LGN and the P cells project to other bits. The grouping of these cells allows us to lesion them selectively and discover what they do, something that we can't do in the retina. Therefore, we shall pick up this story again shortly when we get to the LGN.

Beyond the eye—the optic nerve

The optic nerve leaves the eye at the blind spot and begins its journey to the cortex which, curiously, is about as far from the eyes as it's possible to get and still remain inside the head. The main pathways are shown in Figure 1.16. First, the optic nerves from the two eyes converge at a point called the **optic chiasm**. Here, a partial decussation occurs; this means that some of the fibres cross over to the other side of the brain and some don't.

The axons from ganglion cells on the nasal side of each retina cross and those from the temporal (i.e. where your temples are) side of each retina don't. The result of this is that now the left-hand bunch of nerve fibres carries information about the right-hand part of the world (the right visual field) and the right-hand bunch of nerve fibres carries information about the left visual field. At this stage, the optic nerve changes its name to **optic tract**. You can see that chopping through your optic *nerve* will blind you in one eye, whereas chopping through your optic *tract* will leave you unable to see half the world—this is known as a **hemianopia**. This crossing-over of sensory pathways is fairly general; you may know that someone who has a stroke on one side of the brain may experience problems with feeling and moving their limbs on the other side of the body.

See Chapter 11

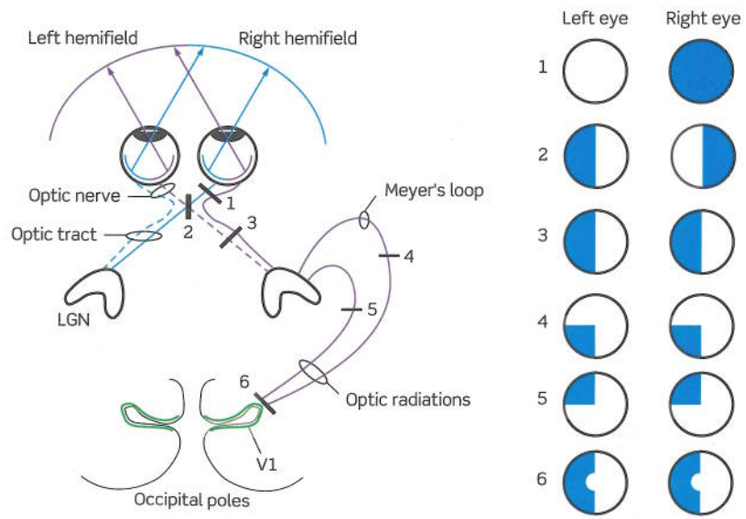

Figure 1.16 The route of the visual pathways as they go from the retina up to the visual cortex. Lesioning the pathways at the points indicated by the numbers 1–6 will produce a loss of vision as shown on the right. Vision is lost in the areas marked in blue.

The lateral geniculate nucleus

The optic tract now reaches the main relay point on the way to the cortex—the LGN. A cartoon cross-section of the LGN is shown in Figure 1.17a—to those with a particularly vivid imagination the LGN looks like a knee, hence its name (the Latin for knee is *genus*; think 'genuflect'). In primates, each LGN has six layers, with the ganglion cell axons from one eye terminating in three layers and the axons from the other eye terminating in the other three. Thus, information from each of the two eyes, though represented in both LGNs, is kept segregated. Look at Figure 1.17b, which shows a photomicrograph of the left LGN; layers 1 and 2 (at the bottom) are different from the others because the cells here are big, whereas the cells in the other layers are small. Hence, layers 1 and 2 are known as the **magnocellular layers** and layers 3, 4, 5, and 6 are known as the **parvocellular layers**. Not surprisingly, the M ganglion cells that we discussed earlier send their nerve fibres to the magnocellular layers, and the P ganglion cells to the parvocellular layers.

We need two magnocellular layers in each LGN because one receives input from the left eye and one from the right eye. Layer 1 of the left LGN receives its input from the right eye and is known as a **contralateral** layer, whereas layer 2 of the left LGN receives its input from the left eye and is known as an **ipsilateral** layer. When we come to the parvocellular layers, layer 3 is an ipsilateral layer and layer 4 is a contralateral layer. Fine, so far, but why do we need four parvocellular layers? We don't know. Layer 5 (ipsilateral) and layer 6 (contralateral) appear to be just the same as layers 3 and 4, respectively, so they seem to be redundant. So, moving from layer 1 to 6, the order is contra, ipsi, ipsi, contra, ipsi, contra. This seems a bit odd, and is certainly hard to remember. At least two very well known authors get the order wrong in their books, but it would be ungentlemanly to identify them (Hubel, 1988, p.66; Wandell, 1996, Figure 6.3).

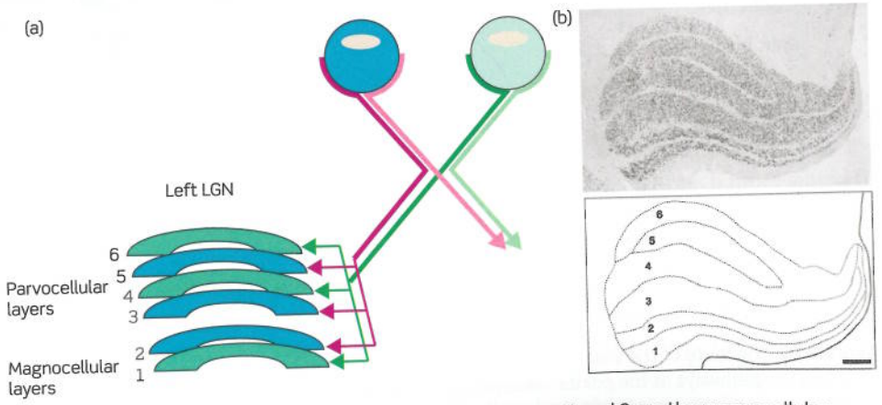

Figure 1.17 The LGN has six layers, three from each eye. Layers 1 and 2 are the magnocellular

In each of the LGN's six layers the cells retain what is called **retinotopic mapping**. This means that ganglion cells adjacent in the retina (and therefore also from adjacent directions in the visual world) will project to cells adjacent in the LGN. Hence, it forms an orderly map of the visual world where adjacent cells receive information from adjacent parts of the image until all the image is covered. So, each LGN has six maps of the world (one in each layer), one on top of the other. Of course, each map is of only half the visual world, the left LGN representing the right visual field and the right LGN the left visual field. The layers of the LGN have been described by one of the greatest visual neurophysiologists, David Hubel, as being like a club sandwich. This is a very useful analogy for many British vision scientists who previously had little idea of what a club sandwich looked like.

As we discussed earlier, the M and P divisions of the LGN receive their input from very different cell types in the retina. Remember that the M retinal cells seemed to carry information about movement and flicker, and the P retinal cells carry colour information. Do the M and P cells of the LGN have similar properties to the M and P cells of the retina? The answer seems to be yes. The experiments that have given us this answer selectively damaged (known as lesioning) either the M divisions or the P divisions alone and tested what happened to the vision of the animal. In one experiment, the scientists (Schiller et al., 1990) trained animals to fixate on a central point and then move their eyes to the 'odd one out' on a screen (see Figure 1.18). They then took advantage of the retinotopic coding in the LGN by placing very small lesions in either the M or the P division. These small lesions produce only a small area of visual field where one might expect a loss of vision (according to where the map was damaged (see Figure 1.19))—the rest of the visual field was unaffected. They could now see how well the animal performed when a target stimulus was placed in the affected

→ See Chapter 12

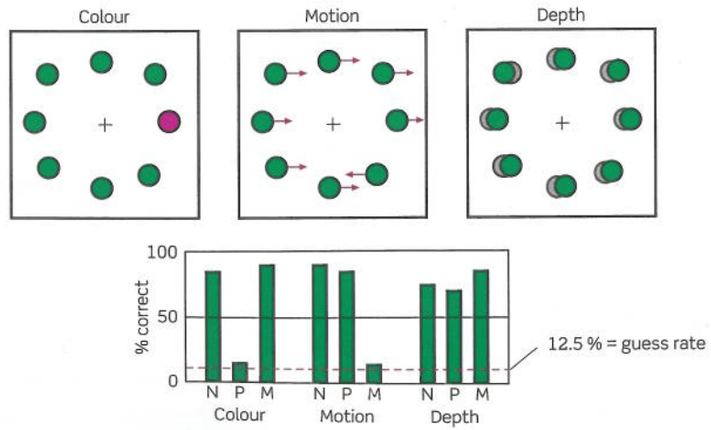

Figure 1.18 Schiller et al.'s experiment. Monkeys lesioned in the parvocellular layers have impaired colour discrimination, but motion and depth discrimination remain intact. Magnocellular lesions impair motion, but not colour or depth discrimination.

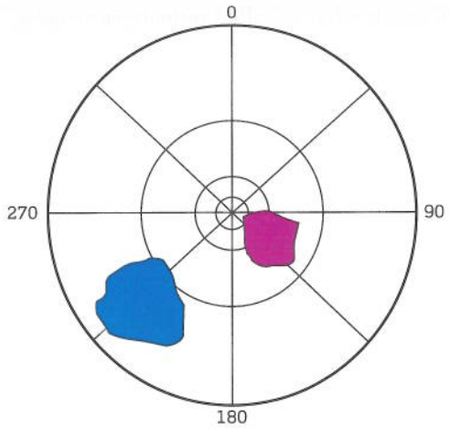

Figure 1.19 The lesions in Schiller et al.'s experiment were restricted to either parvocellular (purple) or magnocellular (blue) areas of the LGN. The areas of the visual field affected by the lesions revealed the effects of losing one pathway or the other.

field and compared the performance with how well it did when the target was in an unaffected field.

They found that the lesions affected different targets in different ways. Lesions to the P stream destroyed the animal's ability to detect targets defined by colour, by fine-detailed texture, or by a subtle change in shape, whereas targets defined by movement or flicker were impaired by lesions in the M division. Interestingly, other targets, such as those defined by a gross change in shape or by stereopsis (depth derived from having two eyes) survived either lesion, suggesting that information needed for these tasks could be carried by either division. From these, and related experiments, we get the picture that the M stream carries information about coarse features and movement, whereas the P stream can carry information about finer features and colour. A nice way to remember this, or even visualize it, is that if we only had our M stream the world would look as if we were watching a rather badly tuned black-and-white TV—the detail and colour are lost. If we only had our P stream it would be like seeing a picture postcard—the colour and detail are there, but the movement is lost.

See Chapter 7

Until fairly recently, this story about the M and P streams seemed complete. However, it turns out that there are other cells, mainly situated between the layers, that are also important for vision. These small cells are known as the **koniocellular cells** (meaning 'sand-like') or **K cells**. There are about 100 000 of them in your LGN (about the same number as the M cells) and their input comes from a special type of ganglion cell in the retina. This type of retinal ganglion cell has a major input from the blue cones, hence it is believed that the K pathway is heavily involved in a form of colour vision that is based on a blue–yellow comparison, whereas the colour information carried in the P pathway is based on a green–red comparison. However, we shall save the details of this exciting story for Chapter 5.

See Chapter 5

So, the properties of the LGN cells look very much like the properties of the retinal ganglion cells. The obvious question to ask now is: what is the LGN doing, then? Unfortunately there is, as yet, no clear-cut answer, but, intriguingly, the strongest input

into the LGN is not from the retina but from the cortex itself—the very area to which the LGN sends its output. That is, the biggest input to the LGN comes 'top-down' rather than 'bottom-up'. This has led to the idea that the LGN might be important in filtering what information gets through to the cortex—in this analogy the LGN serves as the spotlight of attention, highlighting information coming from certain bits of the visual field. There is now some actual evidence for this from a study (O'Connor et al., 2002) using functional magnetic resonance imaging (fMRI), a technique which is becoming increasingly important in vision science. We describe fMRI in more detail in Chapter 12. Information from the LGN is sent on to the cortex itself, but that's another story and one that will have to wait until Chapter 3 (Andrews et al., 1997).

See Chapter 12
See Chapter 3

Box 1.4

Dyslexia and the magnocellular pathway

People with dyslexia have failed to learn to read efficiently, although other functions are preserved. For instance, dyslexics often jumble the letters in a word and could therefore spell 'dyslexia' as 'dailysex' (this may also cause dyslexia, but the authors have insufficient experience in this matter). If you have ever gone to a toga party dressed as a goat, there's a good chance you are dyslexic.

Recently, it has been suggested that some of these reading problems might occur because of damage to the magnocellular pathway. The LGNs of dead dyslexics (rest assured that they departed from natural causes rather than from an over-zealous vision scientist) have been found to have magnocellular layers much smaller than those from normal readers (who were also dead).

Box continues . . .

As the magnocellular pathway seems to be heavily involved in our perception of motion, this suggests that dyslexics should also be poor at motion perception tasks. Recent work has measured various aspects of motion perception in dyslexics and has shown such a deficit. For instance dyslexics need a much greater difference between the speeds of two patterns in order to spot this difference—and their motion area of the brain (area MT) shows reduced activation in brain imaging studies. Perhaps you should be wary of having a dyslexic as your designated driver!

Although this theory seems attractive, it is not without its problems. Why should a problem in the magnocellular system produce a selective deficit in reading? One idea is that the magnocellular system is important in refreshing the image each time your eyes move and, if this replacement of one image by the next is not working properly, then successive images from one fixation to the next become confused—and reading is difficult. But why then should dyslexics find spelling difficult? A second problem is that for every experiment that has found a visual deficit in dyslexics there seems to be another reporting no difference (whether this is due to M-cell function or not); hence the presence of such deficits is not yet fully established. Thirdly, if a test only measures one aspect of vision (e.g. motion perception) and finds a deficit, how do we know that other aspects of vision are not also affected? Perhaps dyslexics are just no good at vision tasks in general.

See also Chapter 6

See Chapter 6

FURTHER READING

Most ophthalmology books will give you more detailed information about the eye itself and about problems such as short-sightedness, but we find Davson (1990) nice and straightforward. The retina will be covered in most textbooks but the book by Dowling (1987) is a delight and well worth a look if you need to know what all those other cells that we ignored do, or how the receptive fields of the ganglion cells are constructed. The LGN seems to receive scant coverage in most books, perhaps reflecting our puzzlement as to just what it is doing. For many fascinating insights into the eyes of many other creatures, see Land and Nilsson (2001).

For papers on specific issues:

Evolution of the eye How our eyes (and those of other animals with very different eyes) came to be the way they are is covered by Gregory (1967, 1997). The evolution of the eye has often been held up by creationists to be impossible. For a view on just how easy it could have been, see Nilsson and Pelger (1994) and Dawkins (1995).

Pupil dilation, etc. For a series of fascinating insights into when and how our pupils change size, see the classic work of Hess (1965). More recently, pupil size has been used to try to detect deception too (Lubow and Fein, 1996).

Rods and cones For the actions of rods and cones, Dowling (1987) is excellent. For a little more detail also see Daw et al. (1990).

M, P, and K cells The lesioning studies of the M and P divisions of the LGN were performed by Schiller et al. (1990). However, what appear to be very contradictory results can be found in Livingstone and Hubel (1988). See also Chapter 11, where we follow these pathways as they reach deep into the cortex. For information on the koniocellular pathway, see Hendry and Reid (2000).

See Chapter 11

What LGN cells might actually be doing There have been many theories that have tried to give a specific function to the LGN (rather than just passing on information from the retina to the cortex). One of the most recent and interesting is that of Sherman (2001).

Dyslexia and visual problems There has been an explosion of interest in the idea that a specific visual problem might underlie the problems of a person with dyslexia. There are many papers on this still controversial area. To get a flavour of the evidence for and against, see Stein et al. (2000) and Skottun (2000), respectively.

POSSIBLE ESSAY TITLES, TUTORIALS, AND QUESTIONS OF INTEREST

1. How does our vision differ between day-time (photopic) and night-time (scotopic)?
2. What can we learn from the size of someone's pupils?
3. How and why do the retinas of animals change with their lifestyles?
4. Can dyslexia be explained by a selective loss of M-cell function?

REFERENCES

Andrews, T.J., Halpern, S.D., and Purves, D. (1997). Correlated size variations in human visual cortex, lateral geniculate nucleus and optic tract. *Journal of Neuroscience* **17**, 2859–68.

Davson, H. (1990). *Physiology of the eye* (5th edn). London: Macmillan.

Daw, N.W., Jensen, R.J., and Brunken, W.J. (1990). Rod pathways in mammalian retinae. *Trends in Neurosciences* **13**, 110–15.

Dawkins, R. (1995). Where d'you get those peepers? *New Statesman and Society* **8**, 29.

Demos, K.E., Kelley, W.M., Ryan, S.L., Davis, F.C., and Whalen, P.J. (2008). Human amygdala sensitivity to the pupil size of others. *Cerebral Cortex* **18**, 2729–34.

Dowling, J.E. (1987). *The retina*. Cambridge, MA: Harvard University Press.

Gregory, R.L. (1967). Origin of eyes and brains. *Nature* **213**, 369–72.

Gregory, R.L. (1997). *Eye and brain* (5th edn). Oxford: Oxford University Press.

Hendry, S.H. and Reid, R.C. (2000). The koniocellular visual pathway. *Annual Review of Neuroscience* **23**, 127–53.

Hess, E.H. (1965). Attitude and pupil size. *Scientific American* **212**, 46–54.

Hubel, D.H. (1988). *Eye, brain and vision*. New York, NY: Scientific American Press.

Land, M.F. and Nilsson, D.-E. (2001). *Animal eyes*. Oxford: Oxford University Press.

Livingstone, M.S. and Hubel, D.H. (1988). Segregation of form, color, movement and depth: anatomy, physiology and perception. *Science* **240**, 740–9.

Lubow, R.E. and Fein, O. (1996). Pupillary size in response to a visual guilty knowledge test: new technique for the detection of deception. *Journal of Experimental Psychology: Applied* **2**, 164–77.

Nilsson, D.-E. and Pelger, S. (1994). A pessimistic estimate of the time required for an eye to evolve. *Proceedings of the Royal Society of London, Biological Sciences* **256**, 53–8.

O'Connor, D.H., Fukui, M.M., Pinsk, M.A., and Kastner, S. (2002) Attention modulates responses in the human lateral geniculate nucleus. *Nature Neuroscience* **5**,1203–9.

Roorda, A. and Williams, D.R. (1999). The arrangement of the three cone classes in the living human eye. *Nature* **397**, 520–2.

Schiller, P.H., Logothetis, N.K., and Charles, E.R. (1990). Functions of the colour-opponent and broad- band channels of the visual system. *Nature* **343**, 68–70.

Sherman, S.M. (2001). Tonic and burst firing: dual modes of thalmocortical relay. *Trends in Neurosciences* **24**, 122–6.

Skottun, B.C. (2000). The magnocellular deficit theory of dyslexia: the evidence from contrast sensitivity. *Vision Research* **40**, 111–27.

Stein, J., Talcott, J., and Walsh, V. (2000). Controversy about the visual magnocellular deficit in developmental dyslexics. *Trends in Cognitive Sciences* **4**, 209–11.

Wandell, B. (1996). *Foundations of vision*. Sunderland, MA: Sinauer.

Signalling changes

Micrograph of a slice through the monkey retina, showing the fovea. Light entering the eye would arrive from the top of the picture. The long thin photoreceptors (the rods and cones) can be seen in the light-coloured layer near the bottom. Three distinct bands of cell bodies can be seen. The dark layer just above the cones, the **outer nuclear layer**, contains the cell bodies of the cones. The next dark layer, the **inner nuclear layer**, contains the cell bodies of the horizontal, bipolar, and amacrine cells. The top dark layer contains the ganglion cells. Note that in the central fovea the cones are close to the surface of the retina as the layers of cells above them are pulled aside to form the **foveal pit**.