

14

How does the timing of neural signals map onto the timing of perception?

DAVID M. EAGLEMAN

Summary

Different features of stimuli are processed in the central nervous system at different speeds. However, such neural time differences do not map directly onto perceptual time differences. How the brain accounts for timing disparities to correctly judge the temporal order of events in the world is the temporal binding problem. I weigh physiological data against new psychophysical findings both within and between modalities. The essence of the paradox is that the timing of neural signals appears, at first blush, too variable for the high accuracy of the psychophysical judgments. I marshal data indicating that ~ 80 ms is an important duration in perception and make the novel suggestion that this number is directly mirrored in the physiology. In recordings from several areas of the primate visual system, the difference between the slowest and fastest latencies based on luminance contrast is 80 ms. If the rest of the brain wants to time outside stimuli correctly, it must account for the fact that the earliest stages of the visual system spread signals out in time. I suggest that the brain waits for the slowest information to arrive before committing to a percept. This strategy only applies to visual awareness; in contrast, the motor system may form its reactions based on the first incoming spikes.

14.1 Introduction

One goal of modern neuroscience is to relate physiological data to perception (Eagleman 2001). How do spikes recorded from single neurons map onto object recognition, brightness perception, or timing judgments? Despite decades of work, there are few good theories uniting the wetware to the perception, and we currently cannot build machines that are visually aware. This chapter attempts to ferret out some relationships between the two domains that can guide our search. My strategy is to come at the problem from the point of view of time. If we can find parallels in the temporal relationships between the physiology and the psychophysics, this may open new inroads into their interaction.

Nervous systems face the problem of feature binding – that is, keeping features of an object perceptually united such that, for example, the redness and squareness do not bleed off a moving red square (Crick & Koch 1990; Engel et al. 1992). The fact that feature binding is usually performed correctly would not come as such a surprise were it not for our modern picture of the mammalian brain, in which different kinds of information are processed in

14 How does the timing of neural signals map onto the timing of perception? 217

different neural streams. Binding requires coordination not only among different modalities (vision, audition, olfaction, etc.) but also among different features within a modality (in vision, for example, color, motion, form identification).

But there is a deeper challenge the brain must tackle, without which feature binding would rarely be possible. This is the problem of *temporal* binding: the assignment of the correct timing of events in the world. The challenge here is that different stimulus features move through different processing streams – and *are processed at different speeds*. I will draw on electrophysiological data below to demonstrate this point. The brain must account for speed disparities between and within its different sensory channels if it is to accurately determine the timing relationships of features in the world.

Many discussions of neural function tacitly rest on the assumption that awareness (what the subject reports) is an online phenomenon, coming about as soon as the leading edge of the represented stimulus reaches a “perceptual end-point” (Zeki & Bartels 1998). However, many demonstrations in the literature suggest that awareness is not an online phenomenon but is instead delayed (Kolers & von Grunau 1976; Dennett 1991; Kinsbourne 1993; Bachmann 1994; Pessoa et al. 1998; Eagleman & Sejnowski 2000a,b,c, 2003, 2007; Eagleman 2008). That is, perceptions are retrospectively attributed after the brain has gathered information from a window of time around an event. I will argue here that this postdictive framework is a necessary component to any theory of awareness and of its necessity to solve the temporal binding problem. I will then discuss what this illuminates about the physiology.

In this chapter I draw mainly on data from the visual system, but the framework proposed here applies to all modalities, and I will give examples where available. I will sketch the relevant neurophysiology of the visual system, illustrating that signals reflecting different stimulus features move through the system at different times. This leads us to ask how visual perception ever gets the timing correct. I will ask why some optical illusions exist and why others do not. Finally, I will suggest that temporal binding can be accomplished via a window of delay that allows for more slowly processed information to participate in the interpretation. In other words, the brain waits to collect the slowest signals. I draw on neurophysiologic measures to suggest that this window should be about 80 ms. This window of delay means that awareness is postdictive, incorporating data from a window of time both before and after the event and delivering a retrospective *interpretation* of what happened (Eagleman & Sejnowski 2000a,b,c; Eagleman 2008). Postdiction is the opposite of prediction, which is the act of extrapolating into the future to guess about events that have not yet happened. I conclude with a discussion of how signals can be delayed and aligned in neural tissue.

14.2 Physiology of the visual cortex

A measure of increasing importance to physiologists is the *latency* of a neuron’s response: how much time passes between the onset of a stimulus and a cell’s first measurable response to it. Figure 14.1 shows that latencies are surprisingly variable across different parts of the

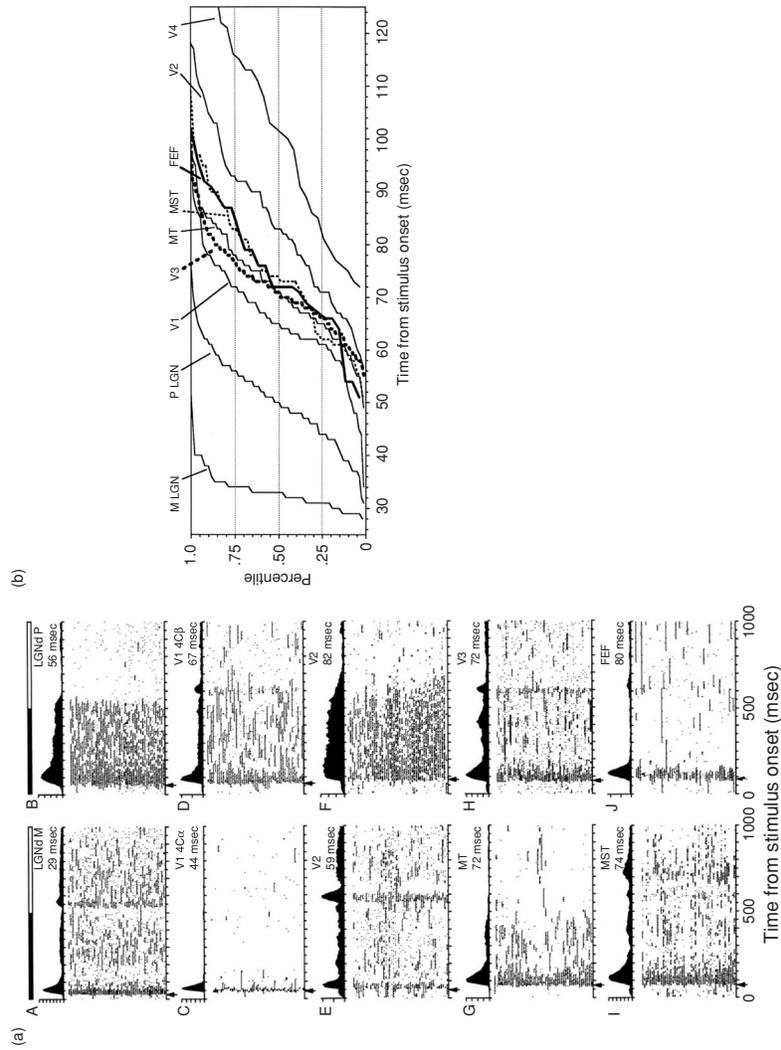


Fig. 14.1 Signal timing across the macaque visual system; data from Schmolesky et al. (1998). There is a wide spread of latencies both between and within areas. How does the visual system get the timing of events correct in the face of this signal variability? (A) Spikes are displayed as ticks in the raster panels, with fitted spike density histograms above. Black bar indicates the time the stimulus was on (0–500 ms). Arrow beneath each raster plot indicates the visual response latency of the cell; the value is also displayed in the top right corner. (B) Cumulative distributions of latencies. Percentage of cells that have begun to respond is plotted as a function of time from stimulus presentation.

14 How does the timing of neural signals map onto the timing of perception? 219

brain – in this case, across different areas of the visual system (Schmolesky et al. 1998). When a stimulus appears in the outside world, the response in the brain is smeared out over a large window. Think of Paul Revere and his colleagues spreading out in different directions over the New England landscape to deliver a message. Some riders are fast, others slow. Almost all of them inspire other riders to saddle up and ride off in different directions. As a result, the colonialists do not all get the message at once; it percolates at different rates to different streets in different townships. As can be seen in Fig. 14.1, there is also a good deal of trial-to-trial variability within a single neuron: in other words, each time you rerun history, each horseman may ride faster or slower than the last time.

What is mysterious about this physiology is the fact that humans have quite good resolution when making temporal judgments. For example, two visual stimuli can be accurately deemed simultaneous to 5 ms resolution, and their order can be assessed with 20 ms resolution (Hirsh & Sherrick 1961). How do the colonies conclude exactly when the British arrived, given the spread of signals in different locations? I will address this paradox by focusing on a well-studied example – the effect of intensity on latency – to drill down to the bottom of the mystery.

14.3 Contrast differences lead to latency differences

Latencies can be modulated in a specific reliable manner by changing the intensity of a stimulus. For example, in a study by Maunsell and colleagues, a monkey sat in a dark room and was sporadically presented with a flash (Maunsell et al. 1999). The flash ranged in luminance from low (0.43 cd/m^2) to high (28 cd/m^2). Figure 14.2A shows the response of cells in dorsal lateral geniculate nucleus (dLGN), a midpoint between the retina and visual cortex. As can be seen from the figure, low luminance stimuli cause smaller, and *later*, responses than higher luminance stimuli. This conclusion holds irrespective of the measure of latency (time to onset, half-peak, or peak; Fig. 14.2B). In other words, at the early stages of the visual system, even before reaching visual cortex, signals are already becoming spread out through time based on stimulus properties.

This fact is true in primary visual cortex as well. The contrast of an oriented bar changes the response latency of neurons in V1 (Fig. 14.2C) (Gawne et al. 1996). In higher visual areas, such as STS, stimulus contrast is the major factor affecting response latency, independent of the response magnitude (Oram et al. 2002).

These data raise a critical question: What does this temporal spread based on intensity mean for perception? Let's look at a few possibilities.

14.4 The online hypothesis

The first possibility we will consider is a popular, enduring, and likely incorrect view. The view is that neural latency differences between two stimuli will translate into perceptual time differences (Purushothaman et al. 1998; Whitney & Murakami 1998; Zeki & Bartels 1998; Patel et al. 2000; Whitney & Cavanagh 2000). I will refer to this as the *latency*

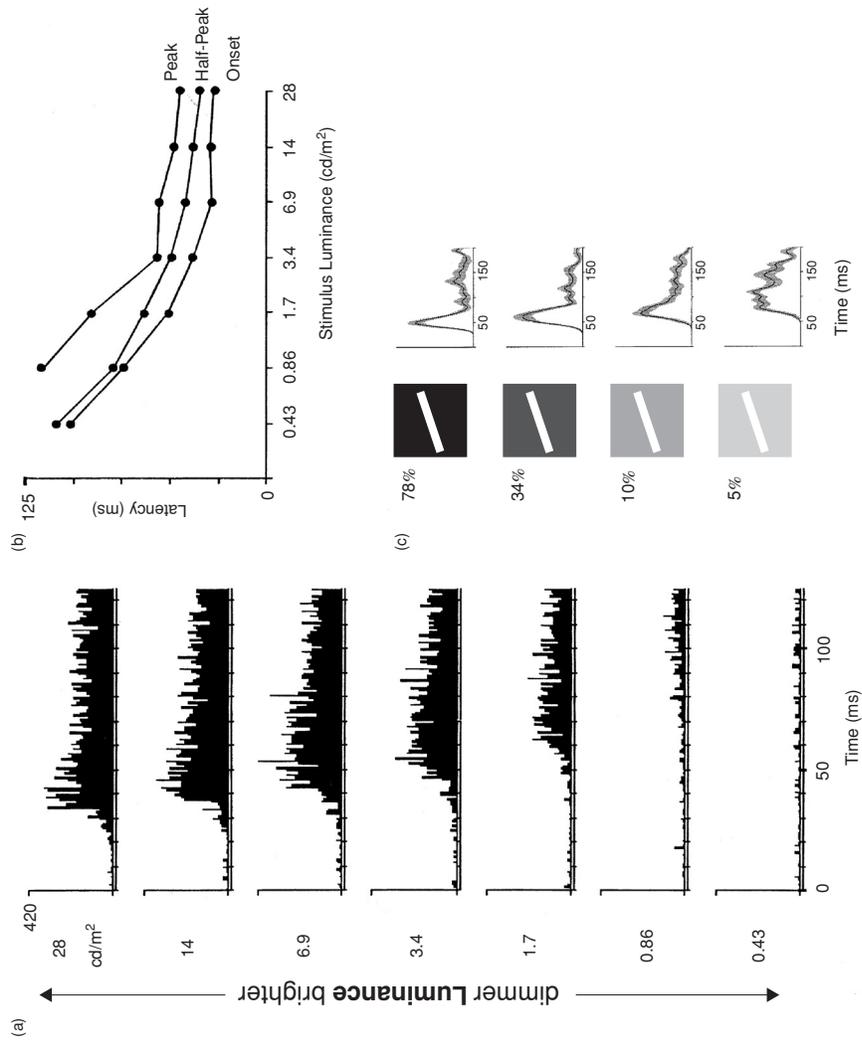


Fig. 14.2 Latency as a function of stimulus intensity. (A) Peristimulus time histograms of single neurons in dLGN in response to flashes of different luminance. Data from Maunsell et al. (1999). (B) The same data plotted differently to emphasize that latencies become shorter as the stimulus gets more intense, irrespective of the method for measuring those latencies. (C) Data from Gawne et al. (1996). Cells in V1 respond with later latency as the contrast of the stimulus is reduced.

14 How does the timing of neural signals map onto the timing of perception? 221

difference or *online* hypothesis. Essentially, this position holds that differences in neural latencies (as measured, say, in LGN or cortex, Fig. 14.2) map directly onto temporal differences in perception. Take as an example two simultaneously appearing stimuli, X and Y, each with different features. If stimulus X causes a cortical response before Y, the online hypothesis states that it is perceived first. The online hypothesis immediately raises some tricky questions (such as where, exactly, is the endpoint at which the rabbit and turtle are compared?). But we'll overlook those for the moment.

Because the online hypothesis is a seductive and pervasive habit of thinking, let's take it seriously for a moment. If it were true that stimuli were perceived online, that is, as soon as the leading edge of information reached some finish line in visual cortex, then neural correlates of perception would be easy to measure: whatever arrives in visual cortex first is perceived first. But some simple thought experiments call the theory into question. We will entertain the online hypothesis, show where it breaks down, and then take a look at its alternatives.

First, if there were illusions from differential latencies from onset, we should expect to suffer motion illusions each time we blink our eyes or turn on the lights. Following either of these events, everything in the visual field has a simultaneous onset. Given that there are bright and dim parts in the visual scene, wouldn't we see illusory movement – the bright areas of the room apparently moving toward the dim bits? It can be easily demonstrated for oneself that this does not occur.

Next, present to yourself a series of photographs flashed rapidly in sequence: a house, a tiger, a car, and so on. Even though each picture contains regions of low and high contrast, it is rarely temporally confused – that is, one does not generally perceive the stripes of the tiger on the house, the headlights on the tiger, and so on. Note that Intraub (1985) reported a temporal dissociation during rapid serial visual presentation, but only for a frame surrounding an outline drawing, and never for the contents of the drawing itself.

To give a more traditional psychophysical example, we turn to a series of optical illusions that should exist but do not. To directly address the online hypothesis, I have engineered a series of simple experiments, shown in Fig. 14.3. First, consider a vertical stack of squares with differing contrasts that move (as a unit) horizontally back and forth across the field of view (Fig. 14.3, top left). The online hypothesis predicts that because the brightest square enjoys the lowest latency, it may be perceived faster, whereas the dimmest square is getting processed most slowly (and so on for the squares in between). Even though the delay is constant, one might expect that the brightest square is always perceived at a position well ahead of the dimmest, and the colinear squares may appear to become noncolinear. However, all observers tested report that no such illusion occurs ($n = 8$). A similarly negative result is obtained with spinning bars with gradient textures (Fig. 14.3, bottom left). At all speeds tested, the bars appear solid, not curved or rubbery.

The theme is repeated with horizontally moving gradient squares (Fig. 14.3, top right). An online model predicts the contraction or dilation of the squares, depending on their direction of movement. That is, if the higher-contrast parts of the square have shorter latencies (and thus faster perception, in the online view), then a bright leading edge and a dim lagging edge

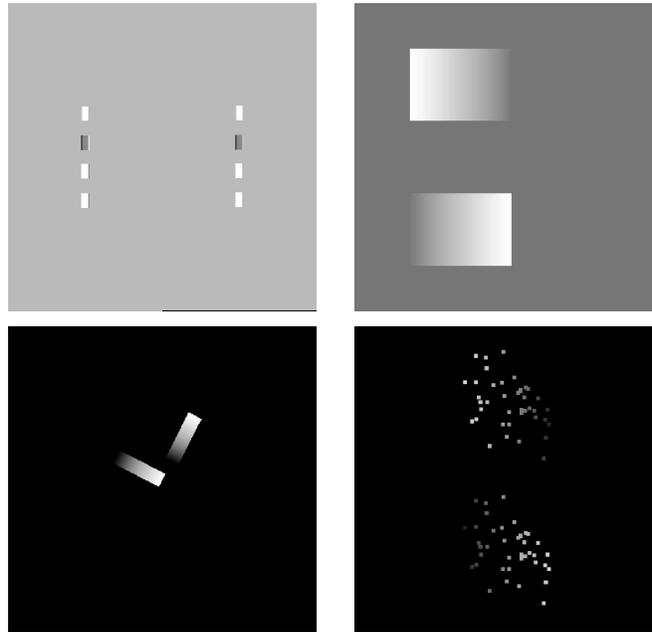


Fig. 14.3 Optical illusions that should exist but do not. Top left: A tower of four squares of differing contrast moves horizontally back and forth. An “online” interpretation of the physiology in Fig. 14.2 would predict that higher contrast squares should appear to be ahead of lower contrast squares. Instead, all observers report no illusion. The squares look vertically aligned, as they should be. Bottom left: Bars with gradients spin clockwise. An online hypothesis would predict that the bar does not look rigid. Again, the expected illusion does not occur. Top right: Two squares with contrast gradients move back and forth horizontally. An online hypothesis would predict that the square moving in the direction of its high luminance should look wider than the other square that moves in the direction of low luminance, which should look narrower. The illusion does not obtain. Bottom right: Same as above, but this time with a moving swarm of flashes. The flashes appear with a gradient of luminances. As above, the online hypothesis predicts a shrinking or stretching of the distribution in the direction of its motion, yet no such illusion occurs.

may appear to stretch out the distribution. Moving in the opposite direction, the square may appear to contract. I constructed a demonstration using two squares with opposite gradients for direct comparison to each other. As the squares moved back and forth, all observers reported there is no perceptual illusion of stretching or shrinking. To rule out the possibility that the deformation of the square was too small, I artificially stretched or constricted the squares by the amount predicted by the physiology in Fig. 14.2A, using the difference between bright and dim signals (the details of this choice will be justified below). Subjects were 100% accurate at detecting the artificially stretched or shrunken squares.

But perhaps one will argue that the edges of the squares give extra information. So in the final demonstration, the moving squares were replaced with random dot distributions (Fig. 14.3, bottom right). The dots were assigned luminances based on a horizontal gradient from low to high. The dots flashed on and off, never appearing in the same relationship

14 How does the timing of neural signals map onto the timing of perception? 223

to one another, and the distribution as a whole moved horizontally back and forth. Again, the distributions did not appear to stretch and shrink in relation to one another, as the online hypothesis would have erroneously predicted. An artificially shrinking or stretching distribution was easily detectable (see online demonstration).

Whatever other effects one might argue are occurring here (e.g., motion capture), the conclusion is the same: the visual system gets the frame-by-frame timing right (see also Kopinska et al. 2003, for similar results). Below I will suggest how it does so.

14.5 Reconciling latency differences with perceptual simultaneity, or, the magic 80 ms: waiting for the slowest signals

The experiments above illustrate that even though stimuli of different luminances give rise to widely different response latencies at the first stages of the visual system, these do not have a direct mapping onto differences in perceptual time. So how are these latencies reconciled?

I suggest that the answer comes from looking at the tasks and resources of the visual system. As one of its tasks, the visual system tries to get the timing of outside events correct. But for its resources, it has to deal with the peculiarities of the equipment that feeds it: the eyes and the thalami. These have their own evolutionary histories and idiosyncratic circuitries – and because of the details of their wiring, signals become spread out in time from the first stages of the visual system (e.g., Maunsell et al. 1999).

So if the visual brain, the recipient of smeared temporal information, wants to get events correct, time-wise, it may have only one choice: *wait for the slowest information to arrive*.

How long would the system have to wait? According to the physiology, I suggest it would have to wait about 80 ms to collect all the information, from the dimmest to the brightest. This number can be read directly from Fig. 14.2A and B: the latency difference between the dimmest and brightest stimuli is 80 ms on the y-axis.

Note that the ~80-ms time window (plus or minus 20 ms) crops up commonly in psychophysics. For example, this window is found in motion integration, motion deblurring (Burr & Morgan 1997), successive pattern integration (Di Lollo 1980), binocular pattern integration (Julesz & White 1969; Ross & Hogben 1974), backward masking (Bachmann 1994; Macknik & Livingstone 1998), and audio–video synchronization (Steinmetz & Engler 1993). In our previous work on the flash-lag effect, we showed an approximately 80-ms window over which a moving object could be manipulated *after* a flash and still achieve an effect on what the viewer reports having seen at the time of the flash (Eagleman & Sejnowski 2000a,b,c).

To my knowledge, this is the first suggestion that the psychophysical 80-ms window is directly mirrored in the physiology. That is, if the brain wants to wait for the slowest information, it must wait about 80 ms. This would allow the visual system to discount latency differences imposed by the early stages, but it has the disadvantage of pushing perception into the past. Counterbalancing the need to collect slow information is the survival advantage to operating as close to the present as possible. In other words, a system would not want to live too far in the past. I suggest that 80 ms is the smallest delay that allows

higher areas of the brain to account for the latencies engendered by the first stages of the system while still operating close to the present. Note that by studying the psychophysics of simultaneity across modalities, Kopinska and Harris (2004) also concluded that the brain waits for the slowest information to arrive – in their case, they suggested a delay of 94 ms.

Among other things, this strategy of waiting for the slowest information has the great advantage of allowing object recognition to be *independent of lighting conditions*. Imagine a striped tiger coming toward you under the canopy of a forest, passing through different patches of sunlight. Neuroscience currently appreciates how difficult the task of object recognition is for the visual system. Now imagine how much harder the task would be if the shadow-play across the tiger caused incoming signals to be processed at different speeds. Different fragments of the tiger would sunder into different locations. Somehow the visual system has evolved to reconcile different information latencies; after all, it is advantageous to recognize tigers regardless of the lighting.

We should note that it has been suggested that latencies could be used by the system as a code (Oram et al. 2002), perhaps for object recognition (Thorpe et al. 2001) or object feature binding (Gawne et al. 1996). For example, if signals arriving along transmission lines arrive in the order 3–6–2–4, that could code for ostrich, whereas 2–7–4–5 codes for carrot. This idea is known as rank order coding. Despite the appeal of its simplicity, I suggest that using latencies for object recognition is untenable in biology – and this is because the luminances from an object are dependent on the lighting conditions. If you learn to recognize a statue of George Washington in the morning, the lighting angle in the afternoon will make the contrast-based latencies entirely different – and a rank order coding network will utterly fail to recognize it. Marrying one’s visual recognition capacities to the particulars of the lighting conditions is not a move Mother Nature seems to have taken.

Finally, an important point that needs to be clarified. The 80-ms window under discussion does not imply that the visual system sees in “chunks.” The idea that the visual system takes discrete “snapshots” has been entertained (Varela et al. 1981; VanRullen & Koch 2003), but the available evidence speaks against it (Kline et al. 2004; Kline & Eagleman 2008). Instead, the 80-ms window we are discussing appears to be a duration over which the visual system waits to collect information, but it can still retain the ability to differentiate events within that window (Blake & Lee 2005). For example, if I were to collect the dot-dash-dot-dot that designates the letter “L” in Morse code, I can know that the temporally smeared information applies to one object and yet still report on the order of dots and dashes. Further, I speculate that an 80-ms window may only be triggered when a perceptual question is asked; this is the topic of future exploration (Eagleman & Dennett, in preparation).

14.6 Latency illusions that do exist, and why

Getting back to the examples in Fig. 14.3, these may strike the reader as a contradiction of other reports in literature. For instance, in the Pulfrich effect, a pendulum appears to rotate in depth when a neutral density filter is placed in front of one eye. Following a suggestion by Fertsch, Pulfrich hypothesized a timing difference between signals from the two eyes

14 How does the timing of neural signals map onto the timing of perception?

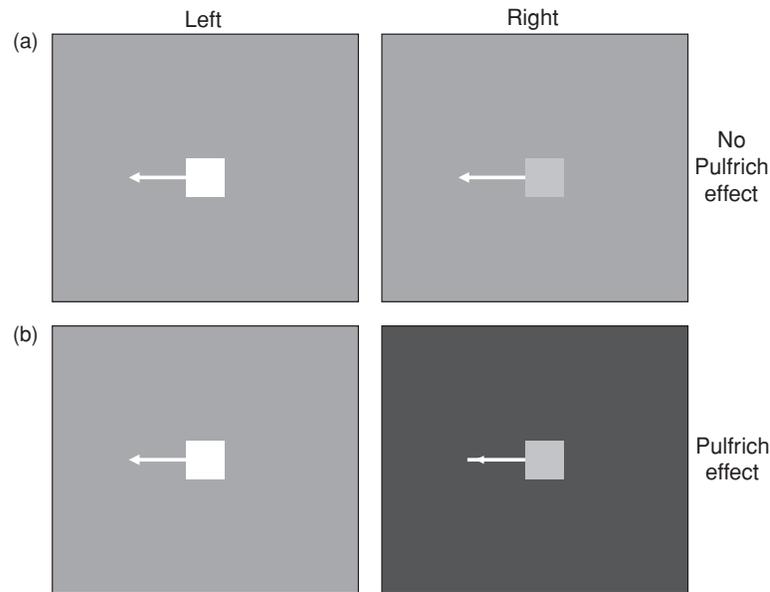


Fig. 14.4 Conditions for the Pulfrich effect. (A) The Pulfrich effect is not obtained simply by reducing the contrast of the moving stimulus to one eye (in this case, the square presented to the right eye). (B) Instead, the effect is only obtained when the presentation to one eye is globally darkened, as with a neutral density filter. We suggest this is because the visual system has evolved to deal with a range of luminance differences in normal vision but has not evolved to deal with the pathologic case of different global levels of illumination between the two eyes.

(Pulfrich 1922; Eagleman 2001). We have every reason to believe that latency differences are the best explanation for the Pulfrich effect, so this is often cited in support of an online model. However, it is critical to note that the Pulfrich effect is a special case. Specifically, our argument is that the visual system has evolved to account for latency differences such that under *normal circumstances* it will not be fooled by variations in luminance. The Pulfrich effect is a totally unnatural stimulus because one retina is reading information at an entirely different mean-luminance level than the other retina. The impulse response function on the dimmer retina is greatly slowed (Purpura et al. 1990), leading to the illusion. Of the many assumptions that the brain makes, an important one is that both eyes are viewing the same luminance in the scene. The visual system has not evolved to deal with deviations from this condition, and it is not surprising that an illusion can be induced.

In support of this argument, I have discovered a fact that appears not to be widely known. When observers are shown a binocular presentation of a square that moves back and forth horizontally, the Pulfrich illusion is *not* produced when one merely lowers the luminance of one of the moving squares (Fig. 14.4A). This negative result tends to come as a surprise to psychophysicists familiar with the illusion. Instead, the effect is only produced when the entire display to one eye is darkened, as with a neutral density filter (Fig. 14.4B). In other words, the effect of a neutral density filter over one eye is not simply to reduce the

luminance of the moving object but more broadly to expose the entire retina to a different mean luminance.

I have found the same to be true of the Hess effect, an illusion in which an offset can be perceived between two horizontally moving dots when one is of high contrast and the other low contrast (Wilson & Anstis 1969). Specifically, I find that this effect is only seen when one uses a neutral density filter over half the screen – simply reducing the contrast of a single dot is insufficient (at least under photopic viewing conditions). Fooling the visual system with a latency difference requires slowing the signals through all or part of the retina – simply changing the luminance of the moving stimulus is insufficient because, as I've argued above, the visual system is equipped to deal with and account for difference object luminances in normal vision.

14.7 Temporally spread signals in the neural tissue

So we have asserted that the brain can keep account of latencies. But how exactly could it know what happened when? To highlight this problem, we can phrase it as a question: If you were a V1 neuron and received a burst of spikes, how would you know if that meant a dim flash occurred 150 ms ago or a bright flash occurred 70 ms ago? There may be at least two strategies the system can employ to take care of this problem with temporally spread, delayed signals: “timestamp” them or physically correct them on the fly (that is, temporally align the signals). Because we do not know which the system does, we will briefly sketch out possible methods for both.

First, I'll address the timestamp model. Even at the single-neuron level, there is more information available than simply the latency. Notice in Fig. 14.2A that the different spike trains have different temporal signatures. For instance, the spike rate at 28 cd/m² has a sharper onset than the spike train at 1.7 cd/m². These structures in the spike timing (or the interspike interval) could in theory allow downstream neurons to distinguish an old dim flash from a recent bright flash merely by the structure through time. This would be one way to reconstruct the actual order of events.

As an alternative, neural circuitry could reconcile differential latencies by physically aligning the timing of signals. Although most of the available data in the field comes from single electrode recordings, it is critical to keep in mind that *populations* of neurons could manipulate both latencies and the variability of those latencies. For example, one way the brain could align signals is by dynamically recruiting more or fewer neurons to speed and slow the passage of signals. For example, a weaker signal (lower spike rate) could recruit more neurons downstream, which could bring further neurons to threshold more quickly. Conversely, faster spike rates activate fewer downstream neurons, such that the timing is slowed. The mechanism for accomplishing this could be fast synaptic depression, with the result that the timing of simultaneous events in the world would be temporally “lined up” at higher stages of the nervous system. In fact, as pointed out by Maunsell et al. (1999), the faster speed of the magnocellular pathways cannot be predicted solely from differences in axon conduction speeds; instead the degree of convergence may be highly regulatory in

14 How does the timing of neural signals map onto the timing of perception? 227

timing issues. For instance, different degrees of convergence (beginning with retinal circuits and continuing through the brain), can speed or retard timing – such that a weaker signal in a highly convergent pathway might lead to faster downstream responses than a faster signal in a pathway with less convergence (Maunsell et al. 1999). So although the details of circuitry are most commonly thought of in terms of connectionist functions (Callaway 1998), the circuitry could also, in theory, lead to changes in timing. If true, this could in the future inspire us to rethink circuitry not only in terms of spatial connections but also what they are implementing temporally.

Although some possibilities for speeding and slowing signals have been sketched, it is critical to remember that there is no theoretical necessity for neural signals to line up temporally for the perception of simultaneity. As in an earlier example, the letter “L” is represented in Morse code by a dot, a dash, and two more dots. When someone receives this temporal signal, it is interpreted as an “L.” The same could hold for perception: signals arriving at different times could be interpreted as simultaneous events in the outside world. In other cases, simultaneously arriving signals might be best interpreted by perception as asynchronous events in the real world. As Uttal stated this point: “The essence of much of the research that has been carried out in the field of sensory coding can be distilled into a single, especially important idea – any candidate code can represent any perceptual dimension; there is no need for an isomorphic relation between the neural and psychophysical data. Space can represent time, time can represent space, place can represent quality, and certainly, nonlinear neural functions can represent linear or nonlinear psychophysical functions equally well” (Uttal 1979).

Finally, we note the importance of massive feedback connectivity in visual awareness (Mumford 1994; Nowak & Bullier 1997; Lamme & Roelfsema 2000). It is in this light that “waiting for the slowest signal” must be finally understood physiologically. Our intuition is that the 80-ms window will not be found in terms of a neural information buffer in a feedforward buffer, but rather in the settling of recurrent networks into a larger pattern. We do not at present know the details of how this works; this is exciting open ground for the future.

14.8 Conclusions

The problem faced by the visual system is a common problem for biological creatures in a temporal world. Understanding the timing of events is critically important, but the signals representing that timing may be spread out in time. This requires a reconstruction of event timing, whether implicitly (symbolic coding) or explicitly (aligning signals in time). Current data are too sparse to arbitrate between these two methods; the nervous system may use either or both.

I have argued that the back of the brain seeks to judge timing of events accurately, but it has to contend with the temporally smeared information sent to it by the eyes and thalamus. The best solution to this problem may be to wait for the slowest information to arrive. Electrophysiology from the primate visual system shows that the window over

which information should be collected is about 80 ms, the latency difference between the fastest and slowest signals. I suggest this is the physiological basis of the 80-ms window observed in a large variety of psychophysics experiments, as detailed above.

Among other things, collecting signals over a window of time allows object recognition independent from lighting conditions, because latency differences based on different luminances can be discounted. In contrast, models that tie themselves to latency differences for object recognition (Thorpe et al. 2001) are unfortunately tied to the details of the lighting, a disadvantageous move for any visual organism.

The argument I have made – that the system waits to collect information over the expected window of time over which it will come streaming in – applies not only in vision, but more generally to all modalities. Therefore, although it is possible to measure an 80-ms window of postdiction in vision (Eagleman & Sejnowski 2000), the breadth of this window may be different in audition and somatosensation. It may therefore be that a unified polymodal perception of the world has to wait for the slowest overall modality. Given conduction times along limbs, this leads to the bizarre but testable suggestion that tall people may live further in the past than short people. The consequence of waiting for temporally spread signals is that perception becomes something like the airing of a “live” television show. Such shows are not truly aired live but are instead delayed by a small window of time in case editing becomes necessary.

Note that the 80-ms window proposed here is the minimum duration that the visual system would have to wait to collect all the information from a visual event. However, it is possible that when including signal travel time and multimodal unification, the total lag between a physical moment and its conscious perception is much longer. Although the total delay between a stimulus and conscious awareness has been proposed to be impossible to know for certain (Dennett 1991), some investigators have estimated the total time-to-awareness in the range of 100–150 ms (Lamme 2003) to 500 ms (Libet et al. 1967).

It must be emphasized that everything I have discussed in this chapter is about visual awareness. It seems clear from preconscious reactions that the motor system does not wait for all the information to arrive before making its decisions. In general, the motor systems can act appropriately with partial or no participation of awareness (Goodale & Milner 1992, 2004) – for example, visual information streams directly to the amygdala, which may have a direct and rapid connection to the motor systems.

This raises a question: What is the *use* of perception, especially given our argument that perception lags reality, is only retrospectively attributed, and is generally outstripped by automatic (unconscious) systems? The most likely answer is that perceptions are representations of information that the brain can manipulate at a later date – these representations can be worked with by cognitive systems the way that tools are handled by the motor systems. For this reason, it is important for the brain to take sufficient time to settle on its best interpretation of what just happened, rather than simply its initial, unfinished interpretation. Its carefully refined picture of what just happened is the only thing it will have to work with later.

14 How does the timing of neural signals map onto the timing of perception? 229

In conclusion, although many neural models in the literature intimate that a well-defined input is neatly mapped onto a particular output as soon as the leading edge of the information reaches an endpoint (e.g., Zeki & Bartels 1998; Patel et al. 2000; Thorpe et al. 2001), the framework presented here highlights the fact that neural dynamics are influenced through time by the ongoing input of sensory information. We hope this starting point will help navigate us to a physiological explanation of visual awareness.

References

- Bachmann, T. (1994). *Psychophysiology of Visual Masking*. Commack, NY: Nova Science.
- Blake, R., & Lee, S. H. (2005). The role of temporal structure in human vision. *Behav Cogn Neurosci Rev* **4**(1): 21–42.
- Burr, D. C., & Morgan, M. J. (1997). Motion deblurring in human vision. *Proc Biol Sci* **264**(1380): 431–436.
- Callaway, E. M. (1998). Local circuits in primary visual cortex of the macaque monkey. *Annu Rev Neurosci* **21**: 47–74.
- Crick, F., & Koch, C. (1990). Some reflections on visual awareness. *Cold Spring Harb Symp Quant Biol* **55**: 953–962.
- Dennett, D. C. (1991). *Consciousness Explained*. Little Brown & Co.
- Di Lollo, V. (1980). Temporal integration in visual memory. *J Exp Psychol Gen* **109**(1): 75–97.
- Eagleman, D. M. (2001). Visual illusions and neurobiology. *Nat Rev Neurosci* **2**: 920–926.
- Eagleman, D. M. (2008). Human time perception and its illusions. *Curr Opin Neurobiol*. In press.
- Eagleman, D. M., & Dennett, D. (In preparation). Computation on a need-to-know basis.
- Eagleman, D. M., & Sejnowski, T. J. (2000a). Motion integration and postdiction in visual awareness. *Science* **287**: 2036–2038.
- Eagleman, D. M., & Sejnowski, T. J. (2000b). Latency difference versus postdiction: Response to Patel et al. *Science* **290**(5494): 1051a.
- Eagleman, D. M., & Sejnowski, T. J. (2000c). The position of moving objects: Response to Kregelberg et al. *Science* **289**: 1107a.
- Eagleman, D. M., & Sejnowski, T. J. (2003). The line-motion illusion can be reversed by motion signals after the line disappears. *Perception* **32**: 963–968.
- Eagleman, D. M., & Sejnowski, T. J. (2007). Motion signals bias position judgments: A unified explanation for the flash-lag, flash-drag, flash-jump and Frohlich effects. *J Vision* **7**(4): 1–12.
- Engel, A. K., Konig, P., et al. (1992). Temporal coding in the visual cortex: new vistas on integration in the nervous system. *Trends Neurosci* **15**(6): 218–226.
- Gawne, T. J., Kjaer, T. W., et al. (1996). Latency: another potential code for feature binding in striate cortex. *J Neurophysiol* **76**(2): 1356–1360.
- Goodale, M. A., & Milner, A. (1992). Separate visual pathways for perception and action. *Trends Neurosci* **15**: 20–25.
- Goodale, M. A., & Milner, A. D. (2004). *Sight Unseen*. Oxford: Oxford University Press.
- Hirsh, I. J., & Sherrick, C. E. (1961). Perceived order in different sense modalities. *J Exp Psychol* **62**: 423–432.

- Intraub, H. (1985). Visual dissociation: an illusory conjunction of pictures and forms. *J Exp Psychol: Hum Percept Perform* **11**: 431–442.
- Julesz, B., & White, B. (1969). Short term visual memory and the Pulfrich phenomenon. *Nature* **222**(194): 639–641.
- Kinsbourne, M. (1993). Integrated cortical field model of consciousness. *Ciba Found Symp* **174**: 43–50.
- Kline, K. A., & Eagleman, D. M. (2008). Evidence against the snapshot hypothesis of illusory motion reversal. *J Vis* **8**(4): **13**, 1–5.
- Kline, K. A., Holcombe, A. O., et al. (2004). Illusory reversal of motion is caused by rivalry, not by perceptual snapshots of the visual scene. *Vision Res*.
- Kolers, P., & von Grunau, M. (1976). Shape and color in apparent motion. *Vision Res* **16**: 329–335.
- Kopinska, A., & Harris, L. R. (2004). Simultaneity constancy. *Perception* **33**: 1049–1060.
- Kopinska, A., Harris, L. R., & Lee, I. (2003). Comparing central and peripheral events: compensating for neural processing delays. *J Vis* **3**: 751a.
- Lamme, V. A. (2003). Why visual attention and awareness are different, *Trends Cogn Sci* **7**: 12–18.
- Lamme, V. A., & Roelfsema, P. R. (2000). The distinct modes of vision offered by feedforward and recurrent processing, *Trends Neurosci* **23**: 571–579.
- Libet, B., Alberts, W. W., Wright, E. W., & Feinstein, B. (1967). Responses of human somatosensory cortex to stimuli below threshold for conscious sensation. *Science* **158**: 1597–1600.
- Macknik, S. L., & Livingstone, M. S. (1998). Neuronal correlates of visibility and invisibility in the primate visual system. *Nat Neurosci* **1**(2): 144–149.
- Maunsell, J. H., Ghose, G. M., et al. (1999). Visual response latencies of magnocellular and parvocellular LGN neurons in macaque monkeys. *Vis Neurosci* **16**(1): 1–14.
- Mumford, D. (1994). Neuronal architectures for pattern-theoretic problems. In: C. Koch & J. L. Davis (eds.), *Large-Scale Neuronal Theories of the Brain* (125–152). Cambridge, MA: MIT Press.
- Nowak, L. G., & Bullier, J. (1997). The timing of information transfer in the visual system. In: J. Kass, K. Rockland, & A. Peters (eds.), *'Extrastriate Cortex' Cerebral Cortex*, Vol. 12 (205–241). New York: Plenum Press.
- Oram, M. W., Xiao, D., et al. (2002). The temporal resolution of neural codes: does response latency have a unique role? *Philos Trans R Soc Lond B Biol Sci* **357**(1424): 987–1001.
- Patel, S. S., Ogmen, H., et al. (2000). Flash-lag effect: differential latency, not postdiction. *Science* **290**(5494): 1051.
- Pessoa, L., Thompson, E., et al. (1998). Finding out about filling-in: a guide to perceptual completion for visual science and the philosophy of perception. *Behav Brain Sci* **21**(6): 723–748; discussion 748–802.
- Pulfrich, C. (1922). Die Stereoskopie im Dienste der isochromen und heterochromen Photometrie. *Die Naturwissenschaften* **10**.
- Purpura, K., Tranchina, D., et al. (1990). Light adaptation in the primate retina: analysis of changes in gain and dynamics of monkey retinal ganglion cells. *Vis Neurosci* **4**(1): 75–93.
- Purushothaman, G., Patel, S. S., et al. (1998). Moving ahead through differential visual latency. *Nature* **396**(6710): 424.
- Ross, J., & Hogben, J. H. (1974). Short-term memory in stereopsis. *Vision Res* **14**(11): 1195–1201.

14 How does the timing of neural signals map onto the timing of perception? 231

- Schmolesky, M. T., Wang, Y., Hanes, D. P., Thompson, K. G., Leutgeb, S., Schall, J. D., et al. (1998). Signal timing across the macaque visual system. *J Neurophysiol* **79**(6): 3272–3278.
- Steinmetz, R., & Engler, C. (1993). Human perception of Media Synchronization. *IBM ENB Tech Rep* **43**(9310). Heidelberg.
- Thorpe, S., Delorme, A., et al. (2001). Spike-based strategies for rapid processing. *Neural Network* **14**(6–7): 715–725.
- Uttal, W. R. (1979). Do central nonlinearities exist? *Behav Brain Sci* **2**: 286.
- VanRullen, R., & Koch, C. (2003). Is perception discrete or continuous? *Trends Cogn Sci* **7**(5): 207–213.
- Varela, F. J., Toro, A., et al. (1981). Perceptual framing and cortical alpha rhythm. *Neuropsychologia* **19**(5): 675–686.
- Whitney, D., & Cavanagh, P. (2000). Latency difference, not postdiction. *Science*.
- Whitney, D., & Murakami, I. (1998). Latency difference, not spatial extrapolation. *Nat Neurosci* **1**(8): 656–657.
- Wilson, J. A., & Anstis, S. M. (1969). Visual delay as a function of luminance. *Am J Psychol* **82**(3): 350–358.
- Zeki, S., & Bartels, A. (1998). The asynchrony of consciousness. *Proc R Soc Lond B Biol Sci* **265**(1405): 1583–1585.
- Zeki, S., & Bartels, A. (1998). The autonomy of the visual systems and the modularity of conscious vision. *Philos Trans R Soc Lond B Biol Sci* **353**(1377): 1911–1914.