2
The Unfinished NCC Project

2.1 Derailment

Some of my colleagues who are familiar with the scientific literature will probably snicker as they skim through the opening chapters. All that high-sounding long-windedness to state the obvious. Isn't an empirical approach grounded in cognitive neuroscience what we have always been doing?

It is true that this approach is really nothing new. In recent times, we sometimes say we are trying to find the neural correlates of consciousness (NCC). This so-called NCC project started in the 1990s. But as with the whole idea of a "revival" of consciousness science, there too is a long and often neglected history well before that (Michel 2019; LeDoux, Michel, and Lau 2020).

Not knowing our history is a problem beyond pedantic concerns of scholarship, because we are prone to repeating the same mistakes. The trouble is that some conceptual issues are hard to avoid, even when we think we are focusing on "straight-ahead" empirical studies of consciousness. In recent years, in fact, we have seen many of these same problems creep up again (Michel and Lau 2020). Amid some exciting calls for *moving on* to more theoretically adventurous endeavors, there is a worry that the empirical project itself could well be derailed. The purpose of this chapter—as well as the next—is to set us back on the right path.

2.2 Mere Correlates?

In cognitive neuroscience, we look for neural mechanisms for psychological phenomena. Rarely do we say we are merely looking for *correlates*. When we say a certain neuronal circuit provides a mechanism for a psychological phenomenon, we mean they are causally or constitutively relevant (Craver 2007). For example, certain neuronal firing patterns in the motor cortex lead to specific patterns of muscular movements. Some other neuronal firing patterns in the hippocampus make possible the formation of certain memories.

In Consciousness We Trust. Hakwan Lau, Oxford University Press. © Hakwan Lau 2022.
DOI: 10.1093/oso/9780198856771.003.0003

Understanding these neuronal circuits is supposed to provide mechanistic accounts of the phenomena.

But for consciousness, things are trickier. Some may think that such mechanistic explanations won't ever be possible. Perhaps certain neural pattern is just *one and the same* as the conscious experience itself (Place 1956). This identity view is somewhat different from the causal view, because logically speaking, a thing cannot cause itself. But to say we are looking for the neural identities or identifiers for consciousness may be too strong; not everyone is happy with the identity view either. This can lead to endless debates about metaphysics. In order to sidestep the issues, the more neutral term of *correlates* was adopted.

But correlates is a very weak notion. As they often say, correlations need not imply causation. Things that aren't directly causally connected may well be correlated. For example, people's consumption of ice cream and use of swimming pools are correlated over the year. But that is not because eating ice cream makes people swim more, or the other way around. Both are just caused by another factor: the heat in the summer. So likewise, many neural activity patterns will be correlated with consciousness, but they too may be neither causally nor constitutively relevant. Perhaps they just reflect the state of being awake, or the ability to produce complex behavior, which isn't quite the same as subjective experiences.

2.3 Necessity and Sufficiency

So there have been some attempts to finesse the definition of the NCC, by borrowing the precise language of necessity and sufficiency from logic and philosophy. The NCC is sometimes said to be the minimally sufficient conditions for a subjective experience to occur, given certain background conditions (e.g., oxygen in the blood or an intact brain stem; Chalmers 2000). This rules out some distal, indirect correlates. That is, if this NCC occurs, we must have the relevant subjective experience (given the background conditions).

The trouble is that we scientists are actually not very good at using this language. Worse still, we often overlook the complexity involved (Miller 2014; Michel and Lau 2020). To appreciate this, one can try explaining the following to their colleagues in neuroscience or cognitive psychology: so the engine of my car is what really drives it. It is both necessary and sufficient for the car to move around.

So far, so good? Shockingly, many would think so, especially if you mention it in passing without warning. But of course, if we think carefully, the engine is not necessary for the car to move around at all. Without the engine, we can push the car. Or maybe it's a hybrid so there is also an electric motor. So even though the engine is the primary causal mechanism, it is not necessary, strictly speaking. Nor is it sufficient. If I leave the engine intact, but take off the wheels, the car wouldn't move around.

The moral here is to not to say we scientists can't think logically. Most of us can, if we pay attention to what the terms "necessity" and "sufficiency" *really* mean. But this is precisely the point: we must always pay attention. Specifically, there are two caveats to bear in mind: First, abolishing the primary causal mechanism (e.g., the engine) may not always abolish the relevant function. This is true for sufficient conditions in general; they aren't irreplaceable. The specification in the NCC definition that they need to be *minimally* sufficient helps to rule out that we are not including within the NCC a large set of mutually redundant conditions. This is to say, we are hoping to home in on the barebones. This may give the ring that these are the "minimal *necessities*," and perhaps this is why scientists confuse themselves at times. But barebones aren't strictly necessary in the logical sense. Just because we do not include redundancy in our definition of the NCC does not mean there's no redundancy in the brain. These barebones (minimally sufficient conditions) could well be replaced by something else—perhaps something clumsier, less minimalist, and less efficient—which may end up doing a similar job. This is a tricky point that we will revisit again in the next chapter.

Second, abolishing some peripheral mechanisms (e.g., the wheels or the gas tank) may well abolish the relevant function, even when the primary causal mechanism itself is intact. In terms of the NCC definition of minimal sufficiency, that's because other background conditions may be necessary. The clause about background conditions exists for good reasons.

As we will see, these two caveats are often neglected, sometimes by some of the original proponents of the NCC definition themselves (Michel and Lau 2020). So even if we don't like to think of the NCC as a causal mechanism (e.g., because of metaphysical considerations), it is often helpful to think of it *almost as if* it is one, just so as to remind ourselves of the two caveats above. It is easier to think clearly this way, in terms of mechanisms, rather than in terms of necessary and sufficient conditions.

2.4 Blindsight & the Primary Visual Cortex

As we mentioned briefly in the introduction, blindsight is the neurological condition in which subjective experience in vision is selectively impaired (Weiskrantz 1997). These patients deny that they consciously see, but when pressed, they can "guess" the identity of simple visual stimuli in front of them correctly. Because blindsight is traditionally found in patients with restricted lesions to the primary visual cortex (V1), it has been suggested that V1 is a good candidate for the visual NCC (Lamme 2001, 2003).

This is supported by some studies involving subjects from the general population too. In these studies, transcranial magnetic stimulation is applied to V1. Under the right stimulation parameters, this disrupts the ongoing neural activity in the area, creating a temporary "virtual lesion" effect. This abolished perceptions which would otherwise be conscious (Pascual-Leone and Walsh 2001; Boyer, Harrison, and Ro 2005; Ro et al. 2010).

On the face of it, these findings may seem to support the local view of consciousness introduced in Chapter 1. But, of course, this is where the second caveat from the last section comes into play. How do we know that a normally functioning V1 is not just a normally required background condition? Most of the visual information from the eyes enters the cerebral cortex through V1. So lesioning V1 may be a bit like cutting off the tube from the gas tank to the engine. It may abolish the relevant functions, but V1 may not be the engine itself (Silvanto 2015).

Turns out, by stimulating other visual areas, researchers could induce conscious visual experience in a blindsight patient without V1 too (Mazzi, Savazzi, and Silvanto 2019). There have also been reports of patients with damage to the visual cortex who experienced vivid hallucinations (discussed in Lau and Brown 2019), or that they can voluntarily engage in conscious mental imagery (Bridge et al. 2012). So the engine does look like it could be elsewhere.

2.5 Content Mismatch in Early Visual Areas

So, lesions alone, real or virtual, may not tell us the full story. Fortunately, there are also many studies measuring neural activity in V1. Some of these studies showed that activity in V1 was in fact correlated with many aspects of the subjective experience of seeing (Tong 2003). But this is not always true.

For example, it is known that we are rather poor at identifying which eyeball is the origin of certain visual information (Schwarzkopf, Schindler, and Rees 2010). If I flash some light briefly to one of your eyes in a controlled laboratory setting, you will most likely see the light but will not be so good at pinpointing which of the two eyeballs received the stimulation. And yet, some cells within V1 are highly sensitive to this information. Many neurons are mainly driven by stimulation to one eye but not the other.

This may be fine for local theorists. They can say that not all activity in V1 is part of the NCC. Some activity in V1 may reflect information of which we are not consciously aware. But still, the visual NCC can be within a subset of activity in V1. Also, some local theorists may say that the NCC is not within V1. Extrastriate areas (Macknik and Martinez-Conde 2008; Fisch et al. 2009; Malach 2011), and/or their feedback projections back to V1 (Lamme 2001, 2003), are perhaps more important.

However, using imaging methods that measure the overall population activity of brain regions, others have found strong nonconscious effects in V1 and beyond. That is, population activity within most visual areas seems to be able to distinguish between stimuli that subjectively look identical (e.g., Jiang, Zhou, and He 2007; Zhang et al. 2012; Huang et al. 2020). One common idea in these studies is to use various "tricks" from vision psychophysics to create a visual illusion of some sort. So the physical stimulus may be presented in a certain way, but subjectively it looks a different way. The question is whether activity in early visual areas tracks the physical stimulus or the subjective percept, under such illusory conditions when the physical and subjective come apart.

One such experiment was based on the double-drift illusion (Liu et al. 2019). A Gabor patch is an abstract stimulus often used by vision scientists. It looks like what is shown in Figure 2.1. When the pattern within the Gabor patch shifts locally, sometimes we mistake the entire patch to be drifting (i.e., relocating) sideways, even when it really isn't. Liu et al. (2019) contains links to video demonstrations of the illusion.

This experiment was conducted using functional magnetic resonance imaging (fMRI), where subjects' brain activity was indirectly measured (via local blood oxygenation changes). By looking at the fine-grained spatial pattern of activity in different brain regions—a method also known as multivoxel pattern analysis (MVPA)—the researchers first trained an algorithm (also called a decoder) to classify the motion based on actual physical drifts (i.e., relocation, not local pattern shifts). After the researchers identified that these spatial patterns of activity represented normal drift signals, they tested whether similar patterns occur during the illusory drifts. They found that this was not the case for any of the visual areas tested. So the visual areas seem to represent the physical movement of the stimulus, not the subjective percept.

In another study also using fMRI, the researchers manipulated the subjectively perceived duration of a simple visual stimulus using sinusoidal gratings (stimuli somewhat similar to a Gabor patch). Again by inducing local pattern shifting within the stimulus, they shortened the perceived stimulus duration subjectively. However, under the illusion, they found that in early visual areas,

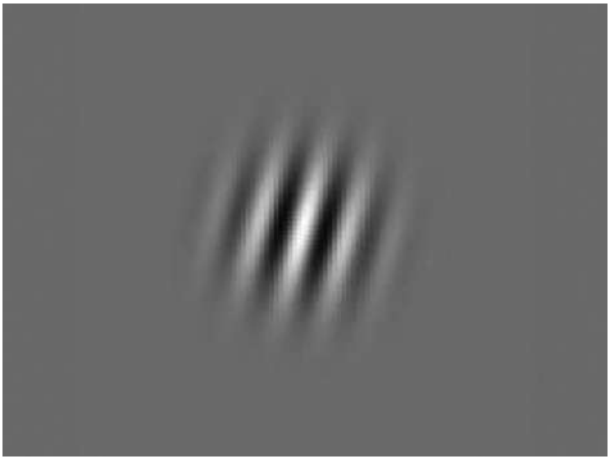

Figure 2.1 A Gabor patch

the signal correlated with the actual physical duration, rather than the subjectively perceived duration (Binetti et al. 2020).

Just for yet another example using a different neuroimaging method, magnetoencephalography (MEG), others and I have scanned subjects while they judged the brightness of some disks. By exploiting another illusion, called the temporal context effect (Eagleman, Jacobson, and Sejnowski 2004), we changed the subjective brightness of a disk, without changing its physical brightness. Typically, we expect certain patterns of activity within the early visual cortex to correlate positively with brightness. But instead, we found the opposite relationship (Zhou et al. 2020).

As an aside, this was a form of adversarial collaboration, perhaps one of the firsts in the field. Although I was more sympathetic to something akin to the global view when I was planning the study, the data were analyzed mainly by researchers in support of a local view (Sandberg et al. 2013). Together, we tried to find correlates in the early visual areas. We just couldn't. Or, we could say we did, but they turned out to be the opposite as expected.

Overall, the temporal aspects of early visual activities may also be problematic if they are treated as candidates for the NCC. Because early visual activities reflect specific perceptual content, the straightforward localist interpretation is that as a certain subjective experience occurs, visual activities should reflect the relevant content at that moment. However, conscious experience takes some time to be fully determined, whereas early visual processes are relatively fast. Our awareness of a stimulus presented at time *t* is known to depend on stimuli presented up a fraction of a second later. Therefore, Michel and Doerig (2021) argued that early visual correlates probably do not have the right time scale to match our conscious experiences.

Relatedly, when two similar stimuli are presented in close succession, the second stimulus is known to elicit weaker activity in the early sensory regions. This phenomenon is sometimes called repetition suppression (Gotts, Chow, and Martin 2012). However, subjectively, we don't tend to perceive the second stimulus as being substantially weaker. Again, this suggests that without some downstream calibration mechanisms, early sensory activities alone are unlikely to always match the content of subjective experiences.

2.6 Distal Cause versus the "Engine" Itself

I do not mean to say that subjective experiences don't correlate with early visual activities at all. Apparently they do, in many studies (Tong 2003). But one may also expect a certain publication bias is at play: vision scientists tend to start with the hypothesis that such correlates can be found within the visual areas, and when they are not, the study may be considered a failure, and accordingly the negative findings may not be reported.

For the sake of argument, let us grant the local theorists that overall, there are likely many more cases where the subjective percept correlates with early visual activity. There is still a problem, though, if this correlation breaks down in a good number of minority cases. This may suggest that the early visual areas are just a distal cause of subjective experience, a bit like the retina itself. Of course retinal activity correlates with features of what we see *most* of the time—but not all the time, which is why we generally don't think of the retina as part of the NCC. So the cases mentioned in the last section, where the subjective percept and early visual activity come apart, may be critical.

There is of course an issue of sensitivity in the measurements. Using neuroimaging, sometimes we fail to observe certain signals. But they may be observable using more sensitive methods like invasive single cell recording. However, in both of the fMRI experiments mentioned (Liu et al. 2019; Binetti et al. 2020), when the researchers could not find the expected correlates within the visual areas, such correlates are found elsewhere (e.g., in the prefrontal cortex)—where sensitivity for neuroimaging measurements is typically weaker (Bhandari, Gagne, and Badre 2018). Therefore, the negative findings in the visual areas seem not to be due to measurement limitations alone.

2.7 Fronto-Parietal Network

So under the double-drift and subjective temporal duration illusions, the subjective percept seems to correlate well with activity in the prefrontal and

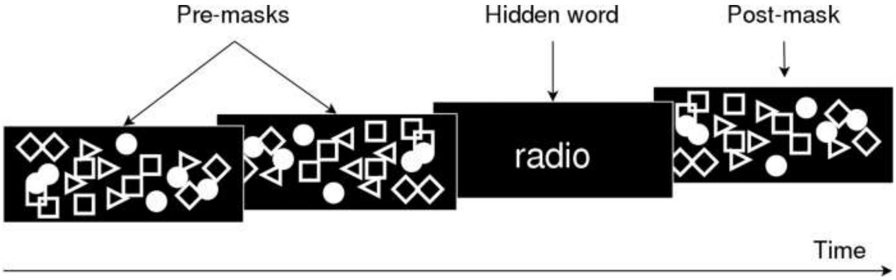

Figure 2.2 The visual masking technique
Reproduced with permission from Dehaene, S., Naccache, L., Cohen, L. et al. (2001). Cerebral mechanisms of word masking and unconscious repetition priming. *Nat Neurosci,* 4, 752–758. https://doi.org/10.1038/89551

parietal areas (Liu et al. 2019; Binetti et al. 2020). This is congruent with a very large body of work accumulated over decades, showing the importance of activity from these higher regions in consciousness (Rees, Kreiman, and Koch 2002; Dehaene 2014; Dehaene, Lau, and Kouider 2017; Odegaard, Knight, and Lau 2017; LeDoux, Michel, and Lau 2020).

In many of these studies, the researchers compared a "conscious" condition where the subjects saw the stimulus (sometimes called the "target"), with a "nonconscious" condition where the subjects did not see the stimulus. Comparing these conditions typically reveals robust and widespread activations in the prefrontal and parietal areas, that is, higher activity for the conscious condition.

For example, Dehaene et al. (2001) compared words that subjects were able to consciously see and recognize, against the same stimuli embedded in some distracting visual patterns, called forward (pre) and backward (post) visual "masks," which made the words invisible and unrecognizable (Figure 2.2). This comparison revealed some differences in the visual areas too. But the difference was more striking in the prefrontal and parietal areas. Under the nonconscious condition, the activity in these areas was near baseline level, as if these regions were engaged exclusively under conscious perception.

2.8 Stimulus Confounder

Global theorists, including Dehaene himself, have taken the results mentioned in Section 2.7 to support their views (Dehaene 2014; Dehaene, Lau, and Kouider 2017). But it could be argued that there was a stimulus confounder in these studies. That is, while the conscious and nonconscious conditions differed in perception, the stimuli presented also differed. So the observed activity could be driven by the physical stimulus difference rather than reflecting subjective experience itself.

There are two common ways to deal with this stimulus confounder. One is to present a stimulus at an intermediate intensity. We can use psychophysical methods to titrate the stimulus so that it will be presented at around detection threshold. At this level, sometimes subjects see it consciously, and sometimes they don't. Then after the experiment, we can sort out the "consciously seen" trials versus the "nonconscious" trials and make the comparison. This way the comparison will involve stimuli presented at the same physical intensity. Many studies employed this strategy and found evidence in support of the global view (e.g., Carmel, Lavie, and Rees 2006).

Another way is to make use of the phenomenon of bistable percepts, in which a static physical stimulus may give rise to different subjective percepts over time. One popular example is binocular rivalry (Blake, Brascamp, and Heeger 2014); see Figure 2.3. This occurs when we present different images to the two different eyes, using laboratory optical instruments. When set up correctly, most subjects will not see a static fusion of both images. Instead, they see one image for a few seconds, and then spontaneously a perceptual switch occurs, after which they see the other image for a few seconds, and then switch again, back and forth and so on. Typically, these perceptual shifts occur automatically, as in, beyond the subject's volitional control. When the subject sees one image but not the other, we say that the seen image is in dominance and the other image is under suppression.

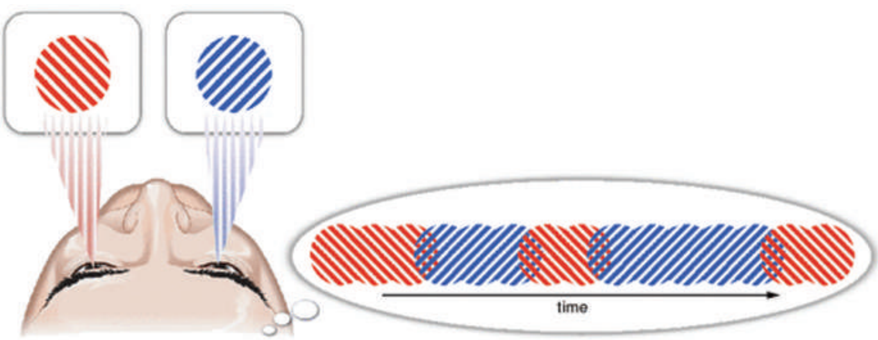

Figure 2.3 Binocular rivalry

Reproduced under a Creative Commons Attribution-NonCommercial 3.0 Unported (CC BY-NC 3.0) from Dieter, K., C., and Tadin, D. (2011). Understanding attentional modulation of binocular rivalry: a framework based on biased competition. *Front. Hum. Neurosci,* 5:155. doi: 10.3389/fnhum.2011.00155

Using binocular rivalry, many studies have found evidence in support of the global view too. With fMRI, many researchers have found an increase in activity in the prefrontal and parietal areas when the perceptual switch occurs (Rees, Kreiman, and Koch 2002). Using methods at a higher resolution, such as invasive neuronal or local field potential recording, it has been reported that activity in these areas can reflect the image in perceptual dominance too (Panagiotaropoulos et al. 2012). Let's say, for a hypothetical example, that the two images in rivalry are left-tilted and right-tilted line patterns, respectively, as shown in Figure 2.3. The logic of these studies is that we can first identify activity in the prefrontal cortex reflecting one image, say the left-tilted line patterns. When the subject sees the left-tilted line pattern, this activity would show an increase. When the same left-tilted line pattern is under suppression, that is when the subject sees the other pattern (the right-tilted lines), this activity would be low.

But others have taken similar studies as support for the local view because early visual activity also tracks the shifts in the subjective percept under binocular rivalry. Although the suppressed image tends to be reflected by some activity in the early visual areas too, the level of activity tends to be lower than that for the dominant image. This difference is in fact rather salient in the visual cortex. However, as we will see next, they may not reflect changes in subjective experience per se.

2.9 The Trickiest Confounder of All

Binocular rivalry elegantly addresses the problem of stimulus confounder. Accordingly, among consciousness researchers it is commonly regarded as an extremely important experimental tool. However, experts on binocular rivalry have also raised serious conceptual issues in interpreting these studies as reflecting consciousness (Blake, Brascamp, and Heeger 2014).

One problem is that when an image is dominant during rivalry, there are other consequences besides our conscious perception of it. For instance, in some studies of binocular rivalry, researchers asked subjects to detect a small change in the two images (Fox and Check 1968; Wales and Fox 1970; Norman, Norman, and Bilotta 2000). Unsurprisingly, subjects could do the task very well, if the change occurred in the dominant image; when we consciously see an image, detecting a small change in it should be easy. However, when the image was in suppression, subjects could also do the detection task above chance. That is, even though the subjects were consciously seeing the dominant image, they could nonetheless detect the small change in the suppressed image. Naturally, this performance was poorer than the performance for the dominant image. So we can think of this as a *performance-capacity confounder*. That is, for an image in dominance, we not only consciously see it but also are able to perform tasks better regarding the image content.

This confounder of performance capacity is important, and yet far from straightforward. In fact, some may think this is not a confounder at all. For example, one may argue that the confounder is only there when we ask subjects to detect the small change. There is no such task in typical binocular rivalry experiments and, therefore, no performance confounder to speak of. But this argument is a bit like saying that we can avoid a confounder by turning a blind eye to it. The emphasis here is on performance *capacity*. Presumably, there is an increased internal perceptual signal strength under perceptual dominance driving this *potential* performance superiority. Even when we don't measure it, the signal difference between dominance and suppression is there, and it could cloud what we may interpret as reflecting subjective experience.

Others may argue that having this superior performance capacity, and a stronger internal perceptual signal, is just a natural consequence of being conscious of the relevant perceptual information. This is what consciousness amounts to. Controlling for such confounders would be like controlling for the length of one's bones as we compare people of different heights; recall our example from Section 1.8 in the Chapter 1.

But there is also a strong disanalogy between these cases. When we compare people with different heights there is simply no way we could control for the length of people's bones. This is because tall people *always* have longer bones. This is what being tall is about. But when it comes to performance capacity, or internal perceptual signal strength, they are not always higher when one is conscious of the relevant perceptual information. We know this from blindsight, in which a lack of conscious experience is nonetheless associated with fairly high performance capacity, sometimes up to above 80% correct in a two-choice discrimination (Weiskrantz 1997). With physically degraded stimuli, which are nonetheless consciously seen, one often performs worse (Persaud et al. 2011). In Section 2.11, we will discuss more examples of how this performance-capacity confounder has been controlled for in actual experiments.

So performance capacity is an important, yet unobvious confounder. In fact, as we will see, it is particularly important *because* it is not obvious.

2.10 Binocular Rivalry as Red Herring

Within the settings of binocular rivalry, it is not so easy to control for performance capacity while assessing subjective experience. But one important study has shed great light on this issue.

Typically in binocular rivalry, we directly ask the subjects to report the percept to assess dominance versus suppression of the images. However, there are indirect methods too. For example, right after an image has been in dominance under rivalry, there are adaptation effects. Adaptation is a well-known phenomenon often exploited in psychophysics (Schwartz, Hsu, and Dayan 2007; Webster 2015). The idea is that after looking at a pattern for some time, an ambiguous stimulus presented at the same location is less likely to look like the pattern. Sometimes we say that the representation for the pattern is *adapted out*. It becomes less likely to win in a perceptual competition, when two alternative percepts are similarly plausible. Therefore, we can indirectly infer which image has been in dominance, based on how a subsequently presented ambiguous stimulus is perceived; it should look less like the image that has been in dominance.

The advantage of using this indirect method of assessing dominance is that we can now test if images rendered invisible may rival too. Zou et al. (2016) used a psychophysical method called flicker fusion to render some images invisible and presented them to the two different eyes. Subjects could not consciously see the images, and, therefore, they could not report what was in dominance or whether there was any rivalry at all. But Zou et al. found that the temporal profile of the adaptation after-effects resembled normal conscious rivalry. An invisible image went into dominance, suppression, and back, in cycles. This must mean that something akin to binocular rivalry happens for the invisible images too. Interestingly, these invisible flicker stimuli activated the visual areas, but not the prefrontal and parietal areas, as measured by their fMRI activity. The prefrontal and parietal areas reacted to the rivalrous stimuli only when they were consciously perceived.

To my mind, these findings from Zou et al. (2016) changed everything regarding the status of binocular rivalry as a tool for studying consciousness (Giles, Lau, and Odegaard 2016). They are very much in favor of the global over the local views. Perhaps these findings are exactly what one would expect, when we consider how the confounder of task performance capacity might have clouded previous results: previous findings of early visual activity reflecting binocular rivalry may just reflect nonconscious rather than conscious signal competitions. These nonconscious signals may compete to drive task performance, but on their own they are not sufficient for consciousness.

2.11 My First Attempt at Taking Down Global Theories

Earlier I mentioned that I am somewhat sympathetic to the global view. But as you will see, this affinity is rather limited. And in fact, it hasn't always been this way.

When I was a young postdoc in London around the mid-2000s, the global view was very much *en vogue*. Being naive and deluded, I took it as my calling to challenge the status quo. In graduate school I understood the point of blindsight, which is that subjective experience and task performance capacity can come apart. So, naturally, I thought we needed to deal with this tricky confounder of performance capacity.

Together with an outstanding undergraduate student, Navindra Persaud, and others, we scanned the blindsight patient GY using fMRI. With a V1 lesion restricted to the left hemisphere, GY was only blindsighted for the right visual field. This means that within the same subject, we can compare blindsight against normal vision (the left visual field). We presented a physically weaker stimulus to the normal, intact hemisphere, and the same stimulus at a higher luminance contrast to the damaged hemisphere. At these psychophysically titrated intensities, GY's performance score on a discrimination task was nearly equal for the two hemifields. And yet, he remained nonconscious of the stimulus in the blind field, while acknowledging that he saw the stimulus consciously about half of the time in the normal field. By comparing stimulation to the different hemispheres, that is, normal vision versus blindsight, we found activations in the prefrontal and parietal cortices, as they are typically found in studies of conscious perception. What was novel though, is that here there was no confounder of task performance capacity, as they were matched between the conditions (Persaud et al. 2011).

Of course, that was from a single patient. But we could do something conceptually similar in subjects from the general population too. Using another "trick" from psychophysics called metacontrast masking, we created conditions under which the subjects' performance in a simple discrimination task was matched. And yet, the target stimuli looked subtly different, as the timing of the masks differed. In one condition, subjects reported that they saw the targets more often than they did in the other condition. Again, comparing the two conditions as we recorded neural activity from the fMRI scanner, we found higher activity in the lateral prefrontal cortex in the condition where subjects said they consciously saw the targets more often (Lau and Passingham 2006).

These findings were somewhat surprising. When I set out to control for performance capacity, I was hoping to challenge the global view. I thought that much of the prefrontal and parietal activations could just reflect one's ability to do the tasks well. In particular, in the metacontrast masking study (Lau and Passingham 2006), the difference in subjective experience was subtle. Subjects only reported seeing the targets consciously about 10% more often in one condition over the other. Despite the fact that it was a visual manipulation (change of mask timing) leading to a visual difference, we couldn't find any activity difference in the visual cortex. This is not to say there was decidedly no difference. With fMRI there is always the issue of sensitivity. Using more invasive methods, I believe we could probably find some differences within the visual cortex too. But it is intriguing that with the limited sensitivity of fMRI, we found an activation in the lateral prefrontal cortex, but not in other areas.

2.12 Why So Hung Up About Performance-Capacity Confounders?

In many ways, the studies described in Section 2.11 are deeply flawed. When we matched for performance capacity, the physical stimulus was no longer matched (Lau and Passingham 2006). To directly control for both performance-capacity and stimulus confounders in the same experiments, one can make use of the phenomenon known as attentional blink. When subjects are asked to detect targets in a rapidly presented series of stimuli, perception is impaired shortly after the detection of a target, as if the subjects need some time to recover. Therefore, when two targets are presented close to each other in time, perception for the second target suffers. Interestingly, this effect seems to be stronger on subjective reports of perception, rather than task-performance capacity itself (Pincham, Bowman, and Szucs 2016; Recht, Mamassian, and de Gardelle 2019). So, like in the study by Lau and Passingham (2006), this provides an opportunity for us to dissociate the two.

Using this task paradigm, Pincham et al. (2016) have measured electrophysical activity on the scalp, and looked at event-related potentials (ERP), a simple and conventional method for analyzing electroencephalogram (EEG) data. In particular, they looked at the ERP time locked to the second target, the visibility of which changed as a function of the temporal distance from the first target, while the physical stimulus itself (for the second target) remained the same. They found that the late-stage ERP component known as P3 reflected subjective visibility better than objective task performance. This is broadly consistent with the claim that subjective visibility is associated with downstream processing beyond the sensory cortices. But as we will see in Section 2.13, this P3 component may be explained by other confounders.

In the case of blindsight patient GY, we were not only comparing conscious versus nonconscious vision. We were also comparing an intact versus a damaged hemisphere, and that was admittedly a huge confounder too (Persaud et al. 2011).

As such, overall, one could rightfully accuse us of merely replacing one set of confounders (task performance capacity) with several others. But turns out, this may not be such a bad thing.

This may sound counterintuitive but having an obvious confounder in one's study is not necessarily so problematic. It may be bad for the individual researcher trying to publish the work. But as a field, we look for converging evidence. Take the case of the stimulus confounder in studies showing that the prefrontal and parietal areas may be involved in conscious perception. For example, Dehaene et al.'s study using visually masked words (2001). It is true that the stimuli were not perfectly matched. But because the confounder was relatively obvious, other studies have dealt with this problem by presenting a constant stimulus at around threshold, as well as binocular rivalry. The converging evidence is that the activations in the prefrontal and parietal cortices under these conditions hold true. As such, these activations are unlikely to be due to the stimulus confounder alone. Otherwise, studies controlling for the confounder wouldn't have shown the same results. Therefore, this stimulus confounder can be considered less fatal to Dehaene et al.'s study. The main conclusion holds after all.

Contrast this with confounders that are so inconspicuous that the absolute majority of studies in the field do not recognize the need to control for them. The converging trends from all of these studies could well be commonly driven by these confounders, and we would be none the wiser. So the main conclusions, *apparently* replicated over and over, could be at risk.

I worry that performance capacity is exactly one such confounder. Among hundreds if not thousands of studies on the NCC, I only know of a few which took the trouble to look into that at all. There is a real danger there. So it is particularly informative that when the task performance capacity confounder was controlled for, the results were clearly in favor of the global over the local view.

This may sound like a funny way to defend one's own flawed studies. They are definitely imperfect, and I don't argue otherwise. But perhaps one could take this as an argument in favor of diversity in thinking within a field; it may be good that different people worry about different confounders and experimental issues, and we see how the evidence converges overall. As we will see in Chapter 3, these studies matching performance capacity did end up leading to new ideas and findings that together form a coherent story. (Otherwise, I would not be discussing them here, would I?)

2.13 Reports?

The above consideration is also relevant for one recent trend in the literature. So in many of the experiments discussed in the chapter, the subjects were required to make a report, either about the stimulus, or about the experience itself. In everyday life, we don't need to do that. So the task of reporting can be considered a confounder too. Specifically, some of the activity in the prefrontal and parietal cortices may be partially accounted for by this reporting demand, rather than reflecting subjective experience per se (Tsuchiya et al. 2015).

While this is certainly a confounder worth controlling for, the trouble is this is sometimes taken as a "new standard," with the "no-report paradigm" becoming a buzz phrase of some sort. But of course, the recognition of this confounder is nothing new. It has been controlled for in neuroimaging studies of the NCC as early as over 2 decades ago (Lumer and Rees 1999). The typical method is to remove the need for reporting and infer subjective experience in some other indirect ways. Alternatively, one asks the subjects to report on a different stimulus, or a different feature of the stimulus, which is irrelevant to the comparison in questions (e.g., Tse et al. 2005; Mante et al. 2013). In one of such studies using binocular rivalry without report, it was shown that prefrontal and parietal fMRI activations survived such controls (Lumer and Rees 1999). And recent studies using methods with higher sensitivity have repeatedly confirmed this finding in both humans (Vidal et al. 2014; Noy et al. 2015; Huth et al. 2016; Taschereau-Dumouchel, Kawato, and Lau 2019) and monkeys (Panagiotaropoulos et al. 2012; Mante et al. 2013; Panagiotaropoulos, Dwarakanath, and Kapoor 2020).

When reports were not required, there are null findings for prefrontal and parietal activations too (Tse et al. 2005; Kouider et al. 2007; Frässle et al. 2014). But in the light of numerous positive findings, individual negative findings need to be interpreted with care. As a matter of basic statistical consideration, the absence of evidence is not evidence of evidence. Especially with traditional mass-univariate fMRI, sensitivity is limited. It is "easy" to obtain a falsely negative result. Therefore, it would be problematic to claim that these activations are due to nothing but reporting, based on individual null findings. It is particularly troubling when this is presented as if it is a novel, trend-setting discovery.

As we have seen in most of the studies discussed in this chapter, it is probably impossible to have all of the confounders controlled for at once within a single experiment (Michel and Morales 2019). Such a "perfect" experiment probably doesn't exist. In particular, it is not clear how one can control for both the stimulus and performance-capacity confounders under the same conditions. Adding in the reporting confounder makes things more challenging still. As such, fixation on a single confounder is unhelpful. When a certain confounder is considered dealt with, it may be strategically more advantageous to move on and address other confounders. This way, as a field, we can provide more meaningful converging answers.

Having said that, I do think some lessons can be learned from this recent focus on controlling for reports. Earlier I cited studies using invasive neuronal recordings (Panagiotaropoulos et al. 2012; Mante et al. 2013; Panagiotaropoulos, Dwarakanath, and Kapoor 2020). From those studies, it seems like the findings from the prefrontal cortex have survived the control for reports without problem. For other invasive recordings (that are not directly neuronal, but concern local field potentials and related signals), the findings too seem to hold up. However, when reports weren't required, the effect size was significantly reduced (Noy et al. 2015). For fMRI, the story seems similar (Huth et al. 2016). With reduced effect size, occasional null results are exactly as expected.

Earlier I mentioned that the P3 component of the ERP may survive the control of both stimulus and task performance-capacity confounders (Pincham, Bowman, and Szucs 2016). But unfortunately, P3 seems to no longer be present after controlling for task relevance and reporting (Cohen et al. 2020). But once again, this may in part depend on data analysis methods. Using multivariate "decoding" approaches on EEG data, rather than ERP, others have found late-stage correlates, likely reflecting prefrontal activity, which survived similar controls (Sergent et al. 2021).

So the negative findings reviewed here do not falsify the global view. They seem all easily explained by the limited sensitivity of some specific methods. With sufficient sensitivity numerous positive findings were reported. However, it does raise the question of whether the activity in the prefrontal and parietal areas during conscious perception is as widespread and strong as previously thought. If such activity is subtle, the notion that it supports "global ignition" and "broadcast" may seem less plausible (Noy et al. 2015).

2.14. Saying No to No-Cognition

Regarding confounders, Ned Block (2019) has recently suggested we take things one step further: We should employ what he calls "no-cognition paradigms," to rule out that our findings on the NCC may just be due to postperceptual cognitive processes.

The basic rationale seems attractive. It is true that when we consciously see something we tend to think about it, remember it, and contemplate relevant thoughts. Some of these processes are not constitutively part of the subjective experience itself. So it would be good to control for them. But as we see in the last section, when formulated as a general prescription, this kind of "no-XYZ" paradigms could end up doing more harm than good.

If the issue is about postperceptual cognitive processes, this confounder has already been dealt with in some studies. In one study (Mante et al. 2013), the monkeys had to focus on one feature of a stimulus (color or motion direction of some moving dots, depending on the experimental block), and ignore the other feature. The task was difficult, with some stimuli that were rather ambiguous. So actively thinking about the irrelevant feature would be disadvantageous, and the behavior data confirmed that such cognitive interference from the irrelevant feature was in fact small. This was conducted while the researchers recorded signals directly from many neurons in the prefrontal cortex. They found that the neuronal activity reflected the task irrelevant feature—almost as well as it did for the task relevant feature. Even if the monkey did think about the irrelevant feature now and then, the cognitive processes involved were unlikely to be so consistently strong to account for the findings.

In other studies, the monkeys did not have to perform any overt task. They just passively viewed a binocular rivalry display, where perceptual dominance was inferred indirectly. Again, prefrontal activity clearly tracked the dominant percept. Block (2020) argued that perhaps post–perceptual cognition could explain the findings: although the monkeys were not required to think about the percept, they might be "bored" and did so anyway. But as Fanis Panagiotaropoulos and colleagues (2020) pointed out, if the monkeys engaged in such thinking out of boredom, one should expect this to occur somewhat randomly, as wandering thoughts are. But the electrophysiological findings were highly consistent over trials.

Also, prefrontal activity during binocular rivalry often occurs prior to the switch (Dwarakanath et al. 2020). Others have shown that disrupting prefrontal activity causally influences perceptual rivalry as well (Vernet et al.

2015). So, such activity is unlikely to reflect cognition in *reaction* to the perceptual switch. Instead, it seems to be part of the causal process (Weilnhammer et al. 2021).

So we can rule out that the findings were merely due to the animals' explicitly thinking about the percept, in ways that are not constitutively part of the subjective experience itself (i.e., postperceptual thinking). But if the demand is that we need to rule out *any* kind of cognitive process, as the phrase "no-cognition" may suggest, we have to be careful. Recall from Section 1.6 of Chapter 1: We cannot define *consciousness* as entirely independent from cognitive access from the outset. Subjective experience may turn out to constitutively involve some degree of cognition. We have to give such empirical possibilities a fair chance.

Or perhaps Ned Block's intention was not to explicitly beg the question, by ruling out any form of cognition from the outset. Perhaps his point was that we should focus on perceptual events that the subjects do not explicitly "notice," in order to minimize the contamination from attention. But the question of whether unnoticed or unattended perceptual stimuli contribute to our subjective experience is controversial. We will address this in Chapter 4.

2.15 Chapter Summary

This chapter is about the first of the five issues setting the global and local views apart: the neural substrate for subjective experience. We've covered a lot of ground. Some summary is in order.

The central point here revolves around the theme of experimental confounders. Some confounders are unfortunately more "fashionable" than others. But trends come and go. We have to think about what lasts and what matters. Because focusing on addressing a certain confounder will favor some findings over others, this can impact our theoretical views as well. So these discussions can get sectarian pretty quick. We need to see the futility in fighting over which confounders are more important. All confounders are important, and they all need to be addressed. But because it is difficult to address all of them at once in a single experiment, the best way is to consider them in turn, one after another, to aim for an overall broad picture.

If we focus on stimulus and report confounders alone, the overall findings may appear to be in favor of local views. Early sensory activity seems to reflect subjective experiences, even when these two confounders are controlled for: for example, in experiments using binocular rivalry.

The problem, though, is with the confounder of performance capacity. In an experiment where performance capacity was explicitly matched, early sensory activity no longer reflected the difference in reported subjective experience between conditions. Unfortunately, too few studies have controlled for performance capacity, so this negative finding alone may seem indecisive. But in nonconscious binocular rivalry, the early visual activity seems to behave just as it did in ordinary conscious rivalry (Jiang, Zhou, and He 2007). This means that such activity is not always conscious. It may very well just reflect nonconscious internal perceptual signals supporting performance capacity rather than subjective experience. And then, there are also some cases of visual illusions, in which such early visual activity did not track subjective perception well at all. So in sum, it is difficult to say that early sensory activity alone is the NCC.

The global view also does fine regarding the stimulus confound. Many binocular rivalry and near-threshold presentation experiments support the role of prefrontal and parietal activity in subjective perception. Compared to the local view, it fares considerably better with respect to the confounder of task-performance capacity. Admittedly the view also faces considerable challenges with the confounder of reports. But some activity, in particular in the lateral prefrontal cortex, survives such controls—even if such activity may be more subtle and less widespread than previously thought.

These neuroimaging and invasive electrophysiological studies give a rather different picture from what one may intuit based on lesion studies alone. The classical neurological phenomenon of blindsight is associated with damage to V1. However, lesions can have distal effects. We have also gone through the conceptual issues involved in the very definition of the NCC. The main lesson is that considering lesion studies in isolation can be misleading. Accordingly, Larry Weiskrantz, who coined the term *blindsight*, also argued that the mechanisms for subjective visual awareness may well reside within the prefrontal cortex rather than V1 (1997).

In Chapter 3 we will discuss more about lesions and focus on damages to the prefrontal cortex in particular.

References

Bhandari A, Gagne C, Badre D. Just above chance: Is it harder to decode information from prefrontal cortex hemodynamic activity patterns? *J Cogn Neurosci* 2018;**30**:1473–1498.

Binetti N, Tomassini A, Friston K et al. Uncoupling sensation and perception in human time processing. *J Cogn Neurosci* 2020;**32**:1369–1380.

Blake R, Brascamp J, Heeger DJ. Can binocular rivalry reveal neural correlates of consciousness? *Philos Trans R Soc Lond B Biol Sci* 2014;**369**:20130211.

Block N. What is wrong with the no-report paradigm and how to fix it. *Trends Cogn Sci* 2019;**23**:1003–1013.

Block N. Finessing the bored monkey problem. *Trends Cogn Sci* 2020;**24**:167–168.

Boyer JL, Harrison S, Ro T. Unconscious processing of orientation and color without primary visual cortex. *Proc Natl Acad Sci U S A* 2005;**102**:16875–16879.

Bridge H, Harrold S, Holmes EA et al. Vivid visual mental imagery in the absence of the primary visual cortex. *J Neurol* 2012;**259**:1062–1070.

Carmel D, Lavie N, Rees G. Conscious awareness of flicker in humans involves frontal and parietal cortex. *Curr Biol* 2006;**16**:907–911.

Chalmers DJ. What is a neural correlate of consciousness? In: Metzinger T (ed.). *Neural Correlates of Consciousness: Empirical and Conceptual Questions*. MIT Press, 2000.

Cohen MA, Ortego K, Kyroudis A et al. Distinguishing the neural correlates of perceptual awareness and postperceptual processing. *J Neurosci* 2020;**40**:4925–4935.

Craver CF. *Explaining the Brain: Mechanisms and the Mosaic Unity of Neuroscience*. Oxford University Press, 2007.

Dehaene S. *Consciousness and the Brain: Deciphering How the Brain Codes Our Thoughts*. Penguin, 2014.

Dehaene S, Lau H, Kouider S. What is consciousness, and could machines have it? *Science* 2017;**358**:486–492.

Dehaene S, Naccache L, Cohen L et al. Cerebral mechanisms of word masking and unconscious repetition priming. *Nat Neurosci* 2001;**4**:752–758.

Dwarakanath A, Kapoor V, Werner J et al. Prefrontal state fluctuations control access to consciousness. *bioRxiv* 2020.01.29.924928. https://doi.org/10.1101/2020.01.29.924928.

Eagleman DM, Jacobson JE, Sejnowski TJ. Perceived luminance depends on temporal context. *Nature* 2004;**428**:854–856.

Fisch L, Privman E, Ramot M et al. Neural "ignition": Enhanced activation linked to perceptual awareness in human ventral stream visual cortex. *Neuron* 2009;**64**:562–574.

Fox R, Check R. Detection of motion during binocular rivalry suppression. *J Exp Psychol* 1968;**78**:388–395.

Frässle S, Sommer J, Jansen A et al. Binocular rivalry: frontal activity relates to introspection and action but not to perception. *J Neurosci* 2014;**34**:1738–1747.

Giles N, Lau H, Odegaard B. What type of awareness does binocular rivalry assess? *Trends Cogn Sci* 2016;**20**:719–720.

Gotts SJ, Chow CC, Martin A. Repetition priming and repetition suppression: A case for enhanced efficiency through neural synchronization. *Cogn Neurosci* 2012;**3**:227–237.

Huang L, Wang L, Shen W et al. A source for awareness-dependent figure-ground segregation in human prefrontal cortex. *Proc Natl Acad Sci U S A* 2020;**117**:30836–30847.

Huth AG, Lee T, Nishimoto S et al. Decoding the semantic content of natural movies from human brain activity. *Front Syst Neurosci* 2016;**10**:81.

Jiang Y, Zhou K, He S. Human visual cortex responds to invisible chromatic flicker. *Nat Neurosci* 2007;**10**:657–662.

Kouider S, Dehaene S, Jobert A et al. Cerebral bases of subliminal and supraliminal priming during reading. *Cereb Cortex* 2007;**17**:2019–2029.

Lamme VA. Blindsight: The role of feedforward and feedback corticocortical connections. *Acta Psychol* 2001;**107**:209–228.

Lamme VAF. Why visual attention and awareness are different. *Trends Cogn Sci* 2003;**7**:12–18.

Lau HC, Brown R. The emperor's new phenomenology? The empirical case for conscious experiences without first-order representations. In A Pautz and D Stoljar (Eds), Blockheads! MIT Press, 2019, Chapter 11.

Lau HC, Passingham RE. Relative blindsight in normal observers and the neural correlate of visual consciousness. *Proc Natl Acad Sci U S A* 2006;**103**:18763–18768.

LeDoux JE, Michel M, Lau H. A little history goes a long way toward understanding why we study consciousness the way we do today. *Proc Natl Acad Sci U S A* 2020;**117**:6976–6984.

Liu S, Yu Q, Tse PU et al. Neural correlates of the conscious perception of visual location lie outside visual cortex. *Curr Biol* 2019;**29**:4036–4044.e4.

Lumer ED, Rees G. Covariation of activity in visual and prefrontal cortex associated with subjective visual perception. *Proc Natl Acad Sci U S A* 1999;**96**:1669–1673.

Macknik SL, Martinez-Conde S. The role of feedback in visual masking and visual processing. *Adv Cogn Psychol* 2008;**3**:125–52.

Malach R. Conscious perception and the frontal lobes: Comment on Lau and Rosenthal. *Trends Cogn Sci* 2011;**15**:507; author reply 508–9.

Mante V, Sussillo D, Shenoy KV et al. Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature* 2013;**503**:78–84.

Mazzi C, Savazzi S, Silvanto J. On the "blindness" of blindsight: What is the evidence for phenomenal awareness in the absence of primary visual cortex (V1)? *Neuropsychologia* 2019;**128**:103–108.

Michel M. Consciousness science underdetermined: A short history of endless debates. *Ergo* 2019;**6**. https://doi.org/10.3998/ergo.12405314.0006.028.

Michel M, Doerig A. A new empirical challenge for local theories of consciousness. *Mind Lang* 2021. https://doi.org/10.1111/mila.12319.

Michel M, Lau H. On the dangers of conflating strong and weak versions of a theory of consciousness. *PhiMiSci* 2020. https://doi.org/ 0.33735/phimisci.2020.II.54.

Michel M, Morales J. Minority reports: Consciousness and the prefrontal cortex. *Mind Lang* 2019, https://doi.org/10.1111/mila.12264.

Miller SM. Closing in on the constitution of consciousness. *Front Psychol* 2014;**5**:1293.

Norman HF, Norman JF, Bilotta J. The temporal course of suppression during binocular rivalry. *Perception* 2000;**29**:831–841.

Noy N, Bickel S, Zion-Golumbic E et al. Ignition's glow: Ultra-fast spread of global cortical activity accompanying local "ignitions" in visual cortex during conscious visual perception. *Conscious Cogn* 2015;**35**:206–224.

Odegaard B, Knight RT, Lau H. Should a few null findings falsify prefrontal theories of conscious perception? *J Neurosci* 2017;**37**:9593–9602.

Panagiotaropoulos TI, Deco G, Kapoor V et al. Neuronal discharges and gamma oscillations explicitly reflect visual consciousness in the lateral prefrontal cortex. *Neuron* 2012;**74**:924–935.

Panagiotaropoulos TI, Dwarakanath A, Kapoor V. Prefrontal cortex and consciousness: Beware of the signals. *Trends Cogn Sci* 2020;**24**:343–344.

Pascual-Leone A, Walsh V. Fast backprojections from the motion to the primary visual area necessary for visual awareness. *Science* 2001;**292**:510–512.

Persaud N, Davidson M, Maniscalco B et al. Awareness-related activity in prefrontal and parietal cortices in blindsight reflects more than superior visual performance. *Neuroimage* 2011;**58**:605–611.

Pincham HL, Bowman H, Szucs D. The experiential blink: Mapping the cost of working memory encoding onto conscious perception in the attentional blink. *Cortex* 2016;**81**:35–49.

Place UT. Is consciousness a brain process? *Br J Psychol* 1956:47: 44–50.

Recht S, Mamassian P, de Gardelle V. Temporal attention causes systematic biases in visual confidence. *Sci Rep* 2019;**9**:11622.

Rees G, Kreiman G, Koch C. Neural correlates of consciousness in humans. *Nat Rev Neurosci* 2002;**3**:261–270.

Ro T, Harrison S, Boyer J et al. Unconscious orientation and color processing without primary visual cortex. *J Vis* 2010;**5**:285–285.

Sandberg K, Bahrami B, Kanai R et al. Early visual responses predict conscious face perception within and between subjects during binocular rivalry. *J Cogn Neurosci* 2013;**25**:969–985.

Schwartz O, Hsu A, Dayan P. Space and time in visual context. *Nat Rev Neurosci* 2007;**8**:522–535.

Schwarzkopf DS, Schindler A, Rees G. Knowing with which eye we see: Utrocular discrimination and eye-specific signals in human visual cortex. *PLOS One* 2010;**5**:e13775.

Sergent C, Corazzol M, Labouret G et al. Bifurcation in brain dynamics reveals a signature of conscious processing independent of report. *Nat Commun* 2021;**12**:1149.

Silvanto J. Why is "blindsight" blind? A new perspective on primary visual cortex, recurrent activity and visual awareness. *Conscious Cogn* 2015;**32**:15–32.

Taschereau-Dumouchel V, Kawato M, Lau H. Multivoxel pattern analysis reveals dissociations between subjective fear and its physiological correlates. *Mol Psychiatry* 2019. https://doi.org/10.1038/s41380-019-0520-3.

Tong F. Primary visual cortex and visual awareness. *Nat Rev Neurosci* 2003;**4**:219–229.

Tse PU, Martinez-Conde S, Schlegel AA et al. Visibility, visual awareness, and visual masking of simple unattended targets are confined to areas in the occipital cortex beyond human V1/V2. *Proc Natl Acad Sci U S A* 2005;**102**:17178–17183.

Tsuchiya N, Wilke M, Frässle S et al. No-Report Paradigms: Extracting the True Neural Correlates of Consciousness. *Trends Cogn Sci* 2015;**19**:757–770.

Vernet M, Brem A-K, Farzan F et al. Synchronous and opposite roles of the parietal and prefrontal cortices in bistable perception: A double-coil TMS-EEG study. *Cortex* 2015;**64**:78–88.

Vidal JR, Perrone-Bertolotti M, Kahane P et al. Intracranial spectral amplitude dynamics of perceptual suppression in fronto-insular, occipito-temporal, and primary visual cortex. *Front Psychol* 2014;**5**:1545.

Wales R, Fox R. Increment detection thresholds during binocular rivalry suppression. *Percept Psychophys* 1970;**8**:90–94.

Webster MA. Visual Adaptation. *Annu Rev Vis Sci* 2015;**1**:547–567.

Weilnhammer V, Fritsch M, Chikermane M et al. An active role of inferior frontal cortex in conscious experience. *Curr Biol* 2021;**31**:2868–2880.e8.

Weiskrantz L. *Consciousness Lost and Found*. Oxford University Press, 1997.

Zhang X, Zhaoping L, Zhou T et al. Neural activities in v1 create a bottom-up saliency map. *Neuron* 2012;**73**:183–192.

Zhou H, Davidson M, Kok P et al. Spatiotemporal dynamics of brightness coding in human visual cortex revealed by the temporal context effect. *Neuroimage* 2020;**205**:116277.

Zou J, He S, Zhang P. Binocular rivalry from invisible patterns. *Proc Natl Acad Sci U S A* 2016;**113**:8408–8013.