

Richard E. Cytowic

## *The Clinician's Paradox*

*Believing Those You Must Not Trust*

**Abstract:** *Clinicians have a convention whereby **symptoms** are subjective statements 'as told by' patients, whereas **signs** are outwardly observable facts. Yet both first-person reports and third-person observations are theory laden and can bias conclusions. Two aspects of the oft-mentioned unreliability of reports are the subject's interpretation of them and the experimenter's assumptions when translating introspective reports into scientifically useful characterizations. Meticulous training of introspectors can address their mischief, whereas investigators can become more attentive to their own theory-laden biases. In the case of hallucinations, for example, ignoring some customary third-person constructs and focusing on the visual experience itself has led to fresh explanations of visual symptoms based on cortical physiology rather than conceptual categories. Constructs that historically have ignored the subject's state of mind are also problematic; an example is the so-called resting state during metabolic brain imaging, long believed to reflect a blank mental slate.*

*Introspective reports, not accepted literally but properly interpreted and revised by investigators as necessary, are legitimate sources of data.*

When practising as either a neurologist or neuropsychologist, a great deal of data upon which I draw inferences and make decisions is clinical, understood as meaning generally subjective and hence unreliable in the eyes of an ideal objectivist. Clinical convention conceptualizes *symptoms* — such as pain, vertigo or forgetfulness — as subjective statements 'as told by' patients, whereas *signs* — such as paralysis, seizure or agrammatical errors — are objective, outwardly observable facts.

In practice, the clinician's milieu is muddled. Not much of what we have to work with is purely objective because we are often not around to witness matters first hand and so must reconstruct events, sometimes from second-hand sources. Looking at the other extreme, we realize that even the most subjective

Correspondence: Richard E. Cytowic MD, 4720 Blagden Terrace NW, Washington DC 20011-3720, USA. Email: [Richard@Cytowic.net](mailto:Richard@Cytowic.net)

*Journal of Consciousness Studies*, 10, No. 9–10, 2003, pp. ??–??

symptoms, such as pain, are coupled to objective signs (in this case, say, increased sympathetic activity). Neurologists and kindred clinicians, like me, wonder if the subjective–objective dichotomy that so exercises philosophers is overwrought. The distinction does not appear sharp to us because we operate within a paradox: we must rely on patient reports, yet cannot trust them fully.

Part of the problem is that patients frequently *interpret* events instead of *reporting* them straightforwardly as one would wish ideally. As this is an occupational fact, clinicians have traditionally learned to look beyond face value when translating first-person accounts into a third-person understanding of how the nervous system operates. In so doing, however, examiners are prone to bias introspective reports with their own theory-laden assumptions.

Consider an anecdote and two responses to it, which I caricature as the objectively oriented American and the subjectively oriented British. The anecdote is that a well-known newspaper editor falls at home and now seems ‘different’ to his worried family. The editor protests that he is fine and wants to go home. The Yank, whose schedule runs late as usual, pops into the exam cubicle, hears the wife’s statement, ‘he’s confused’, and responds, ‘let’s get a scan’, dispatching the patient into the tunnel of technology without further ado.

Our Brit is schooled in the ‘clinical method’ as exemplified by the clinicopathological approach taught at a renowned institution such as Queen Square. The editor’s reputation and implied level of education predisposes the Brit favourably. He judges the patient’s comportment and conversation to be normal; in fact, his vocabulary sounds quite rich. What could the wife mean? Looking for evidence of confusion but finding none, the Brit’s empirical knowledge that vocabulary does not falter until late in the course of many brain pathologies, when physical signs may be already apparent, makes him discount the smooth talk and probe further. Asked the reason for falling, the editor says, ‘I lost my glasses . . . I was going to the garage’. Incongruent remarks are exactly the kind of signal our Brit has learned that signify something amiss in cognitive or language domains (e.g. a dementia or aphasia). When drawn out of him, the editor’s narrative turns implausible, fragmented, and shows that he does not grasp the meaning of the immediate situation. A social history from the wife uncovers heavy drinking and skipped meals. Next, the suspicion of anterograde and retrograde amnesia is quickly confirmed by a simple procedure, as is the presence of ataxia when tandem walking is tested — our clinician is now confident that ataxia caused the fall and that the mental disorder has gone unnoticed until now, a textbook example of ‘decompensation’.

Unlike his American colleague who asked for an objective test straightaway, the British clinician applies his knowledge base and examination manoeuvres flexibly as the clinical situation unfolds, and so uncovers Wernicke-Korsakoff syndrome in the editor. His prepared mind was not misled by statements, but deduced a diagnosis (Latin: *dia* = through + *gnosis* = knowledge) that the editor is a so-called maintenance alcoholic with thiamine deficiency causing brain damage that conforms to a recognized syndrome. Should this clinician request a scan, he will be alert for specific lesions in the diencephalon or note that cerebral

atrophy exists in a *recognized pattern* involving the vermis, interhemispheric fissure and temporal fissures (Victor *et al.*, 2000).

When the Yank's scan report comes back saying, 'mild atrophy', he may set off on an expensive fishing expedition of more tests to follow up this objective but unhelpful finding, whereas the Brit may be satisfied with his diagnosis and search no further, although possibly at the cost of failing to uncover a different etiology for the editor's condition. This possibility is small, although time will tell. Note also that the disinterested radiologist, assumed to be objective in the sense of being a neutral third party, is like the Yank in having no context 'in what way' the patient is 'confused', as the requisition slip states. His anatomical interpretation of the scan is limited to a general statement. Provided context, however, he will *view the image differently* now that he knows what to look for, perhaps saying that the pattern of atrophy observed is 'consistent with alcoholic cerebral degeneration', a clinically more useful statement to all concerned.

These caricatured scenarios surrounding a common syndrome raise issues of preparation, expectation, assumptions and social context. Such factors can influence both subjective and objective observations. Examples supporting the absolute fallibility of verbal reports are easy to come by. For example, split-brain individuals demonstrate the existence of cognition not accessible to language as well as serving as an instance wherein verbal reports can actively mislead both patient and clinician (Zaidel *et al.*, 2003). Counterexamples claiming verbal reports as crucial to scientific understanding are also plentiful. Dreaming during REM sleep is an obvious one: if no one had awakened sleepers during this objective yet enigmatic EEG event, the meaning of REM sleep would have remained undeciphered.

Although my scenarios above refer to the intellectual process by which clinicians make medical judgments (an area with its own literature), I will not focus on it. Rather, I was invited to comment on *the necessity of clinicians and experimenters to adopt a different interpretation from that of their subject*. This implies that clinical reports can never be accepted verbatim; rather, their meaning must be drawn out and elicited. To use Henry James's term, they are 'rendered', involving the clinician as much as a participating artist as he is a detached amanuensis. After all, there must be some reason behind the phrase 'art of medicine'. Does it help, or get in the way?

I consider that (1) introspections are subject to two kinds of biases: (a) subjects interpret their experiences instead of factually reporting them, and (b) investigators fail to recognize that their assumptions and interpretations are theory-laden. As a corollary, scientists cannot accept introspective reports literally. Accordingly: (2) (a) can be addressed by training subjects; (3) both (a) and (b) can be lessened by using a script; and (4) other sources of so-called objective data (e.g. neural metabolism) are also biased by (b), with no reason to suppose that (b) is any more problematic for introspective reports than it is for other kinds of data.

It follows that introspective reports, not accepted literally but properly interpreted and revised by investigators as necessary, are legitimate sources of data.

Jack and Shallice (2001) comment: ‘Introspective evidence always arrives interpreted . . . all descriptions of experience, no matter how basic, carry implicit theoretical commitments of one sort or another.’ They concur with the general statement, typically said with regard to objectivity, that observation in support of theories is itself always theory-laden. There are ways to address this when assessing first-person reports, particularly those involving perceptual issues where the facts of physiology are often counterintuitive to common sense (so-called ‘folk psychology’ accounts of vision, for instance, are ignorant of the vertical meridian or the quadrants in the visual fields, indeed of the very concept of a visual field).

Consider the experience of visual hallucinations. In describing Heinrich Klüver’s work starting in the 1920s and spanning many decades, I recounted how he was initially ‘frustrated by the vagueness with which subjects described their experience, their eagerness to yield uncritically to cosmic or religious interpretations, to ‘interpret’ or poetically embroider the experience in lieu of straightforward but concrete description, and their tendency to be overwhelmed and awed by the ‘indescribability’ of their visions’ (Cytowic, 1997, p. 29). He intuited that the novelty of what subjects saw and the vivid colours grabbed their attention more than did the spatial configurations. Because people ordinarily have no need to describe their conscious experience, they do not pay attention to its subtle features. By training his observers, however, Klüver (1942; 1966) eventually did succeed in identifying three categories of visual pathology. Among these are his well-known ‘form constants’ encompassing four varieties of consistent hallucinogenic configuration: gratings and honeycombs, cobwebs, tunnels and cones, and spirals. (His other two categories regard variation in the number, size and shape of objects, and variation in their spatial and temporal relations.)

Others replicated and extended Klüver’s work (Siegel and Jarvik, 1975; Horowitz, 1964; 1975). Recently, ffytche and Howard (1999) confirmed these form constants and other visual experiences across a range of clinical conditions, showing how ‘similarities between the visual experiences of unrelated clinical and experimental contexts’ are meaningful in neuroscientific terms. As Klüver foreshadowed, ‘the occurrence of these [identical] symptoms in aetiologically different conditions suggests that we are dealing with some fundamental mechanisms involving various levels of the nervous system’. To achieve this, ffytche and Howard ignored customary third-person constructs of perception that are typically used to describe visual pathology — these being examples of investigator bias (b). For instance, they did not distinguish between hallucinations (experiences that have no afferent input) and illusions (experiences arising from falsely interpreted afferents), or between hallucinations wherein insight is preserved (i.e. patients know that they are hallucinating, as in the Charles Bonnet syndrome) and those where insight is lost (e.g. in schizophrenia, where patients cannot grasp the unreality of their visual experience). By ignoring these traditional theoretical constructs of how patients are ‘supposed to’ experience their hallucinations and focusing instead on the final common pathway of the visual experience itself, the authors were *able to identify three new categories of abnormal*

*vision*. Noting next that they observed identical abnormalities in etiologically different conditions and experimental settings, they related each new category of visual experience to known visual cortex physiology and were able to offer a biologically based classification of positive and negative visual symptoms — something the customary third-person constructs had failed to do.

These examples stress two aspects of the oft-mentioned unreliability of reports. The more familiar is the subject's bias (a): I use the term 'embroider' to describe how they explain, interpret or rationalize their experience based on assumptions usually very different from those of neuroscience. The other side concerns the experimenter's assumptions (b), either individually or as a member of the collective neuroscience community, when translating introspective reports into scientifically useful characterizations.

The opportunity to embroider is especially ripe in synesthesia because the experience is so ineffable, hence difficult to convey. Subjects often resort to metaphor when describing 'What it is like'. For example, my index case MW, in whom taste and smell evoked tactile perceptions predominantly in the hands and face, described the taste of spearmint as 'cool glass columns'. Was he being metaphorical, or verbally interpreting a sensory experience by translating tactile sensations into images I could understand? I sought to tease this out by asking him to describe the exact sensations he felt, rather than 'what it is like'. When given the stimulus, MW rubbed his fingertips together and moved his hand through the air as if palpating an actual object, saying:

I feel a round shape. There's a curvature behind which I can reach, and it's very, very smooth. So it must be made of marble or glass, because what I'm feeling is this incredible satiny smoothness. There are no ripples, no little surface indentations, so it must be glass because if it were marble, I would be able to feel the roughness of the stone or the pits in the surface. It's also very cool so it has to be some sort of glass or stone material because of the temperature. What is so wonderful is the absolute smoothness of it. I can run my hand up and down, but I can't feel where the top ends. I feel that it must go on up forever. So the only thing I can explain this feeling as is that it's like a tall, smooth column made of glass . . . There is this funny sort of feeling of being able to reach my hand into this area. It's very, very pleasant. (Cytowic, 2002, p. 33).

MW experienced simple tactile qualities that he sensed as identical over time if given the same stimulus. Querying other synesthetes about their experience together with test-retest situations subsequently confirmed this stability, thus leading to a general statement that became one of my five diagnostic criteria for the trait, namely that 'synesthetic percepts are generic and consistent', meaning that what is experienced is never complex or pictorial, but elementary — blobs, lattices, cold, rough, sour, zigzag or geometrically simple — and essentially invariable (Cytowic, 2002, pp. 67–9). Teasing out this phenomenal feature therefore contributed to synesthesia's nosology (the classification of diseases), and as an axiom it has held up over time.

Getting past (a)-type biases to clarify 'what it is like' led me to concurrently rethink the mechanism of synesthesia from a wholesale transgression of sense

modalities into one another's territories (as it had always been described) to the coupling of very specific attributes (what I originally called 'sensory fragments' but later recognized as the philosopher's 'qualia') in anomalous combinations: the whoosh of a furnace ignition, for example, always being *heard* as 'a stack of pink, green, and steely blue lines' about a foot in front of the subject 'moving up to the right as they fade' over the span of three seconds. Note the specificity of the description and the multiple sensations within it.

The issue of qualia led me to the binding problem (by one definition, synesthesia is anomalous binding) and similar contradictions that synesthesia presented to established concepts, among them modularity and functionalism. In fact, the history of synesthesia has been an increasingly unavoidable contradiction of traditional ideas about how the nervous system is organized. The typical response of contemporary scientists was denial of synesthesia as a real perceptual phenomena — an extreme example of (b)-type bias. Cytowic (2002a) summarizes how the inadequacy of traditional hierarchical models of brain organization led me to alternative models, such as selectively distributed systems as well as how the phenomenon may help elucidate the neural basis of metaphor, whereas Gray *et al.* (2002) discuss synesthesia's challenge to functionalism.

Had I accepted MW's descriptions at face value or as metaphoric, rather than probing to tease out the sensory qualities he was struggling to convey, synesthesia might have remained relegated to a mere curiosity rather than being recognized today as an anomaly with important implications for concepts of how the brain is organized (Cytowic, 2002; Smilek and Dixon, 2002; Ramachandran and Hubbard, 2001). Rather than second guessing experiential reports, the focus of those who are currently exploring the phenomenon has been on trying to understand synesthetes' reports and seek behavioural correlates of the subjective claims. This requires interplay between first person and third person accounts. A dialogue or structured questioning between clinician and subject constitutes a second person relation between shared knowledge about experience. This kind of feedback sometimes leads to further self-observations from a synesthete and the subsequent deployment of additional third person tests.

Theoretical constructs ((b)-type biases), though evidently derived from third-person observation, nonetheless have historically shortchanged first person accounts. My own awareness of the subject-examiner mismatch reaches back to ophthalmologic training in the 1970s. I was struck by how often patients spoke of 'seeing things' that they could describe in detail sufficiently nuanced so as to sound similar to other patients' 'things', yet the neuro-ophthalmologic exam revealed nothing amiss despite our using plenty of equipment. Faculty and trainees assumed that every optical symptom had a physical correlate that we could observe through our optics. This conceit never had to be taught explicitly because it was just taken for granted. Looking back, that confidence is breathtaking, but many (b)-type assumptions appear so 'obvious' that their bias is never questioned. Patients did not want to hear, 'I don't see anything wrong with your eyes', because it left their questions unanswered, and their persistent asking left us frustrated. We each had different expectations. Perhaps my sensitivity to

the mismatch between subject–observer expectations left me open to MW's experience instead of dismissing his synesthesia as impossible, as my peers automatically had (Cytowic, 2002, p. 64). (b)-type bias can therefore influence the very kind of data an investigator will take up as interesting.

Can we dismiss subjects' interpretations without dismissing either them or their experience, thereby leaving us empty handed of potentially valuable data? Yes, nondismissive disregard is a learnable skill so long as we acknowledge that knowledge has limits and we see conventional wisdom as porous to new ideas rather than fixed and impervious. It is another bias not to realize that human nature makes us emotionally attached to our ideas (as in the above denial of synesthesia being real because traditional models could not account for it). Training ourselves is the compliment to training introspectors. By letting experiential reports guide data analysis, for example, the pathophysiological loci of the above-mentioned 'spontaneous visual phenomena' (as those visual 'things' are now called) was eventually determined (Lepore, 1990), just as ffytche and Howard discovered new explanations for visual hallucinations by this same approach a decade later (see above).

Reporting on my experience in synesthesia, I cautioned: 'Though synesthetes are often dismissed as being poetic, it is *we* who must be cautious about unjustifiably interpreting their comments' (Cytowic, 1997, p. 24). Whereas training introspectors is one way to address their bias (Varella, 1996; Varella and Shear, 1999), examiners still need to become more attentive to their theory-laden assumptions that influence diverse sources of data. One theoretical construct that illustrates how an assumption can gradually change is the 'resting state' observed during metabolic brain imaging.

For twenty years it was accepted fact that subjects' brains were essentially blank when not performing the experimental stimulus task. That is, the metabolic landscapes obtained while quietly lying in the apparatus were taken as 'zero'. These resting states were even regarded as comparable between subjects. In the earliest days of measuring regional cerebral blood flow (rCBF), techniques based on following the dilution of a tracer in the bloodstream (the standard Fick principle of physiology) required experimenters to hold a 12-gauge needle and enormous syringe in the carotid artery while repeatedly shouting to the subject, 'Relax!'. So far as I can determine from those who actually conducted such measurements, no one considered if or how arousal in a terrified subject might influence the results.

When rCBF assessment changed over to the xenon<sup>-133</sup> inhalation technique in the early 1970s, it took about ten years to recognize that the elevated frontal activations routinely seen in resting landscapes were artifacts of the testing environment due to novelty and emotional arousal. Only recently have we even begun to take into account what subjects actually do as they lie in today's sophisticated scanners supposedly occupied by the experiment. We now admit that subjects grind their teeth, daydream, plan, are emotionally aroused, or engaged in a range of shifting cognitive states that have nothing to do with the experiment, may be contrary to it, and unquestionably affect the images obtained. Subjective mental

states produce metabolic landscapes that need factoring rather than being dismissed as irrelevant to the experimental activation. We are beginning to use scripts in an attempt to isolate a given cognitive process, but this approach is not yet completely adequate.

The blank slate supposition also extended to animal work. At a World Congress of Neurology, a presenter discussing unit recording of parietal neurons in monkeys reaching for targets actually said: 'The monkey is *told to sit there and not do anything*'. Perhaps more remarkable is that few people caught the implicit assumption that the monkey's brain is wholly devoted to the experimental task *and nothing else*. The phrase 'behaving monkey' continues to appear in many papers, but can we even define 'behaving monkey' in a satisfactory way? As Uttal (2001) argued in his provocative book, many purported cognitive functions that are not circularly defined cannot be adequately defined at all. Leopold and Logothetis (this issue) address subjective states in monkeys.

Objectivists seem always to pine for Nagel's 'view from nowhere', that Platonic somewhere that is detached from everyone's bias (Nagel, 1986). A favourite perceptual example is the colour red supposedly being determined by viewer-independent red wavelengths of a specific measure. Though trying to separate the seer from what is seen in this manner bypasses the question of whether we can know if viewers of light of such-and-such a wavelength have the same visual experience even though they may all name it 'red' (the privacy of experience argument), this example is mistaken for a more fundamental reason. First, the basis of colour vision is widely misunderstood. It is emphatically not wavelength-based. If it were, then the colour of objects would continually change because the wavelength composition of reflected light varies with the illumination. Yet a banana looks yellow whether we view it indoors in fluorescent or incandescent light, or outdoors in bright sun or shade. *Colour constancy* is a fundamental issue in vision science, a problem that Edwin Land solved thirty years ago (Land, 1974; Cytowic, 2002, pp. 254–60, 327–36) yet an elucidation not widely known.

Colour is a property that brains assign to surfaces, not a property inherent in light and not a property inherent in objects. Area V4/V8 calculates a ratio (which remains constant as illumination varies) between wavelength composition reflected from a point and that from all surrounding points, and assigns an object a constant colour. (This means that physical mind-independent objects do not literally have dispositional properties causally corresponding to experiential states representing colour. Further, synesthetes who experience colour in response to hearing spoken words activate their left V4/V8 (Nunn, *et al.*, 2002).) As fascinating and pregnant as colour perception is for discussion, my point is that it is always a particular someone's brain performing this ratio-taking, a brain that has a unique synaptic matrix, hence a unique knowledge base determined partly by genetically determined axonal connections and more so by experience-induced modifications in their synaptic strengths. The deep intermediate processing (i.e. everything between input and output) in human brains allows identical stimuli to trigger disparate responses depending on context, experience, consequences and

so forth (Mesulam, 2000). So the privacy argument may matter in ways its proponents never realized or intended.

An attempt to eliminate the subjective role of a human observer in gathering empirical data is very much a twentieth-century phenomenon, first in physics then in psychology. In the nineteenth century, psychophysicists such as Gustav Fechner tried to formulate laws regarding sensation and perception, taking as a given that mental states exist. Even today's craze for functional brain imaging, which is supposed to be anatomically objective, starts with what one wants to verify objectively: the subject's state of mind. Because the nineteenth century Brodmann maps that people use to point to this and that spot in the brain differ in absolute coordinates from person to person, fMRI interpretations are probably unreliable when subjects are averaged in a study, yet this happens routinely. So 'red wavelength' propositions — the idea of subject-independent measurement — unfortunately do lead to (b)-type errors even in the case of our latest sophisticated technology. 'Who' gets scanned matters. Another problem with functional imaging that has received scant attention is that whereas it has impressively advanced understanding of the functional landscape in human brains, it has contributed nothing to our understanding of connectivity among various structures. We currently understand very little in the human compared to our knowledge of connectivity in other species (Mesulam, 2000). If the Brodmann maps vary from person to person, connectivity probably does as well. These issues need to be determined rather than assumed one way or the other.

The experimenter may observe, but the subject has access to experience. I have accounted only for biases (a) and (b) from subject and experimenter respectively, and my conclusion stated at the beginning — that introspective reports, not accepted literally but properly interpreted and revised by investigators as necessary, are legitimate sources of data — follows only if these are the sole sources of bias. Critics of introspection typically invoke the privacy and consequent unverifiability of reports as their main reason for rejecting them as legitimate sources of scientific evidence. It is not my intent to cover all methodological issues, but this objection strikes someone with a clinician's point of view as more academic than practical.

The clinician's milieu perhaps reminds us daily that reports are our bread and butter, sometimes the only thing we have to go on. They are not so absolutely unverifiable and private that we never have some shared point of reference. In the clinician's experience, subject and object are commingled, separating fully only in the categories that human minds invent. Concepts and models — whether of brains, cognition or parallel universes — are by definition reductions of reality produced by human minds trying to grasp relations among too many variables. A common mistake is reifying the model and confusing it with reality. When having to deal with reality — whether in the form of sick patients or experimental subjects — the clinician by temperament is likely to forgo conceptual arguments and suggest that science is better served by striving for clear expression and understanding what both subject and experimenter 'mean by' their reports, their instructions and whatever else gets said. There is no reason to suppose that

investigator bias is any more problematic for introspective reports than it is for other kinds of data, such as cerebral metabolism.

Should we trust the subject, as the editors of this volume ask? I think so, but not literally. The assumed biases of both subject and investigator can and should be turned into a useful methodology. We should listen with the mind of someone on whom nothing is lost, render the data artfully, and try not to judge too quickly.

### References

- Cytowic, R.E. (1997), 'Synesthesia: phenomenology and neuropsychology', in *Synaesthesia: Classic and Contemporary Readings*, ed. S. Baron-Cohen and J.E. Harrison (London: Blackwell), pp. 17–42.
- Cytowic, R.E. (2002), *Synaesthesia: A Union of the Senses*, 2nd edn. (Cambridge, MA: MIT Press).
- Cytowic, R.E. (2002a), 'Touching tastes, seeing smells — and shaking up brain science', *Cerebrum*, **4** (3), pp. 7–26.
- Cytowic, R.E. (2003/1993), *The Man Who Tasted Shapes*, revised edn. (Cambridge, MA: MIT Press).
- ffytche, D.H. and Howard, R.J. (1999), 'The perceptual consequences of visual loss: "positive" pathologies of vision', *Brain*, **122** (pt. 7), pp. 1247–60.
- Gray, J.A., et al. (2002), 'Implications of synaesthesia for functionalism: theory and experiments', *Journal of Consciousness Studies*, **9** (12), pp. 5–31.
- Jack, A.I. and Shallice, T. (2001), 'Introspective physicalism as an approach to the science of consciousness', *Cognition*, **79** (1–2), pp. 161–96.
- Horowitz, M.J. (1964), 'The imagery of visual hallucinations', *Journal of Nervous and Mental Diseases*, **138**, pp. 513–23.
- Horowitz, M.J. (1975), 'Hallucinations: an information processing approach', in *Hallucinations: Behavior, Experience and Theory*, ed. R.K. Siegel and L.J. West (New York: Wiley), pp. 163–96.
- Klüver, H. (1942), 'Mechanisms of hallucinations', in *Studies in Personality*, ed. Q. McNemar and M.A. Merrill (New York: McGraw Hill).
- Klüver, H. (1966), *Mescal and Mechanisms of Hallucination* (Chicago: University of Chicago Press).
- Land, E. (1974), 'The retinex theory of color vision', *Proceedings of the Royal Institute of Great Britain*, **47**, pp. 23–58.
- Lepore, F.E. (1990), 'Spontaneous visual phenomena with visual loss: 104 patients with lesions of retinal and neural afferent pathways', *Neurology*, **40**, pp. 444–7.
- Mesulam, M.-M. (2002), 'Behavioral neuroanatomy', in *Principles of Behavioral and Cognitive Neurology*, 2nd edn., ed. M.-M. Mesulam (New York: Oxford), pp. 1–120.
- Nagel, T. (1986), *The View From Nowhere* (New York: Oxford University Press).
- Nunn, J.A., et al. (2002), 'Functional magnetic resonance imaging of synesthesia: activation of V4/V8 by spoken words', *Nature Neuroscience*, **5** (4), pp. 571–5.
- Ramachandran, V.S. and Hubbard, E.M. (2001), 'Synesthesia: a window into perception, thought, and language', *Journal of Consciousness Studies*, **8** (12), pp. 3–34.
- Siegel, R.K. and Jarvik, M.E. (1975), 'Drug induced hallucinations in animals and man', in *Hallucinations: Behavior, Experience and Theory*, ed. R.K. Siegel and L.J. West (New York: Wiley), pp. 81–62.
- Smilek, D. and Dixon, M.J. (2002), 'Towards a synergistic understanding of synaesthesia: combining current experimental findings with synaesthetes' subjective descriptions', *Psyche*, **8** (01), [<http://psyche.cs.monash.edu.au/v8/psyche-8-01-smilek.html>]
- Uttal, W.R. (2001), *The New Phrenology: The Limits of Localizing Cognitive Processes in the Brain* (Cambridge, MA: MIT Press).
- Varela, F.J. (1996), 'Neurophenomenology', *Journal of Consciousness Studies*, **3** (4), pp. 330–49.
- Varela, F.J. and Shear, J. (ed. 1999), *The View From Within: First-Person Approaches to the Study of Consciousness* (Exeter: Imprint Academic).
- Victor, M., Ropper, A.H. and Adams, R.D. (2000), *Principles of Neurology*, 7th edn. (New York: McGraw Hill), or see any standard neurology textbook.
- Zaidel, E., Iacoboni, M., Berman, S. and Bogen, J.E. (2003), 'The callosal syndromes', *Clinical Neuropsychology*, 4th edn., ed. K.M. Heilman and E. Valenstein (New York: Oxford University Press), ch. 14.