

Neural Synchrony, Attention, and Unified Consciousness

Abstract:

A central issue in philosophy and neuroscience is the problem of unified visual consciousness. This problem has arisen because we now know that an object's features (e.g., its color, texture, shape, etc.) are represented in different areas of the visual cortex (Felleman & Van Essen, 1991). At the same time, there are very few, if any, direct neural connections between specific visual areas, such as those that represent color and motion (Zeki, 2003). So how do we experience unified object representations? Some neuroscientists propose that neural synchrony is the mechanism that binds an object's features into a representational unity (e.g., see Crick and Koch, 1990, 1994; Singer, 1996; von der Malsburg, 1996, 1999; Roelfsema, 1998; Engel, 2003). However, I argue on both empirical and philosophical grounds that neural synchrony fails to explain the unity of visual consciousness. In the concluding section, I discuss some recent evidence for an approach to consciousness that takes seriously both cognitive neuroscientific data *and* phenomenal data.

Key Words: attention, binding, consciousness, feature ambiguity, neural synchrony, unity

1. Introduction

A central issue in philosophy, cognitive psychology, and related neurosciences is the problem of unified visual consciousness. This problem has arisen because we now know that an object's features (e.g., its color, texture, shape, etc.) are represented in different areas of the visual cortex (Felleman & Van Essen, 1991). At the same time, there are very few, if any, direct neural connections between specific visual areas, such as those that represent color and motion (Zeki, 2003).¹ So how do unified object representations arise in visual consciousness? Some prominent neuroscientists propose that neural synchrony is the mechanism that binds an object's features into a representational unity (e.g., see Crick and Koch, 1990, 1994; Singer, 1996; von der Malsburg, 1996, 1999; Roelfsema, 1998; Engel, 2003). On this account, an object's features are said to be bound together (and thus appear to consciousness as a unity) because the underlying neurons of those features fire in a correlated fashion. As von der Malsburg observes: 'If,

during a time interval, the signals on a set of neurons are found to be significantly correlated, the set is interpreted as being bound during that interval' (1996, p. 137).

However, I argue on both empirical and philosophical grounds that neural synchrony fails to explain the unity of visual consciousness. In the concluding section, I discuss some recent evidence for an approach to consciousness that takes seriously both cognitive neuroscientific data *and* phenomenal data.

The paper is organized as follows. In section two, there is a brief discussion of the visual system and the object feature binding problem. In section three, I utilize the work of Crick and Koch (1990, 1994) to elaborate the neural synchrony approach to binding. In section four, I discuss the limitations of neural synchrony as a binding mechanism. I conclude with a discussion of recent empirical evidence that implicitly favors a wider, non-reductive approach to consciousness that takes phenomenal/first-person data seriously.

2. Visual Areas and the Object Feature Binding Problem

Contemporary neurobiology indicates that object representations of visual consciousness rely upon the specialized activity of neuronal assemblies distributed throughout the visual hierarchy (Felleman & Van Essen, 1991; Zeki, 1994). As one gazes upon a maple tree, for instance, the maple's retinotopic pattern is rapidly laid out on the surface area of V1. Neurons in areas V1 and V2 respond to primitive visual features, such as edges and contours, in a variety of orientations (Grosz, Shapely, & Hawken, 1993). Once a relatively stable figure of the maple has been established and segregated from its background in these lower areas, attentional mechanisms send information about the

maple downstream for further processing in specialized areas of the ventral and dorsal pathways, also known as the 'what' and 'where' systems (Mishkin & Ungerleider, 1982). The ventral system's neuronal pathway projects from the occipital lobe to the inferior temporal lobe (IT); it plays the functional role of processing information about an object's shape, color, and texture. The dorsal system's neuronal pathway projects from the occipital lobe to the parietal lobe (7a); it performs the functional role of processing information about an object's location and size. The middle temporal and middle superior temporal lobes (MT and MST) process movements associated with individual objects or collections of objects against stable or relatively stable background conditions. Thus, visually perceiving the maple's features on a blustery day would at least involve the relevant firing patterns of neuronal subassemblies in areas V1-4, IT, 7a, MT and MST.²

Neuropsychological evidence confirms this distributed view of visual feature representations as well. For example, damage in V4, an area of the ventral system, produces achromatopsia, i.e., color blindness (Meadows, 1974); damage in IT produces associative agnosia, i.e., the inability to identify shapes (Efron, 1968); and damage in MT produces akinetopsia, i.e., motion blindness (Zhil, von Cramon, & Mai, 1983; Zeki, 1992). Moreover, various studies implicate the dorsal system as underlying our spatial attention capacity (Posner & Peterson, 1990; see also Corbetta, et al., 1995), spatial representational capacity (Colby, 1991; Robertson, 1999), and the ability to discriminate within and between objects (Elgy, Driver, & Rafal, 1994; Robertson, 1999). Thus, what we currently know about the visual system is that it relies upon several specialized subsystems distributed throughout the visual hierarchy. Although this functional

specialization has enabled the brain to process information efficiently, it also implies several vision-related binding problems.

One such problem is the object feature binding problem. This problem has arisen because we now know that an object's features (e.g., its color, texture, shape, motion, etc.) are represented in different areas of the visual cortex. At the same time, there are very few, if any, direct neural connections between specific visual areas, such as those that represent color and motion (Zeki, 2003). Furthermore, it is unlikely that there could be a convergence zone that underlies what we visually experience because of the combinatorial capacity problem: the possible combinations of features that are represented throughout our lives would seem to far exceed the neural machinery with which we are equipped (e.g., see Crick & Koch, 1990, 1994; Singer, 1996; von der Malsburg, 1996, 1999).³ So, how do we explain the conscious unity of an object's features, if the distinctive features of an object are represented by neurons distributed throughout the visual hierarchy?⁴ Let us now consider the neural synchrony approach to unified visual consciousness.

3. Neural Synchrony Approach to Unified Visual Awareness

Some neuroscientists allege that the conscious unity of an object's features is explained in terms of synchronized neuronal activity in the relevant parts of the visual cortex (Crick and Koch, 1990; see also von der Malsburg, 1996; Singer, 1996). Neuronal synchronies are phase-locked and exhibit oscillations in the 40 to 75 Hz range (Crick & Koch, 1990, p. 271).⁵ Neurons are phase-locked in the sense that if some neurons within a given population slightly speed up or slow down, the remaining neurons within that population

will also speed up or slow down to maintain the same speed with the others. In effect, the representational unity of a visual object's features allegedly corresponds to a subset of neuronal events in the visual cortex. As Crick and Koch note, '[a]t any moment consciousness corresponds to a particular type of activity in a transient set of neurons that are a subset of a much larger set of potential candidates' (1990, p. 266). What Crick and Koch mean by 'corresponds to' is that unified visual consciousness is wholly explainable in terms of, and is therefore reducible to, neural activity. Incidentally, their commitment to reductionism as a method of explanation is not limited to the area of visual consciousness. For example, Crick claims that 'each of us *is* the behavior of a vast, interacting set of neurons' (1994, p. 203; my italics). Metaphysically speaking, their view is consistent with type-type identity theory, which says that the mental is reducible to the neural (cf. Place, 1956; Smart, 1959; cf. Armstrong, 2000).

To substantiate the claim that neural synchrony is the mechanism that explains the unitary character of an object's features at the level of consciousness, Crick and Koch (1990, 1994, 2003) and other neuroscientists appeal to experiments conducted on cat primary visual cortex (see also von der Malsburg, 1996). This research indicates that neurons exhibit synchrony in response to a moving bar. The experimental results also reveal that correlated firing among a group of neurons is much stronger in response to a single moving bar as opposed to a pair of moving bars (1990, p. 271). I might add that there are some implicit advantages to the neural synchrony view. First, it is compatible with the functional specialization of our visual anatomy and thus does not posit a convergence zone to account for unified object representations. Second, neural synchrony has also been observed between neurons whose locale is in entirely different

hemispheres of the brain (Engel, et al., 1991). Thus, the neural synchrony view is consistent with an account of large-scale binding, not simply local binding (see Varela & Thompson, 2003).

Advocates of the neural synchrony approach recognize, however, that neuronal synchrony has important challenges that it must overcome if it is to successfully account for the unity of an object's features at the level of consciousness. One such challenge that it must address is the feature ambiguity problem. The feature ambiguity problem can be articulated as follows. Imagine that one is viewing two objects at the same time, for instance, a red circle and a blue square in a child's toy box. Presumably the stimulus features of these two objects will activate separate neuronal assemblies simultaneously. However, this implies a problem for the theory of neuronal synchrony. It is unclear how synchronized neuronal activity could know how to assign the appropriate shape with the appropriate color when competing objects are viewed simultaneously. Along with Crick, one might ask, 'How could the brain know which color to put with which shape? In other words, if awareness corresponded merely to rapid (or sustained) firing, the brain might easily confuse the attributes of different objects' (1994, p. 210).

Crick and Koch (1990, 1994) invoke the functional role of attention to solve this problem.⁶ It is proposed that attention operates in a serial or sequential fashion, marking the features of individual objects at distinct times, by causing individual neuronal assemblies to fire in synchrony. These time-indexed neuronal assembly firings are supposedly the correlates of a particular visual object and its respective features. On this view, 'what matters is not just the average rate of firing of a neuron but the exact moments at which each neuron fires' (1994, p. 211). The distinct subpopulations of

neurons that correlate with the representational features of the red circle (namely, its shape and color) fire vigorously at one time (say, t_1) for approximately 100 milliseconds; and the distinct subpopulations of neurons that correlate with the representational features of the blue square fire vigorously at another time (say, t_2) for approximately 100 milliseconds. Hence, the neurons that underlie the representational features of the red circle and the blue square never fire at the same temporal interval. By attending to the representational features of each object at separate times, Crick and Koch infer that conscious awareness depends on and is the result of this disambiguating process: ‘This form of transient binding probably depends on a serial attentional mechanism, sometimes called the spotlight of attention’ (1990, p. 269).

As it happens, in a recent paper Crick and Koch (2003) summarize their position concerning the role attention plays in binding an object’s features together when competing objects are simultaneously present before one’s visual field:

Several objects/events can be handled simultaneously—more than one object/event can be attended to at the same time—*if there is no significant overlap in the cortical neural network*. . . . If there is such an overlap, then (top-down) attention is needed to select one of them by biasing the competition among them. This [attention based] approach largely solves the classical binding problem, which was mainly concerned with how two different objects/events could be ‘bound’ simultaneously. On this view, the ‘binding’ of the features of a single object/event is simply the membership in a particular coalition (2003, p. 123; author’s own italics).

Even on this recently proposed cellular coalition (i.e., assembly) model, visual attention is still required to select (or ‘mark’) a group of features as belonging to its respective object when two or more objects create significant overlap in the cortical hierarchy. Much like their previous theorizing, they propose that visual attention selects a coalition of cells when competition among neural groups occurs, and that it is in virtue of this

attentionally selected coalition that the perceived unity of an object's features is produced: 'In general, at any moment the winning coalition is somewhat sustained, and embodies what we are conscious of' (Crick & Koch, 2003, p. 121). In other words, Crick and Koch implicitly assume that if their neural account of consciousness passes the test of object feature ambiguity, then it should be able to explain object feature binding.

4. Limitations of Neural Synchrony: Empirical and Philosophical Considerations

Recent empirical evidence strongly suggests that neural synchrony is not enough to explain unified object representations in visual consciousness. First, neural synchrony probably cannot account for the enduring character of object representations in visual consciousness because of its fleeting nature. For example, O'Reilly and colleagues (2003) point out that once an object is removed from a subject's visual field and the neuronal assembly that subserves that object's features ceases firing, there is no indication of that assembly's neurons being bound and 'yet we can have enduring representations of bound features, so somehow this problem needs to be addressed' (2003, p. 171). What O'Reilly and colleagues have in mind is the experience of the continuous (or diachronic) character of unified percepts. Phenomenologically speaking, the conscious mind does not perceive the world's objects/events in discrete temporal intervals. Rather, the experiential unity of a scene's component objects/events bears the distinct phenomenal feature of continuity, the unbroken stream of consciousness that James (1890) and Dainton (2000) have so deftly characterized in their phenomenological analyses. The perceived unity of objects in visual space seems to be experienced along with a feeling of the continuity of time. There is, in a sense, a coalescence of space and time in the experience of object perception. There are different forms of temporal unity

associated with object unity in visual consciousness. These forms are expressible through an asymmetry claim: unity over time implies unity at a time, but unity at a time does not necessarily imply unity over time. For example, the visual subject can be aware of the distinct features of an object at separate times in the initial phase of the perceptual cycle (see Moutoussis & Zeki, 1997; Zeki & Bartels, 1998)⁷; but the visual subject can also be aware of the coalescence of those features as unified wholes over time. Now, since the experience of the representational unity of an object's features persists beyond the subpopulation of cells that fire in synchrony, the experience of such diachronic unity is irreducible to neural synchrony.

Second, temporal synchrony could play a functional role other than binding an object's attributes together. For example, it could be that neural synchrony fundamentally explains the flexibility of our cognitive architecture rather than its binding capacity. Neuronal flexibility underpins an account of how specialized neurons (e.g., those representing color, texture, motion, etc.) can be recruited to participate in distinct assemblies of neurons at successive times. Due to combinatorial capacity limitations, the brain has likely developed an efficient strategy to represent the individual features of objects when two or more objects share some of the same features in a given visual scene. For example, the very same neurons that represent the color red for one object (O1) could represent the color red for another object (O2). This is possible because the neurons that represent red for O1 could fire in synchrony at some distinct time (say, t_1) and then the very same neurons could fire in synchrony at a later time (say, t_2) in order to represent red for O2. The brain likely developed this flexible strategy (in the form of successive chains of synchronized neural assemblies) in order to overcome combinatorial

limitations. But providing an account of cognitive flexibility is not equivalent to providing an account of binding. Therefore, the very mechanism that underlies cognitive flexibility—namely, successive coalitions of neurons firing in synchrony—is not the same mechanism that explains binding.

Third, the experimental research has indicated that neural synchrony results primarily from ‘moving stimuli and is notoriously difficult to measure with stationary stimuli’ (Hummel & Biederman, 1992, p. 509). Thus it is difficult to see how neural synchrony could be an adequate binding mechanism, since some of the objects we perceive in both natural and artificial environments (e.g., mountains, boulders, houses, etc) are *stationary*.

Fourth, the experimental data used to confirm the temporal correlation hypothesis were obtained from both anaesthetized and awake animals, namely, cats and monkeys (see Gray & Singer, 1989; Engel et al., 1990; Shadlen & Moshvov, 1999). In fact, Crick mentions that much of the experimental work is performed on ‘anesthetized animals who are not conscious’ (1994, p. 15). That is to say, synchronized neuronal firings in the 40 Hz range occur in response to visually presented objects in the brains of both conscious and unconscious animals. We, therefore, cannot claim that neural synchrony is the distinctive neural correlate of unified visual consciousness.⁸ Moreover, neuronal synchrony correlates strongly with preattentive awareness, i.e., unconscious activity in V1 (see Gray, 1999; Luck and Beach, 1998). It seems to me that a critical question naturally follows from this data: *How can synchronized neuronal firings in the 40 to 70 Hz range be the direct neural correlates of consciousness if such ‘specialized’ neuronal activity correlates with preattentive awareness (i.e., unconscious activity) as well?* What

this experimental data intimates is that consciousness (C) and synchronized neuronal firings (SNFs) cannot be identical, since SNFs also correlate directly with unconsciousness (UC). In other words, identity is a symmetrical relation. So, if C is fully explainable in terms of SNFs, then C and SNFs will share all and only the same properties. Since C cannot share the property of UC, C and SNFs cannot be identical. Or, if we think of the above relation as holding a transitive relation, then we end up with absurdity: if C and SNFs are neurobiologically equivalent and SNFs and UC are neurobiologically equivalent, then C and UC are neurobiologically equivalent. Evidently, the temporal correlation hypothesis explains too much; the same neuronal mechanism (SNFs) allegedly explains both stimulus-related binding correlated with consciousness and stimulus-related binding correlated with unconsciousness.

Fifth, we discovered that Crick and Koch invoke the notion of an attentional mechanism in order to solve the feature ambiguity problem. This seems like a plausible inference to draw. Nonetheless, Crick and Koch implicitly assume that if their neural theory possesses the explanatory tools requisite to explain object feature disambiguity (OFD), then it should also provide a viable account of object feature unity (OFU). They appeal to attentional mechanisms in order to explain object feature disambiguation, and then infer that membership in a neural coalition, which is selected by attentional mechanisms, is enough to explain object feature unity. As they put it, 'If there is such an overlap [in the cortical neural network], then (top-down) attention is needed to select one of them [i.e., a coalition of cells] by biasing the competition among them. . . . On this view, the 'binding' of the features of a single object/event is *simply* the membership in a particular coalition' (2003, p. 123; my italics). This inference, however, is not warranted.

Even if an explanation of feature disambiguation could show how an object's representational features are selected ('marked' or 'labeled') when competition occurs, this would not automatically guarantee an account of how such features, once selected, are bound together to form a unified object of consciousness. Which is to say: feature disambiguation does not entail feature binding. To further clarify and motivate this point, an illustration may be helpful. Imagine a grid of lights, such that each light is labeled to represent an individual feature of an object. To simplify matters, the total number of objects that can be represented by the grid is five. The labels that represent an object's features at any given moment are distributed throughout the grid-like system; the label for 'red' is above the label for 'circle' and the label for 'motion' is above and to the right of the label for 'red', and so forth. Since an object's features are represented (via labels) in a distributed fashion, we need a mechanism to disambiguate an object's respective group of features when competing object features light up at the same time. One way to know which features belong to their respective objects is by means of a selecting (or 'marking') mechanism. In the case of the grid, an electrical switch causes ('selects') the correct coalition of labeled lights to turn on (or 'fire' more brilliantly than competing sets of lights) at a distinct time. Notice, however, that *simply because we have a mechanism that selectively activates the correct set of features distributed throughout the grid-like system, we have yet to explain how those features, once selectively activated, are bound together to form a single unitary object as opposed to an unconnected group of features.* Thus, binding the *correct* set of features together (when competition occurs in the cortical neural network) entails disambiguation, but disambiguation per se does not entail binding. For why not suppose that if the selectively activated subpopulations of neurons

that underlie the representational states of red and of circle were identical to (or merely strongly correlated with) such subpopulations, that consciousness would actually look like a theatre of separate features minus object unity?

To provide further weight to this criticism, evidence from neuropsychology strongly suggests that attentional mechanisms are not necessary for binding, though they seem to play a central role in feature disambiguation. For example, a variety of experiments have been performed on normal subjects indicating that selecting the correct set of representational features when competition occurs in the cortical hierarchy depends upon the focus of attention (Treisman, 1996, 2003). In one experimental set-up, subjects were briefly shown two colored letters at the same time: a green T and a red O. The experimenters found that when the focus of attention was prevented subjects would experience illusory conjunctions: the subjects reported seeing ‘a red T when a green T and a red O’ were presented at the same time (Treisman, 2003, p. 99; see also Treisman & Schmidt, 1982). Consequently, red, rather than green, was bound to T and the result was an experience of illusory conjunctions. What is interesting is that these normal subjects have still performed the function of binding, albeit of an illusory conjunction sort. Therefore, focal attention is *not* necessary for binding. If binding can take place independent of the disambiguating role of attention, then an account disambiguation (via attention) does not guarantee an account of binding. Hence, OFD does not necessarily entail OFU. Nevertheless, the functional role played by attention implies a *normative* dimension to the disambiguating process: the perceptual subject’s ability to select the *correct* set of features when competition occurs in the cortical hierarchy is directly tied to the subject’s attentional capacities.

Finally, even if binding were to occur because of some yet to be discovered neural mechanism, one could argue that an explanatory gap would still remain between neural binding and conscious experience. In this respect, Chalmers argues that Crick and Koch do not explain why binding itself should go together with conscious experience: if we are not sure why binding brings about experience, then “telling us a story about oscillations cannot help us Crick and Koch’s theory gains its purchase by assuming a connection between binding and experience, and so can do nothing to explain that link” (2000, p. 387). In other words, suppose that we constructed a robot that could bind together different pieces of visual data by way of its silicon and metal bits, or quasi “neural nets.” Binding, in the case of the robot, involves displaying individual features of visual information on its internal monitor in a unified way. But despite the robot’s capacity to simulate human behavior and to bind different features of visual data together on its internal monitor, what the robot lacks is an *experience* of the data bound together. Even though experience entails binding, binding does not necessarily imply experience (see my 2002; cf. Treisman, 2003). Therefore, there is no necessary connection between binding and experience. The concepts invoked to describe binding by neural activity do not logically entail experience and thus a gap arises between what neural mechanisms can explain and the nature experience.

Crick and Koch could reply that they are merely attempting to explain *how* the features of an object appear to consciousness in a unitary way. Thus addressing deep metaphysical issues (such as *why* experience arises from neural activity) falls outside the scope of their more modest aims. This type of rejoinder, however, could imply that providing an explanation of the *unity of experience* is not the same as providing an

explanation of the *experience of unity*. Thus, experience itself may *not* be vulnerable to neural explanations after all. This latter claim, if true, would clearly fly in the face of Crick and Koch's commitment to neural reductionism. So perhaps it is in this sense that Chalmers's criticism is warranted.

5. Conclusion

Because the relation between cognitive neuroscientific concepts and the concepts of experience is asymmetric, it seems that any balanced theory of consciousness (which is not swayed by the winds of neural reductionism or eliminativism) will have to accommodate both personal (i.e., the phenomenological/first-person experience of unity) and subpersonal (i.e., neurobiological/third-person) aspects. In other words, if consciousness is not merely a theoretical posit (to be explained by a matured neuroscience in the future), but counts as a genuine explanandum, then we will have to adopt a wider methodological approach that takes seriously both cognitive scientific data *and* phenomenal experiential data as 'reciprocal constraints' towards understanding the nature of consciousness (see Varela & Thompson, 2003; p.267; see also Chalmers, 1996, 2000; Hurley, 1998; and my 2002). What role could phenomenological/first-person data play in a theory of consciousness? We might consider, for example, how the person's interpretations of ambiguous figures supply implicit evidence for this wider methodological approach. A well-known ambiguous figure is the Necker cube. While viewing a Necker cube, many persons consistently experience the perception of the cube

flipping back and forth between alternative orientations. To elaborate how this occurs, I will utilize the experimental research of Kelso and colleagues (1995; see also Varela and Thompson, 2003). Kelso and colleagues (1995) devised an experiment in which persons were asked to view a Necker cube in eight different spatial orientations. These orientations were presented randomly to the persons involved in the experiment. They were instructed to push a button each time they detected a change in the cube's spatial orientation. For the most part, there was no consistent pattern exhibited by the participants' response to observed switching times for each spatial orientation. However, as the orientation of the cube approached that of a square (oriented at 80 degrees), the orientation was perceived for a longer period of time without switching. In light of these results, one could infer that different interpretations of the Necker cube initiated by the person will influence slightly the neuronal bias that underlies the cube reversal (see Varela & Thompson, 2003, p. 277-278). This suggests that ambiguous figures are experienced differently depending on how they are interpreted by the person. Hence, higher-level interpretations carried out by the person act as top-down influences on the lower-level neural assemblies responsible for generating differing orientations. The person, in turn, experiences the result of these higher-level influences on lower-level processes.

Acknowledgements: I would like to thank Jesse Prinz, David Chalmers, and Kenneth Williford for helpful comments on earlier drafts. Parts of this manuscript were presented at *Toward a Science of Consciousness*, the sixth biennial Tucson conference on consciousness, sponsored by center for consciousness studies at the University of Arizona (April 2004); and *Neurophilosophy: The State of the Art*, held at the California Institute of Technology (June 2005). I thank all of the participants at these conferences for collegial attitudes and useful suggestions.

References

- Bayne, T., & Chalmers, D. (2003). What is the unity of consciousness? In A. Cleeremans (Ed.), *The unity of consciousness: Binding, integration, and dissociation* (pp. 23-58). Oxford: Oxford University Press.
- Bergen, J. R., & Adelson, E. H. (1988). Early vision and texture perception. *Nature*, 333, 363-364.
- Bravo, M., & Blake, R. (1990). Preattentive vision and perceptual groups. *Perception*, 19, 515-522.
- Braun, J. (1993). Shape-from-shading is independent of visual attention and may be a 'texton.' *Spatial Vision*, 7, 311-322.
- Brook, A. (1994). *Kant and the mind*. Cambridge: Cambridge University Press.
- Chalmers, D. (1996). *The conscious mind: In search of a fundamental theory*. Oxford: Oxford University Press.
- Chalmers, D. (1998). The puzzle of conscious experience. In P. Van Inwagen & D. Zimmerman (Eds.), *Metaphysics: The big questions* (pp. 333-341). Oxford: Blackwell Publishers.
- Chalmers, D. (2000). Facing up to the problem of consciousness. In B. Cooney (Ed.), *The place of mind* (pp. 382-400). Belmont: Wadsworth.
- Colby, C. (1991). The neuroanatomy and neurophysiology of attention. *Journal of Child Neurology*, 6, 90-118.
- Corballis, M. (1988). Recognition of disoriented shapes. *Psychological Review*, 95, 115-123.
- Corbetta, M., Shulman, G., Miezin, F., & Petersen, S. (1995). Superior parietal cortex activation during spatial attention shifts and visual feature conjunction, *Science*, 270, 802-805.
- Crick, F., & Koch, C. (1990). Towards a neurobiological theory of consciousness. *Seminars in the Neurosciences*, 2, 270-281.
- Crick, F. (1994). *The astonishing hypothesis*. New York: Charles Scribner's Sons.
- Crick, F., & Koch, C. (2003). A framework for consciousness. *Neuroscience*, 6, 119-116.

- Dainton, B. (2000). *Stream of consciousness: Unity and continuity in conscious experience*. London: Routledge.
- Duncan, J. (1984). Selective attention and the organization of visual information. *Journal of Experimental Psychology*, *113*, 501-517.
- Eccles, J. (1994). *How the self controls the brain*. New York: Springer-Verlag.
- Elder, J., & Zucker, S. (1993). The effect of contour closure on the rapid discrimination of two-dimensional shapes, *Vision Research*, *33*, 981-991.
- Elgy, R. Driver, J., & Rafal, R. (1994). Shifting visual attention between objects and locations: Evidence from normal and parietal lesion subjects, *Journal of Experimental Psychology: General*, *123*, 161-177.
- Engel, A., Konig, P., Gray, C., & Singer, W. (1990). Stimulus-dependent neuronal oscillations in cat visual cortex. *European Journal of Neuroscience*, *2*, 588-606.
- Engel, A. (2003). Temporal binding and the neural correlates of consciousness. In A. Cleeremans (Ed.), *The unity of consciousness: Binding, integration, and dissociation* (pp. 132-152). Oxford: Oxford University Press.
- Felleman, D., & Van Essen, D. (1991). Distributed hierarchical processing in the primate visual cortex, *Cerebral Cortex*, *1*, 1-47.
- Gray, C., & Singer, W. (1989). Stimulus specific neuronal oscillations in orientation columns of cat visual cortex. *Proceedings of the National Academy of Sciences*, *86*, 1698-1702.
- Gray, C. (1999). The temporal correlation hypothesis of visual feature integration: Still alive and well, *Neuron*, *24*, 31-47.
- Grosovsky, D., Shapely, R., & Hawken, M. (1993). Macaque V1 neurons can signal illusory contours, *Nature*, *365*, 550-552.
- Holmes, G. (1918). Disturbances of visual orientation, *British Journal of Ophthalmology*, *2*, 449-468.
- Hubel, D., & Wiesel, T. (1968). Receptive fields and functional architecture of monkey striate cortex, *Journal of Psychology*, *166*, 106-154.
- Hummel, J., & Biederman, I. (1992). Dynamic binding in a neural network for shape recognition, *Psychological Review*, *99*, 480-517.

- Humphreys, G., Cincel, C., Wolfe, J., Olson, A., & Klemm, N. (2000). Fractionating the binding process: Neuropsychological evidence distinguishing binding of form from binding of surface features, *Vision Research*, 40, 1569-1596.
- Humphreys, G. (2003). Conscious visual representations built from multiple binding processes: Evidence from neuropsychology. In A. Cleeremans (Ed.), *The unity of consciousness: Binding, integration, and dissociation* (pp. 114-131). Oxford: Oxford University Press.
- James, W. (1890). *Principles of psychology*. New York: Holt.
- Julesz, B. (1975). Experiments in the visual perception of texture, *Scientific American*, 232, 34-43.
- Keele, S., Cohen, A., Ivry, R., Liotta, M., & Yee, P. (1988). Tests of a temporal theory of attentional binding, *Journal of Experimental Psychology*, 14, 444-452.
- Kosslyn, S., & Koenig, O. (1995). *Wet mind: The new cognitive neurosciences*. New York: The Free Press.
- Lane, R. (2000). Neural correlates of conscious emotional experience. In R. Lane & L. Nadel (Eds.), *Cognitive neuroscience of emotion* (pp. 345-370). New York: Oxford University Press.
- Levine, D., Maini, R., & Calvanio, R. (1988). Pure agraphia and Gerstmann's syndrome as a visuospatial-language dissociation: An experimental case study, *Brain and Language*, 35, 172-196.
- Luck, S., & Beach, N. (1998). Visual attention and the binding problem: A neurophysiological perspective. In R. Wright (Ed.), *Visual attention* (pp. 455-478). New York: Oxford University Press.
- Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information*. New York: W.H. Freeman.
- Meadows, J. (1974). Disturbed perception of colours associated with localised cerebral lesions, *Brain*, 97, 616-632.
- Mishkin, M., & Ungerleider, L. (1982). Object vision and spatial vision: Two cortical pathways, *Trends in Neurosciences*, 6, 414-417.
- Moutoussis, K., & Zeki, S. (1997). Functional segregation and temporal hierarchy of the visual perceptive systems. *Proceedings of the Royal*

- Society of London B.*, 264, 1407-1414.
- O'Reilly, R., Busby, R., & Soto, R. (2003). Three forms of binding and their neural substrates: Alternatives to temporal synchrony. In A. Cleeremans (Ed.), *The unity of consciousness: Binding, integration, and dissociation* (pp. 168-192). Oxford: Oxford University Press.
- Peterhans, E., & von der Heydt, R. (1989). Mechanisms of contour perception in monkey visual cortex, *Journal of Neuroscience*, 9, 1149-1763.
- Peterson, M. (2001). Object perception. In Goldstein (Ed.), *Blackwell handbook of perception* (pp. 169-203). London: Blackwell.
- Place, U. (1956). Is consciousness a brain process? *British Journal of Psychology*, 47, 44-50.
- Posner, M., & Petersen, S. (1990). The attention system in the human brain, *Annual Review of Neuroscience*, 13, 25-42.
- Rafal, R. (1996). Balint's syndrome. In T. Feinberg & M. Farah (Eds.), *Behavioral neurology and neuropsychology* (pp. 337-356). New York: McGraw-Hill.
- Ramachandran, V. (1988). Perception of shape from shading, *Nature*, 331, 163-166.
- Ricci R., Vaishnavi S., & Chatterjee, A. (1999). A deficit of intermediate vision: Experimental observations and theoretical implications. *Neurocase*, 5, 1-12.
- Rock, I., & Brosgole, L. (1964). Grouping based on phenomenal proximity, *Journal of Experimental Psychology*, 67, 531-538.
- Roelfsema, P. (1998). Solutions for the binding problem, *Zeitschrift fur Naturforschung*, 53, 691-715.
- Roland, P., & Gulyas, B. (1995). Visual memory, visual imagery, and visual recognition of large field patterns by the human brain: Functional anatomy by positron emission tomography', *Cerebral Cortex*, 5, 79-93.
- Sekuler, A., & Palmer, S. (1992). Perception of partly occluded objects: A microgenetic analysis, *Journal of Experimental Psychology*, 121, 95-111.
- Shadlen, M., & Movshon, A. (1999). Synchrony unbound: A critical evaluation of the temporal binding hypothesis, *Neuron*, 24, 67-77.

- Singer, W. (1996). Neuronal synchronization: A solution to the binding problem? In R. Llinas & P. Churchland (Eds.), *The mind-brain continuum: Sensory processes* (pp. 101-130). Massachusetts: MIT Press.
- Smart, J. J. (1959). Sensations and brain processes. *Philosophical Review*, 68, 141-156.
- Sun, J., & Perona, P. (1997). Shading and stereo in early perception of shape and reflectance, *Perception*, 26, 519-529.
- Tanaka, K. (1993). Neuronal mechanisms of object recognition, *Science*, 262, 685-688.
- Treisman, A., & Schmidt, H. (1982). Illusory conjunctions in the perception of objects. *Cognitive Psychology*, 14, 107-141.
- Treisman, A. (1996). The binding problem. *Current Opinion in Neurobiology*, 6, 171-178.
- Treisman, A. (2003). Consciousness and perceptual binding. In A. Cleeremans (Ed.), *The unity of consciousness: Binding, integration, and dissociation* (pp. 95-113). Oxford: Oxford University Press.
- von der Heydt, R., Perterhans, E., & Baumgartner, G (1984). Illusory contours and cortical neuron responses, *Science*, 224, 1260-1262.
- von der Malsburg, C. (1996). The binding problem of neural networks. In R. Llinas & P. S. Churchland (Eds.), *The mind-brain continuum: Sensory processes* (pp. 131-146). Massachusetts: MIT Press.
- Zeki, S. (1980). The representation of colors in the cerebral cortex, *Nature*, 284, 412-418.
- Zeki, S. (1994). *A vision of the brain*. Oxford: Blackwell Science.
- Zeki, S., & Bartels, A. (1998). The asynchrony of consciousness. *Proceedings of the Royal Society of London B*, 265, 1583-1585.
- Zeki, S. (2003). The disunity of consciousness. *Trends in Cognitive Neuroscience*, 7, 214-218.
- Zihl, J. von Cramon, D., & Mai, N. (1983). Selective disturbance of movement vision after bilateral brain damage, *Brain*, 106, 313-340.

¹ Moreover, many binding theorists doubt that there could be a convergence zone that underlies what we visually experience because of the combinatorial problem: the possible combinations of features that are represented throughout our lives would seem to far exceed the neural machinery with which we are equipped (e.g., see Crick and Koch, 1990, 1994; Singer, 1996; von der Malsburg, 1996, 1999).

² It is important to note that I use the terms neuronal assembly, group, coalition, and population interchangeably. A neuronal assembly will be composed of individual subassemblies that correlate with specific features of any particular visual experience. Estimates vary. But the individual subassemblies that subservise or correlate with a single object-feature of experience will typically have a quantitative makeup ranging anywhere from 10 to the second power to 10 to the sixth power (Crick and Koch, 1998).

³ Presumably, this is why the brain evolved specialized subsystems. As Crick and Koch observe, 'One reason for these multiple areas is that to handle all activity in one single very large neural net, with everything connected to everything else, would make the brain both cumbersome and prohibitively large' (1990, p. 267).

⁴ In addition to the problem of object feature binding, there are other forms of binding that require an explanation as well, such as binding the parts of an object's shape into an integrated whole. Shape is an essential feature of an object whose parts exhibit determinate intra-spatial relations. The representation of a dog, for instance, involves establishing the spatial relations internal to a dog's body parts: the ears are 'on top of' the head; the legs are 'connected to' the torso, etc. A successful theory of part binding would explain how the parts of an object's shape are segregated from the background and then correctly bound together at higher levels of the visual process (Treisman, 1996; cf. Gray, 1999). Scenes are comprised of objects located in space; so explaining how each object is bound to its present location is also essential. This form of binding is known as 'location binding' (Treisman, 1996, p. 171). Since recognizing objects is a function of the ventral system and locating them is a function of the dorsal system, it is essential to show how the visual system connects 'what' properties to 'where' properties. I might also mention that location binding is a species of inter-spatial binding, since its principle function involves binding objects to their present locations relative to other objects. For example, there is an external spatial relation that holds between a vertically oriented cup and a table's horizontal surface when the cup is on *top of* the table (cf. Hummel & Biederman, 1992). Given combinatorial possibilities, the cup could have been to the left of, or below the table, etc. Bayne and Chalmers have aptly dubbed the experience of this external aspect of spatial relatedness *spatial unity*: 'We can say that two conscious states are *spatially unified* when they represent objects as being part of the same space' (2003, p. 25). What they are describing is *inter-spatial* unity, or what Brook refers to as a

‘global representation’ (1994, p. 133). Moreover, there is also the problem of explaining how information from sensory and non-sensory modalities becomes integrated. How does one explain the sensory-motor binding processes involved in coordinating the perceived location of a rapidly approaching baseball with the trajectory of one’s swinging bat? Visual experience rarely, if ever, occurs in isolation from other perceptual phenomena. A person’s unified phenomenal experience of a live musical concert, for example, would likely correlate with a plethora of functional roles performed by distributed neuronal constellations within several brain modalities that subserve vision, audition, tactile stimulation, emotion, and mood. It is now known that emotional experience is associated with several distinct neuronal regions (e.g., the rostral anterior cingulate cortex, an emotion processing subsystem connected to other emotion subsystems, such as the amygdala, orbitofrontal cortex, sectors of the anterior cingulate cortex, and paralimbic structures) and typically involves somatic, perceptual, and/or recall based cognitive activities (see Lane, 2000; see my 2002, 2005). Thus, cognition and experience typically involve many modalities at any given time. This suggests that any forthcoming solutions to unimodal binding problems would still leave unanswered the explanatory question of multimodal binding. The aforementioned problems are important binding problems in their own right, but space will not permit a thorough analysis of all of them here. For the purposes of this paper, I will focus on the object feature binding problem. Let us now consider Crick and Koch’s neurobiological explanation of the unity of an object’s features in visual consciousness.

⁵ A 40-Hertz oscillation repeats every 25 milliseconds.

⁶ See also Hummel and Biederman, 1992; Tanaka, 1993; von der Malsburg, 1996, 1999.

⁷ For example, recent data strongly suggests that the perception of color is occurs before the perception of motion by approximately 80 ms; locations, however, are perceived before the perception of colors (Zeki, 2003).

⁸ A further inference that can be drawn is that neural synchrony fails to explain the experience of visual unity at a time, since it cannot count as the distinctive neural correlate of unified visual consciousness *per se*. In light of the conclusion drawn in the previous paragraph, neural synchrony can neither explain unity at a time nor unity over time.