The Journal of Mind and Behavior
Summer 1996, Volume 17, Number 3
Pages 251–266
ISSN 0271–0137

How We Get There From Here: Dissolution of the Binding Problem

Valerie Gray Hardcastle

Virginia Polytechnic Institute and State University

On the one hand, we think that our conscious perceptions are tied to some stage of whatever processing stream we have. On the other hand, we think that our conscious experiences have to resemble the computational (or brain) states that instantiate them. However, nothing in our alleged stream resembles our experienced perceptions. Hence, a conflict. The question is: How can we go from what we know about neurons, their connections, and firing patterns, to explaining what conscious perceptual experiences are like? No intuitive answer seems plausible. Our perceptual experiences are complex and unified; however, brains divide their processing tasks into small chunks and segregate those smaller pieces across the gray matter. In this essay, I conjecture that what corresponds to our visual perceptions are higher order patterns of bifurcation in an attractor phase space. If I am correct, then the problem of the explanatory gap in philosophy, the binding problem in psychology, and the problem of perception in neuroscience disappear. If the traditional computational perspective is wrong, and sensory processing is not piece-meal, step-wise, and segregated, then there is no need for something in the head to tie things together.

The question that at once confuses and delights philosophers of mind is: How do we get there from here? How do we get consciousness, with all its attendent phenomenology, from the soggy gray matter of brains? The real difficulty in answering that question is not that the answer has not been worked out yet; it is that we do not even know what an answer should look like. Furthermore, when anyone proposes some solution, someone else immediately claims that the solution might explain something else (such as visual processing or attention), but it does not account for the qualitative character of consciousness. The common complaint of any theory of consciousness is

Requests for reprints should be sent to Valerie Gray Hardcastle, Ph.D., Department of Philosophy, Virginia Polytechnic Institute and State University, Blacksburg, Virginia 24061–0126.

that it is never clear why it should be *that* that is conscious. One hard problem in explaining consciousness is articulating why the answer given is a reasonable answer to the question, What is consciousness?, and not a reasonable answer to some other question, such as, How do our visual systems work? (cf., Chalmers, 1996).

There are several versions of this conflict throughout cognitive science. The "explanatory gap" in philosophy points to the discrepancy between any third-person objective description of our physical world (including brains) and our first-person subjective experiences (cf., Levine, 1987). Any explanation approved by our objective sciences will not be able to touch our subjective experiences. A gap in explanation between what we want to explain and what we can explain given our intellectual resources will always remain. The "binding problem" in psychology and the "problem of perception" in neurophysiology both highlight the difficulty in explaining how our consciousness experiences can appear unified, while everything we know about the brain suggests that its processing is very distributed (cf., Treisman, 1986; Treisman and Schmidt, 1982). Our percept of a blue triangle is of an object that is both blue-ish and triangular, and we do not confuse it with a percept of a red square, which appears both reddish and rectangular. If color and shape are processed separately in the brain (and if the information streams never come together), then why do we not get confused? How do we know which color to attach, or bind, to which shape?

The same basic conflicting assumptions drive all of these puzzles. We believe (1) that some stage of our information processing underwrites our conscious experience (though many disagree on what that stage is), and (2) that our conscious experiences must resemble in some fashion the computational or brain states which instantiate them. However, it is clear that nothing in the known information processes of the brain resembles our experienced perceptions. Hence, the puzzles.

In this paper, I argue that these puzzles depend on a particular view of mental processing. I show why this view is probably incorrect and suggest an alternative dynamical systems perspective that dissolves these similar worries concerning explaining consciousness. As will be clear in the discussion, some versions of the puzzle are confused (viz., the version that drives Flanagan's [1992] "mysterians"). Other less cantankerous versions, though, are worthy of careful consideration.

Marr's Paradox

Perhaps the clearest description of the puzzle is given by Akins and Winger (1994) in their discussion of "Marr's Paradox" (see also Akins, in press). Though it is now widely agreed that Marr's theory of vision is false, his theory still provides a nice introduction to the problems involved in conscious perceptions because the theory is so clearly articulated. Moreover, Akins and Winger believe that what they say of Marr's theory will apply *mutatis mutandis* to any other computational theory of perception.

Let us consider how Marr's (1982) theory goes. David Marr was interested in visual processing as shape recognition. To that end, he postulated a hierarchy of four processing levels that extract and compute progressively more abstract object shape information from our activated two-dimensional retinal array. Level one would correspond to the retinal array itself, conceived of as a grid of light intensities. This gives us our most primitive representational information from which we actively construct our perceptions of objects. The second level is a "primal sketch." Here we calculate the spatial organization of the various points of light intensity, things like zero-crossings, discontinuities, boundaries, groups, and so on. The third level is a "2.5 D sketch" which moves from the zero-crossings and boundaries of the primal sketch to representing edges, orientation, and depth. Akins and Winger (1994) compare the representations at this level to "line [drawings] with dotted lines (discontinuities in object surface) and arrows (surface orientation)" [p. 4]. Finally, there is the completed "3 D sketch" which uses the information from the previous level to calculate the actual shapes and spatial orientation of objects. The completed three-dimensional "objects" are represented then in terms of "volumetric primitives" attached to "stick figures" [p. 4].

What is important to notice about this hierarchical theory is that *none of the representations at any of the levels corresponds to what we consciously experience in the visual domain.* Our visual experiences are rich, continuous, unitary, and semantically loaded. They do not *seem like* the output of any of the stages of visual processing.[1] Here then comes the "paradox": nothing in Marr's theory resembles our experienced perceptions. (Both Akins and Winger suggest that retinal images might resemble our perceptions, but that seems wrong to me. My retinas respond to a two-dimensional array of various wavelengths impinging on them; I experience three-dimensional objects projected out into space.)

Moreover, this claim is true for *any* hierarchical theory of vision thus far developed. Nothing in any processing hierarchy resembles our richly unified phenomenology. Hence, the puzzling question: How can we explain what conscious perceptual experiences are like, given what we know about the brain, its structure, and its computations? Our intuitions, such as they are, fail us completely. On the one hand, our experiences of percepts are of

[1]Tye (1991) argues that a version of the 2.5 D sketch corresponds to our phenomenology. Akins (in press) though argues convincingly that revising the third level does not overcome the paradox.

diverse qualia bound together into (complex) units; on the other hand, scientific investigations tell us that brains divide their processing tasks into small, manageable pieces, segregate those tasks across the gray matter, and then leave it at that.

Dennett's Solution

Dennett (1991) hints that the two conflicting assumptions which lead to Marr's "paradox" — that conscious perceptions are tied to some stage of our information processing stream and that conscious experiences resemble the computational states that instantiate them — are both false.[2] Therefore, there is no conflict. He argues that our intuitions of a "Cartesian theater" that there must be a single place in psychological processing that is conscious — is fundamentally misguided. There is no such thing. Instead, the content of many different processing streams might become conscious, depending upon where we probe. We do not process information and then parade the results in some place; rather, we only represent things once, and, depending upon what we do with the information then, it might be conscious (or not) at that time. In essence, Dennett replaces our metaphorical theater with a metaphorical probe. However, this move does not solve our puzzle regarding how something like a brain could ever produce something like consciousness. Moreover, Dennett's proposed "solution" does not work for exactly the same reasons the traditional pictures do not work, for his Multiple Drafts Model is but a more complex version of the same old story.

Consider: either Dennett's "probe" is some sort of moveable Cartesian theater with the powers to produce our experiences in all their unified richness, or it is much weaker. If it is the former, then we have not rid ourselves of Cartesianism at all — we have just made it a traveling side show instead of Broadway. If it is the latter, then consciousness would depend heavily on the (content of the) processing stream probed. In this case, we are back to puzzle again, for in no processing step do we find anything that resembles what we see or hear. If we want to overcome the explanatory gap and explain the binding problem and the problem of perception, we are going to have to be a bit more radical.

How do we solve the difficulty? *Something* is conscious in our brain. However, that nothing in the processing stream obviously resembles our experiences should not disturb us. Few, so far as I know, seriously believe that the representations in our brains *exactly* resemble our phenomenal experi-

[2]At least, this is Akin's reading of Dennett.

ences.[3] That is, no one believes that our experience of a blue triangle corresponds to something in our brains that is blue-ish in hue and shaped like a triangle. Nevertheless, once we deny this very naïve view, we are bereft of clues about what *counts* as a resemblance. It is unclear even the general *sort* of isomorphisms we should be looking for. We are still left with an explanatory difficulty.

This, of course, is the mystery that drives the purported problem of the explanatory gap, as well as shows fallacy in Flanagan's mysterians' reasoning. The mysterians claim that since we cannot even imagine how we would recognize a solution to the problem, one must not exist. It would be far better to claim, I think: so much the worse for our intuitions. There are many things I can't image — five-dimensions, superstrings, the beginning of the universe, infinity, the size of a quark — but my personal inadequacies do not stop science. We should not be moved by what we can or cannot imagine a priori.

But even if we dismiss the philosophical problem of the explanatory gap as a failure of imagination, there are other, perhaps less gripping, difficulties of a similar stripe. The binding problem is one such example. Our visual experiences are detailed, consolidated, and constant; however, the visual system in the brain is fundamentally distributed and operates in a piece-meal fashion (see Ramachandran and Churchland, 1993, for discussion).

In response, Dennett has suggested that we do not perceive what we think we do. Our visual experiences are a lot more serial and fragmented than we think. Others (myself included) have made similar suggestions, distinguishing our actual experiences from our judgments of those experiences (cf., Churchland, 1984; Hardcastle, 1995). We still might be right; however, this fact alone is not going to solve the binding problem. We do experience richly bound perceptions. We know this because we can make mistakes in binding features together. So, how do we do this if there is no theater?

Dennett errs in his analysis when he claims that we do not re-represent things to ourselves, for we do that continuously. The first thing that should strike anyone studying the brain are the massive feedback loops. There are as many things traveling back downstream as there are up. For example, scientists have identified 305 pathways interconnecting the thirty-two cortical visual areas in the macaque monkey (whose visual system is very much like our own). This is about 30% of the number of pathways one would find if our visual system were completely interconnected (Van Essen, Anderson, and Felleman, 1992). Nevertheless what we find is that any forward flow of information is "generally paired with reciprocal pathways that have patterns

[3]The Editor of *The Journal of Mind and Behavior* pointed out to me that Köhler's concept of radical isomorphism hypothesizes an identity between percept and brain state. Nevertheless, Köhler is one on a very short list.

suggesting feedback from a higher to a lower area" (Van Essen et al., 1992, p. 419). It is probably a mistake to talk about or assume a processing "stream" with "higher" and "lower" levels at all. More accurate is to envision dynamical loops exhibiting resonating frequencies and the visual system operating as a single complex unit.

Moreover, if we think along these lines, we can also see a second mistake in Dennett's presentation. He claims that any contentful representation could become conscious if appropriately probed. However, I live in a world that is object-filled and me-centered. Those are the only sorts of visual representations I have. My color experiences have shapes, and vice versa. And if there are no shapes out there, then my brain does its best to convince me that there are. I cannot be aware of any other sort of visual representation. But, if there are not any separate streams that process information in a stepwise fashion, if there are only unitary dynamical loops disturbed by sensory input, then it should not be surprising that there is only one sort of experienced perception (three dimensional objects radiating away from me), even if there is no "theater."

In what follows I conjecture that what corresponds to our visual perceptions are higher order patterns of bifurcation in an attractor phase space. If I am correct, then the binding problem as traditionally conceived disappears. If the traditional computational perspective is wrong, and sensory processing is not piece-meal, step-wise, and segregated, then there is no need for some theater or other to tie things together. In this case, Dennett would ultimately be right about the need to dismiss our intuitions in this matter; he just fails to do so himself.

The Binding Problem

As I explained above, the binding problem refers to how our minds know to join the perception of a triangular shape with the perception of the color blue to give us the single, united experience of a blue triangle when we are in fact confronted with a blue triangle. Why does the blueness not spill over just enough onto surrounding input so that we would be unable to differentiate the triangle from its surrounding context clearly? Why is it that we cannot separate the color from the shape so that we can perceive each independently? If our visual input is confined to a two dimensional array of light intensity differences, why is our experienced world one of discrete objects?

The binding problem has a venable history in philosophy, first appearing in its modern guise in Hume as he speculated on the rules that our minds must follow in uniting simple impressions into more complex ideas. He recognized that the rules of association alone could not be enough: incoming stimuli is always changing, yet we manage to experience ideas as constant across time.

Somehow our "faculties of imagination" step in and fill the gap between stimuli impressions and later memories and ideas. Kant too recognized that mere spatial contiguity and temporal conjunction would not unite certain incoming stimuli into bound impressions at the exclusion of others. (This is the now familiar figure–ground problem in psychology.) As did Hume, Kant concluded that our minds must add something to our perceptions so that our experiences would be of a three-dimensional, object-filled world.

This history recapitulates itself in contemporary cognitive science. Treisman (1986) speculates that before we can assign a semantic content to some stimuli input — before we can see the blue triangle *as* a blue triangle — we first extract "features" from the incoming light pattern array and then construct a temporary representation of the object or event before we can match that temporary representation with stored memories of similar items. Only then, only after a match is found, can we assign a meaning to our perceptual experiences. At the same time, however, cognitive scientists recognize that the brain is radically segregated (Hubel and Livingstone, 1987; Livingstone and Hubel, 1987). As is commonly remarked, the brain processes the "what" and the "where" in separate and distal locations (Van Essen, 1985). Indeed, among the "what" information the brain computes, it responds to edges, colors, and movements using different neuronal pathways (Ramachandran, 1990; Ramachandran and Anstis, 1986).

How does the mind go from modularized outputs to a single and united conscious percept? The simplest and most obvious answer is that the disparate information must come together through some central cortical information exchange. However intuitively pleasing, though, this reply must be false. As far as we can tell, we do not have true association areas in our cortices. There are no neurons that light up in response to our most sophisticated interpretations — the so-called "grandmother" neurons. There are not "convergence zones" where information is pooled and united (Damasio, 1989). Still, the visual features that we extract and compute over separately have to come together in some way since our experiences reflect these features and mnemonic contents joined together into a single experienced unit. If we could figure out how the brain accomplishes this, then we might be able to generate a serious response to Flanagan's mysterians. If we could localize where and explain how the brain unites the segregated outputs, then we would thereby localize and explain some of the processes underlying consciousness.

As with Hume and Kant, cognitive scientists recognize that the story of visual perception told thus far is incomplete. The brain must rely on something besides physical connectedness among cortical areas to generate united percepts. But what? Association, even in the head, is not enough. What would be enough?

Freedom in Time

Time is a likely candidate, for mental and biological functions operate on two different and complementary time scales. There is the *psychological* time scale which characterizes mental processing and is ordered on tenths of a second; there is also the faster *neuronal* time scale characterizing cellular fluctuations and ordered on only thousands of a second. The mean activity of a set of neurons evolves on the longer psychological time scale, but the activity of the individual cells that comprise the set operates on the faster time scale (von der Malsburg, 1987, p. 424). Entraining the fluctuations of coactivated neurons could be the "glue" that binds a distributed firing pattern into a single unit.

Perhaps we should tie bound conscious perceptions to the particular synchronized assemblies of neurons in the cortex. Sensory input would disturb our cortex's natural but noisy spontaneous behavior by forcing various neuronal groups to fire in stable "attractor" based patterns. Without incoming stimuli, our neurons fire randomly and spontaneously. This activity is largely incoherent and produces a background tapestry of noise against which we receive sensory inputs. These inputs shake the neurons up, as it were, and push them to fire rhythmically and in unison. If we map the activity of each neuron (or set of neurons) along a single axis in a multi-dimensional space, then when the neurons oscillate together in their firing patterns, the activity of the neurons would seem to be "attracted" to a small corner of the huge phase space. If we can think of each neuron as representing the values a feature or sub-feature could take in perception, then the various attractor spaces comprised of sets of active neurons would correspond to sets of activated features grouped together in a meaningful unit.

Moreover, since each neuron has preferred ways of firing, there will only be a few attractor regions relative to the dimensionality of the huge phase space. Hence, the object regions in phase space appear to entrain or "attract" stimuli–specific neuronal activity so that when the neurons' natural firing patterns are disturbed, they will fire in the oscillatory patterns keyed to these attractors. These "attractors" could then serve as our neurophysiologically realized "interpretation" of the incoming stimuli.

The hypothesis is that this larger scale order is similar to the sort of order which emerges in the scroll waves of concentration in the Zhabotinske reaction, in the formation of Bénard cells of convection in fluid dynamics, or in populations in ecology, forms of order also explained by attractor patterns in non-linear dynamical systems (Segal, 1965; Winfree and Strogatz, 1983; Yorke and Li, 1975). One advantage to using this sort of analysis is that we can describe the sorts of paths the oscillating neurons take within the multi-dimensional phase space using rather simple nonlinear equations. Since attractors usually occupy only a fraction of the dimensions, their trajectories through phase space are not that complex. As a result, these attractors effectively have a low degree of freedom — there are only a relatively few places in the huge multi-dimensional space that they could be. The sort of order that the low degrees of freedom entail exists within a vast potential disorder — the very high degrees of freedom found in the multi-dimensional space itself. This order out of disorder is the hallmark of nonlinear dynamical analysis.

Taking this sort of perspective means that conscious states could be mapped to the transient firing patterns of groups of neurons whose behavior has a particular mathematical description (see Nunez, 1981, 1989, 1993; see also Kinsbourne, 1988, for a different theory of the same type). And if correct, we should find a higher level structural pattern emerging from lower level neuronal behavior tied to our perceptual experiences. Although I grant this story can seem quite fanciful, remarkably enough, we do seem to find exactly the correspondences I am hypothesizing. Let me give an example of how neural oscillatory patterns index at least the very simple conscious visual perceptions of drug-induced visual hallucinations (cf., Ermentrout, 1979).

Regardless of cause, visual hallucinations are remarkably similar in their early stages. Most persons undergoing an hallucination experience simple geometric shapes, such as a gratings, lattices, checkerboards, cobwebs, funnels, tunnels, or spirals (Klüver, 1967; Siegel, 1977). If we could find some pattern of activity that occurs in the cortices experiencing visual hallucinations that is absent from those that are not, we will have picked out a neural correlate of the hallucinatory perceptual pattern. What is different about the neuronal firing pattern in organisms on hallucinogenic drugs?

There are few studies that record the firing patterns of humans during hallucinatory experiences (though see Winters and Wallach, 1970). However, EEG waves have been recorded over the cortex of hallucinating cats. These show a higher ordered, synchronous, oscillatory pattern of the sort I have been discussing (Adey, Bell, and Dennis, 1962). What is even more interesting is that the patterns recorded over cat cortex are *geometric transformations* of the patterns the cats are presumably seeing.[4] Most significant is that the transformations follow the topographically organized path we can trace from retina to cortex. Though the neuronal projections from retina to cortex are organized topographically, the path twists and turns along the way. If we were to untwist and unturn the cortical pathway, we would find that the patterns recorded moving over cortex (rolls and waves) exactly correspond to what we hallucinate (spirals and grids; see Figure 1).

[4]Ermentrout (1979) details the mathematics.

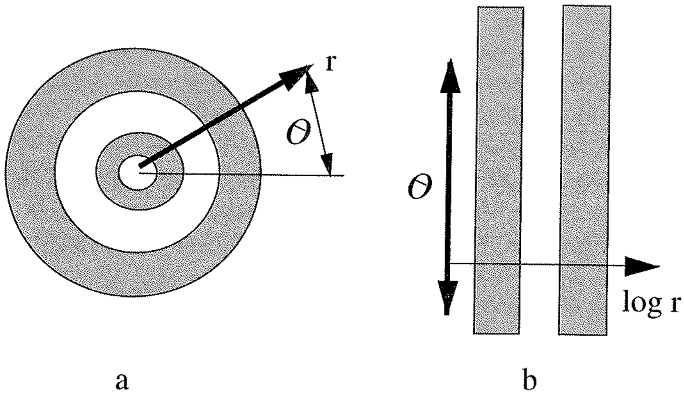

Figure 1: (a) Appearance of a typical "funnel" pattern (a set of concentric circles) in an early hallucinatory experience; (b) The funnel pattern as it appears in EEG recordings of cortex: a set of bars equal in size, arranged in parallel to the y-axes. Notice that the cortical pattern is a simple transformation of a set of concentric circles. (based on Ermentrout, 1979)

Of course, most of our visual experiences are much more complicated than these simple experiences, so we would not expect an EEG recording of normal perception to show such obvious regularities. Instead, we would expect to find more complicated oscillatory patterns. And such patterns are beginning to be uncovered. For example, areas in which stimulus-evoked oscillations have been found include the rabbit olfactory cortex (Freeman, 1978, 1979a, 1979b; Freeman and Schneider, 1982; Freeman and Viana Di Prosco, 1985), the cat, hedgehog, and rat olfactory systems (Adrian, 1942; Bressler and Freeman, 1980; Gault and Leaton, 1963), auditory cortex and the hippocampal areas in the cat (Basar, 1972, 1983; Basar and Ozesmi, 1972; Basar and Ungan, 1973), the hippocampus of guinea pigs (Charpak, Pare, and Llinas, 1992), and the monkey and cat visual cortices (Basar, Rosen, Basar–Eroglu, and Greitshus, 1987; Freeman and van Dijk, 1987; Kruger and Mayer, 1990). Similar higher order resonances have also been found in humans. EEG recordings show the oscillatory patterns above our association, limbic, and motor areas. The auditory pathway, cerebellar cortex, and neocortex can be provoked into an oscillatory pattern when presented some stimuli to respond to (Basar, 1972, 1980; Basar, Gönder, and Ungan, 1976; Basar–Eroglu and Basar 1985; Kreiger and Dillbeck, 1987; Turbes, 1992; see also Freeman and Barrie, 1994). Preliminary results using magnetic field tomography in primary motor and sensory cortices thus far corroborate the brainwave data (Llado et al., 1992; Ribrary et al., 1991).

Freeman (1994a, 1994b, 1994c, 1995) maps these sorts of patterns into attractor phase space, which would make them of same type as the simple patterns found over hallucinating cat cortex. He and Baird both extrapolate from their research on olfactory systems and argue that these stable patterns create temporary meaningful interpretations of stimuli input, as well as create dynamic associative memories. Systems "recognize" an input as belonging to a particular category when the spatial pattern of sensory synaptic input pushes the noisy resting state of the system into a stable oscillatory pattern (Baird, 1991a, 1991b, 1991c; Freeman and Baird, 1987). The perceptual experience then corresponds to oscillatory pattern, which in turn is part of a set of possible patterns concentrated around an attractor. Each attractor would then comprise a single perceptual category (e.g., "is a house," "is food").

Freeman and Baird argue that we should think of the oscillatory patterns as either inhomogeneous limit cycles or weakly chaotic attractors. Regardless though of the attractor's dynamical properties, adopting this perspective means that *things like EEG patterns are the relevant information-bearing states in a biological system.* Only ensembles of neurons perform psychologically relevant computations; hence, the details of the individual neuronal events are not important in explaining psychological phenomena. In sum, it may be that the macroscopic spatial and temporal patterns are true macrostates (Basar, 1988; Basar, Basar–Eroglu, Röschke, and Shütt, 1989; Basar–Eroglu, 1989; Basar–Eroglu and Basar, 1985; Freeman, 1995) and the dimensionality of space for the higher level patterns is the space of "information" processing.

If this view is correct, then the sequence of patterns that the oscillatory networks pass through over time, or the bifurcations from one attractor to another, constitute the "computation" of the system (Dynes and Delautte, 1992; Eggemont, 1991). These oscillations entrain the "micro-activity" of individual neurons in various areas of cortex into a well defined macrostate, while, at the same time, gating irrelevant background "noise." Different brain areas would be discrete systems in which an influx of energy (the sensory input) causes its current internal state to destabilize and bifurcate into a new ordered state that itself becomes momentarily stable. Inputs that the system has already learned trigger one pattern more than the others, allowing that pattern to win a "growth competition."

The mental computation that emerges from bifurcation dynamics and resonant frequencies is very different from a classical von Neumann computing description in which the syntactic structure of mental "sentences" is the currency of cognition. First, this type of system dynamical approach is inherently multi-level (whereas the traditional computer metaphors for the mind operate on only a single level of analysis). The underlying microstructure actually determines which pattern will emerge, but the higher level series of attractor bifurcations through which the network goes are the actual percepts. Second, the oscillatory patterns as described change their shape as the system learns more about its environment. Indeed, what the patterns look like depend upon the connections among individual neurons that previous activity and the corresponding memory have forged. Consequently, there is no constant "syntax" over which functional algorithms can operate. Third, the patterns of oscillation that emerge as we learn to recognize and categorize stimuli are the basic units of cognition. We do not find a feed-forward processor conjoined with some sort of CPU or other connectionist "teacher" function propagating backwards. The oscillations arise spontaneously in the brain in virtue of the massive feedback connections — there is nothing controlling or directing the cortices how to behave.

The Binding Problem Revisited

How exactly would this new view now affect the binding problem? Attractor dynamics and resonant frequencies entail a different conception of information processing than what we traditionally see — descriptions of single cell interactions and individual connectivities. If we consider entire cortical regions as a single network system acting as one resonating unit, then a vehicle for uniting disparate feature-bits would be unnecessary. Instead, the higher level statistical aggregate of firing neurons would unite sets of feature primitives. In other words, binding may not require, strictly speaking, a *process* to tie bits of perceptual information together; instead, binding might be determined solely by the connectivity among neurons between and across the various topographical maps in the cortical areas and the firing responses inherent in the neurons themselves. The oscillatory patterns which emerge as we learn to recognize and categorize stimuli is the basic stuff of cognition, and our perceptual states just are the bifurcations that the network goes through.

Though the underlying microstructure does determine which pattern will emerge, our current techniques for single cell recordings do not allow us to chart the larger scale order as a function of smaller scale interactions. Hence, trying to distinguish the microprocesses which lead to a certain bifurcation in the informational phase is not possible. Indeed, since what we see on the larger scale reflects merely the statistical aggregation of stochastic microprocesses, charting individual cellular behavior might be analogous to following the activity of a single particle in a Bénard cell in order to understand convection. The whole might not be greater than the sum of its parts, but it is certainly different.

Perhaps, as important as the lower level single-cell research is in understanding how neurons behave and how neural networks are actually constructed, it is not the appropriate level of analysis for perceptual experiences. Instead, the true story may be found in system dynamics. The way to describe visual perceptions may be in terms of large scale network behaviors. The *interaction* of the neurons is what is crucial, and we can track the trajectory of the oscillatory macro-patterns through state space in virtue of nonlinear mathematical functions.

The puzzle of explaining consciousness is laid to rest, or, at least, it is set aside. Our inability to point to some place in our information processing stream that resembles our conscious experiences could be an artifact due solely to adopting a von Neumann-style model of computational perception. If we take known neuronal architecture seriously, then we are forced to think in terms of recurrent processing loops in a continually active brain. Perceptions become particular, activated, dynamic, oscillatory loops. Do these loops then *resemble* our phenomenal experiences? They might structurally, at least. Though serious mysterians may rest content gaping at consciousness's bizarre properties, if we assume materialism, then it follows that something in the brain is consciousness. And once we accept that conclusion, we should step aside and let science do its work (though perhaps with appropriate sniping from philosophers). A priori intuitions about what counts as a proper answer should not regulate scientific inquiry. *Informed* intuitions might, but thus far, our intuitions are far from that.

References

Adey, W.R., Bell, F.R., and Dennis, B.H. (1962). Effects of LSD-25, psilocybin, and psilocin on temporal lobe EEG patterns and learned behavior in the cat. *Neurology, 12*, 591–602.

Adrian, E.D. (1942). Olfactory reactions in the brain of the hedgehog. *Journal of Physiology, 100*, 459–473.

Akins, K. (in press). Review of Michael Tye's *The Imagery Debate*. *Philosophical Review*.

Akins, K., and Winger, S. (1994, June). *Ships in the night: Churchland and Ramachandran on Dennett's theory of consciousness*. Paper presented at the Society of Philosophy and Psychology Annual Conference, Vancouver.

Baird, B. (1991a). A bifurcation theory approach to the programming of periodic attractors in network models of olfactory cortex. In D.S. Touretzky (Ed.), *Advances in neural information processing systems 1* (pp. 459–468). San Mateo, California: Morgan Kaufman.

Baird, B. (1991b). Associative memory in a simple model of oscillating cortex. In D.S. Touretzky (Ed.), *Advances in neural processing systems 2* (pp. 68–75). San Mateo, California: Morgan Kaufman.

Baird, B. (1991c). Hierarchical sensory-motor architecture of oscillating cortical area subnetworks. *IJCCN Proceedings*, 1–6. Seattle: Washington.

Basar, E. (1972). A study of the time and frequency characteristics of the potentials evoked in the acoustical cortex. *Kybernetik, 10*, 61–64.

Basar, E. (1980). *EEG–brain dynamics: Relations between EEG and brain evoked potentials*. Amsterdam: Elsevier.

Basar, E. (1983). Synergetics of neural populations: A survey on experiments. In E. Basar, H. Flohr, H. Haken, and A.J. Mandell (Eds.), *Synergetics of the brain* (pp. 183–200). New York: Springer.

Basar, E. (1988). EEG dynamics and evoked potentials in sensory and cognitive processing by the brain. In E. Basar (Ed.), *Dynamics of sensory and cognitive processing by the brain* (pp. 30–55). New York: Springer.

Basar, E., Basar–Eroglu, C., Röschke, J., and Shütt, A. (1989). The EEG is a quasi-deterministic signal anticipating sensory cognitive tasks. In E. Basar and T.H. Bullock (Eds.), *Brain dynamics* (pp. 43–71). New York: Springer.

Basar, E., Gönder, A., and Ungan, P. (1976). Important relation between EEG and brain evoked potentials: II. A system analysis of electrical signals from the human brain. *Biological Cybernetics, 25*, 41–48.
Basar, E., and Ozesmi, O. (1972). The hippocampal EEG activity and a systems analytical interpretation of averaged evoked potentials in the brain. *Kybernetik, 11*, 45–54.
Basar, E., Rosen, B., Basar–Eroglu, C., and Greitshus, F. (1987). The associations between 40 Hz-EEG and the middle latency response of the auditory evoked potential. *International Journal of Neuroscience, 33*, 103–111.
Basar, E., and Ungar, P. (1973). A component analysis and principles derived for the understanding of evoked potentials in the brain: Studies in the hippocampus. *Kybernetik, 12*, 133–140.
Basar–Eroglu, C. (1989). *Eine vergleichende Studie corticaler und subcorticaler ereigniskrelierter Potentiale des Katzengehirns.* Habilitationasschrift: Medizinische Universität zu Lübeck.
Basar–Eroglu, C., and Basar, E. (1985). An analysis of field potentials of the cat cerebellum. *Pflügers Archiv, 403*, R55.
Bressler, S.L. and Freeman, W.J. (1980). Frequency analysis of olfactory system EEG in cat, rabbit, and rat. *Electroencephalography Clinical Neurophysiology, 5*, 19–24.
Chalmers, D. (1996). *The conscious mind: In search of a fundamental theory.* New York: Oxford University Press.
Charpak, S., Pare, D., and Llinas, R. (1992). Entorhinal cortex (E) generates 40 Hz hippocampal oscillations. *Society for Neuroscience Abstracts, 18*, 917.
Churchland, P.M. (1984). *Matter and consciousness.* Cambridge, Massachusetts: The MIT Press.
Damasio, A.R. (1989). The brain binds entities and events by multiregional activation from convergence zones. *Neural Computation, 1*, 123–132.
Dennett, D. (1991). *Consciousness explained.* Cambridge, Massachusetts: The MIT Press.
Dynes, S.B., and Delautte, B. (1992). Phase-locking of auditory-nerve discharges to sinusoidal electric stimulation of the cochlea. *Hearing Research, 58*, 79–90.
Eggemont, J.J. (1991). Rate and synchronization measure of periodicity coding in cat primary auditory cortex. *Hearing Research, 56*, 153–167.
Ermentrout, R. (1979). A mathematical theory of visual hallucination patterns. *Biological Cybernetics, 34*, 137–150.
Flanagan, O. (1992). *Consciousness reconsidered.* Cambridge, Massachusetts: The MIT Press.
Freeman, W.J. (1978). Spatial properties of an EEG event in the olfactory bulb and cortex. *Electroencephography and Clinical Neurophysiology, 44*, 586–605.
Freeman, W.J. (1979a). Nonlinear dynamics of paleocortex manifested in the olfactory EEG. *Biological Cybernetics, 35*, 21–32.
Freeman, W.J. (1979b). EEG analysis gives model of neuronal template matching mechanism for sensory search with olfactory bulb. *Biological Cybernetics, 35*, 221–230.
Freeman, W.J. (1994a). Chaos in the CNS: Theory and practice. In R.J. Greenspan and C.P. Kyriacou (Eds.), *Flexibility and constraint in behavioral systems* (pp. 177–184). New York: John Wiley and Sons, Limited.
Freeman, W.J. (1994b). Neural networks and chaos. *Journal of Theoretical Biology, 171*, 13–18.
Freeman, W.J. (1994c). Role of chaotic dynamics in neural plasticity. In J. van Pelt, M.A. Corner, H.B.M. Uylings, and F.H. Lopes da Silva (Eds.), *Progress in brain research, volume 102* (pp. 319–333). Amsterdam: Elsevier.
Freeman, W.J. (1995). Chaos in the brain: Possible roles in biological intelligence. *International Journal of Intelligent Systems, 10*, 71–88.
Freeman, W.J., and Baird, B. (1987). Relation of olfactory EEG to behavior: Spatial analysis. *Behavioral Neuroscience, 101*, 393–403.
Freeman, W.J., and Barrie, J.M. (1994). Chaotic oscillations and the genesis of meaning in cerebral cortex. In G. Buzaki et al. (Eds.), *Temporal coding in the brain* (pp. 13–37). Berlin: Springer–Verlag.
Freeman, W.J., and Schneider, W. (1982). Changes in spatial patterns of rabbit olfactory EEG with conditioning to odors. *Psychophysiology, 19*, 44–56.
Freeman, W.J., and van Dijk, B.W. (1987). Spatial patterns of visual cortical EEG during conditioned reflex in a rhesus monkey. *Brain Research, 422*, 267–276.

Freeman, W.J., and Viana Di Prosco, G. (1985). EEG spatial pattern differences with discriminated odors manifest chaotic and limit cycle attractors in olfactory bulb of rabbits. In G. Palm and A. Aertson (Eds.), *Brain theory* (pp. 97–119). Berlin: Springer–Verlag.
Gault, F.P., and Leaton, R.N. (1963). Electrical activity of the olfactory system. *Electroencephalography and Clinical Neurophysiology, 15*, 299–304.
Hardcastle, V.G. (1995). *Locating consciousness*. Philadelphia: John Benjamins Press.
Hubel, D.H., and Livingstone, M.S. (1987). Segregation of form, color, and stereopsis in primate area 18. *The Journal of Neuroscience, 7*, 3378–3415.
Kinsbourne, M. (1988). Integrated field theory of consciousness. In A.J. Marcel and E. Bisiach (Eds.), *Consciousness in contemporary science* (pp. 239–256). Oxford: Clarendon Press.
Klüver, H. (1967). *Mescal and the mechanisms of hallucinations*. Chicago: University of Chicago Press.
Krieger, D., and Dillbeck, M. (1987). High frequency scalp potentials evoked by a reaction time task. *Electroencephalography and Clinical Neurophysiology, 67*, 222–230.
Kruger, J., and Mayer, M. (1990). Two types of neuronal synchrony in monkey striate cortex. *Biological Cybernetics, 64*, 135–140.
Levine, J. (1987). *The Nature of Psychological Explanation* by Robert Cummins, a critical notice. *The Philosophical Review, 96*, 249–274.
Livingstone, M.S., and Hubel, D.H. (1987). Psychophysical evidence for separate channels for the perception of form, color, movement, and depth. *The Journal of Neuroscience, 7*, 3416–3468.
Llado, F., Ribrary, U., Ioannides, A., Volkman, J., Joliot, M., Mogilner, A., and Llinas, R. (1992). Coherent oscillations in primary motor and sensory cortices detected using MEG and MFT. *Society for Neuroscience Abstracts, 18*, 848.
Marr, D.C. (1982). *Vision: A computational investigation into the human representation and processing of visual information*. San Francisco: Freeman.
Nunez, P.L. (1981). *Electric fields of the brain. The neurophysics of EEG*. New York: Oxford University Press.
Nunez, P.L. (1989). Generation of human EEG by a combination of long and short range neocortical interactions. *Brain Topography, 1*, 199–215
Nunez, P.L. (1993). *Neocortical dynamics and human EEG rhythms*. New York: Oxford University Press.
Ramachandran, V. (1990). Visual perception in people and machines. In A. Balke and T. Troscianko (Eds.), *AI and the eye* (pp. 21–77). New York: John Wiley and Sons Limited.
Ramachandran, V., and Anstis, S.M. (1986). The perception of apparent motion. *Scientific American, 254*, 101–109
Ramachandran, V., and Churchland, P.S. (1993). Filling in versus finding out. In B. Dahlbom (Ed.), *Dennett and his critics* (pp. 28–52). New York: Blackwell.
Ribrary, U., Ionnides, A., Singh, K.D., Hasson, R., Bolton, J.P., Lado, F., Mogilner, A., and Llinas, R. (1991). Magnetic field tomography of coherent thalamocortical 40-Hz oscillations in humans. *Proceedings of the National Academy of Sciences of the United Sates of America, 88*, 11037–11041.
Segal, L.A. (1965). The nonlinear interaction of a finite number of disturbances to a layer of fluid heated from below. *Journal of Field Mechanics, 21*, 359.
Siegel, R.K. (1977). Hallucinations. *Scientific American, 237*, 132–140.
Treisman, A. (1986). Features and objects in visual processing. *Scientific American, 254*, 114b–125.
Treisman, T., and Schmidt, H. (1982). Illusory conjunctions in the perception of objects. *Cognitive Psychology, 14*, 107–141.
Turbes, C.C. (1992). Brain state matrix: Computational approaches to neocortical limbic rhythms. *Society for Neuroscience Abstracts, 18*, 848.
Tye, M. (1991). *The imagery debate*. Cambridge, Massachusetts: The MIT Press.
Van Essen, D.C. (1985). Functional organization of primate visual cortex. In A. Peters and E.G. Jones (Eds.), *Cerebral cortex* (pp. 259–329). New York: Plenum Publishing.
Van Essen, D.C., Anderson, C.H., and Felleman, D.J. (1992). Information processing in the primate visual system: An integrated systems perspective. *Science, 255*, 419–423.

von der Malsburg, C. (1987). Synaptic plasticity as basis of brain organization. In J.-P. Changeux and M. Konishi (Eds.), *The neural and molecular bases of learning* (pp. 411–432). New York: John Wiley and Sons.

Winfree, A.T., and Strogatz, S.H. (1983). Singular filaments organize chemical waves in three dimensions. *Physica, 8D*, 35–49.

Winters, W.D., and Wallach, M.B. (1970). During induced states of CNS excitation: A theory of hallucinosis. In D.H. Efron (Ed.), *Psychotomimetric drugs* (pp. 193–214). New York: Raven Press.

Yorke, J.A., and Li, T.-Y. (1975). Period three implies chaos. *American Mathematical Monthly, 82*, 985–992.

Zeki, S. (1993). The visual image in mind and brain. *Scientific American, 267*, 68–77.