

Mental Representations: A computational-neuroscience scheme

Marius Usher

Department of Psychology, University of Kent at Canterbury
Kent, CP2 7NP, UK
e-mail: M.Usher@ukc.ac.uk

Ernst Niebur

Krieger Mind/Brain Institute and Dept. of Neuroscience
The Johns Hopkins University, Baltimore, MD 21218, USA
e-mail: niebur@jhu.edu

December 14, 1998

Abstract

We discuss a series of problems facing referential theories of mental representations and we propose a scheme based on neurophysiological principles that should meet previous limitations. According to this scheme, mental representations are brain traces linked to stimuli in the environment, via a causal but probabilistic process of *categorical perception*, and fluctuations in activity reflect fluctuations in the confidence-level of perceptual and cognitive hypotheses. The scheme provides an explanation for cases of misrepresentation and is consistent with the abundance of recurrent connections in the cortex, which play an important role in mediating a process of interpretation and of binding of relational properties via temporal synchronisation.

Introduction

Despite the fact that mental representations are a central concept in cognitive psychology and in neuroscience, a basic understanding of their nature seems to be lacking. Indeed, all theories for mental representations, proposed in the philosophy of mind literature, seem to suffer from difficulties when accounting for the reference of mental states (Crane, 1995; Cummins, 1991). In addition, recent neurophysiological data indicates that the classical scheme based on a bottom-up hierarchy of encoding cells is inadequate; neural representations seem to be distributed and most of the input is contributed by feedback and lat-

eral connections rather than by afferent connections (Douglas, Martin, & Whitteridge, 1988; Douglas & Martin, 1991). This seems therefore to contradict a simple scheme where neural representations are simple mappings implemented via bottom-up projections from sensory stimulation to internal responses (see Peschl & Riegler and von Stein, this volume). Moreover, even when tested under very constrained and controlled laboratory conditions, neural responses are noisy (Softky & Koch, 1993) and spontaneous spatio-temporal activation patterns influence subsequent responses to stimuli (Arieli, Sterkin, Grinvald, & Aertsen, 1996). The theoretical implications of those findings, in relation to neural codes, are yet unclear despite a very stimulating debate (Softky, 1995; Shadlen & Newsome, 1995; Ferster & Spruston, 1995).

For these reasons, a re-evaluation of the classical theories of neural representation seems now to be needed. One approach is to give up on the referential characterisation of mental representations: that they refer to (or are about) objects and events in the external world/environment. As we will argue below, this is likely to make the whole concept of representation non-explanatory and unable to provide the required mediation for psychological processes. The strategy followed here is different. We will show that a consistent scheme for referential neural representations in line with the neurophysiological data mentioned above can be construed. This will require, however, revision

of the mechanisms by which representations are thought to operate in the brain (a key factor is to replace deterministic with probabilistic processes of reference) and to refine the effective powers of such representations (which become active computational processes.)

Our paper is organized as follows. First we briefly discuss the motivation for preserving a referential theory of mental representations. Second, we review conceptual and neurophysiologically based arguments that are problematic for all the present theories of referential mental representations. Third, we develop a scheme based on *cathegorical perception* and show that it can overcome most of these problems. Finally, we will show (by computer simulations) that neural representations generated according to this scheme, are active processes which can mediate contextual computations and perform binding of relational properties.

The need for referential representations

There are two main motivations for preserving the referential characteristic of mental representations. First, from the philosophical perspective, a major characteristic of mental states is their *intentionality* (Brentano, 1874; Searle, 1983); this is the property (special to mental states) of referring (or *being about*) some possible (but not necessarily actual) state of events in the external world. It is beyond the scope of this paper to engage in a metaphysical discussion on the status of this "external state": "things in themselves" or "actuality of experience" (see Peschl and Riegler, this volume). It is enough to notice in this regard that mental representations are about something *external* to themselves and inter-subjective (to avoid solipsism and to explain communication, mental representations of different agents need to be able to refer to the same entities).¹

Second, the concept of representation is central in psychology as it allows to explain the complex cognitive apparatus observed in human per-

¹An external reality (or environment) in which agents exist, and interact with, is a simple but not unique way to achieve this. An alternative scheme could be based on Leibniz' monade-type metaphysics, according to which experiences of various agents are synchronised and kept consistent by "divine" intervention. Even then however, mental representation will need to be about (but different from) the synchronized properties of the monades.

formance. In this regard, complex behavior is successfully explained in terms of an economical data structure that reflects semantic organisation (Collins & Quillian, 1969; Collins & Loftus, 1975). While it is universally accepted that such a semantic organisation reflecting essential characteristics in the environment of an agent will enable it to perform actions with survival value, survival is definitely not a necessary condition for possessing meaning or representations (a person being kept captive and killed is able to represent to herself the unpleasant circumstances on which she, unfortunately, has no control). Moreover, even if we neglect such non-survival circumstances, equating representations with ability for survival (as proposed in some constructivist approaches presented in this volume) but not with a reflection of regularities within the environment, results in a loss of explanatory power over how the survival behavior (as well as all other behaviors) is mediated.²

Conceptual problems for mental representations

The problem of reference

Simply put, the problem is to explain how can brain states (or any other object) represent (refer or point to) objects in the world (other than by convention³). There are in fact only two principles that have been seriously considered by philosophers to account for a naturalistic and reductive explanation of mental representation: *resemblance* and *causation* (Fodor, 1987; Cummins, 1991; Crane, 1995). However, theories based on both of these principles, have been shown to suffer from severe difficulties (Crane, 1995; Cummins, 1991).

According to the resemblance principle, mental representations represent (refer to specific objects) in virtue of a similarity type of relation between the object and the brain trace. Even when more

²Unfortunately, some connectionist studies find it enough to demonstrate that a network can be trained to perform a task. Without additional explanation (such as by examining the structure of internal representations and the way they contribute to task performance), the results are not less mysterious than the phenomenon being explained.

³While convention is a reasonable explanation for many kind of representations, say, words, such a solution is totally unsatisfactory for brain states (which unlike words should be an instance of "underived" representation), due to infinite regression.

modern version of similarity relations (based on homomorphisms rather than on simple images) are used, it seems unlikely that similarity is sufficient for obtaining mental representations. This can be illustrated with the following example used by Putnam (1981). An ant moves across the sand leaving a trace which by pure accident resembles the shape of Winston Churchill. It seems quite compelling that despite its similarity this accidental trace does not represent the person of Winston Churchill, in itself. Thus resemblance cannot be a sufficient condition for representation⁴.

The causation principle seems a much better candidate for a theory of representation. Accordingly, X represents Y if and only if X is caused by Y⁵. This seems indeed to be the missing ingredient in the previous example: the trace of the ant has no causal relationship to the person of Winston Churchill, while a perceptual brain state is caused by the presence of its perceptual object. However, a theory of representation based on causation alone is still insufficient, as it suffers from two ubiquitous problems: the problem of error and the problem of interpretation.

The problem of mis-representation

To illustrate the former, notice that a valid theory should be able to account for errors in the representation. For example, under short exposure conditions, one might misperceive the letter "A" for an "H", or when looking for one's dog at night, one might misjudge a passing cat for the expected dog. The problem of mis-interpretation is that the underlying brain state (mediating the recognition of the dog) is causally linked with the passing-cat. Thus according to the causality principle the dog-state represents a cat instead of misrepresenting a dog. It is thus not possible to account on the basis of causation alone for the difference between veridical and erroneous perception. There are two type of mis-representation, unsystematic and systematic (often mediated by context as in the dog-

⁴One might also consider the fact that the words "Winston Churchill" spoken by a human being who is familiar with modern history do represent the person of Churchill despite of bearing no similarity to him. This could be taken to imply that "resemblance" is not a necessary condition, either; note however that word-representations are derived (i.e., they rely on convention).

⁵A technical detail used sometimes, is to require causation based on natural regularity (or "reliable indication") (Crane, 1995). As this will make no difference in the following discussion we use only the simple causation label instead.

example above), which need to be accounted for.

The problem of interpretation

An important characteristic of perceptual processes is that they are interpretative. For example a Necker cube (or other ambiguous figures; see e.g. Gregory, 1998) can be perceived in one of two exclusively ways. This may be a result of the economical constraints the cognitive system is subject to; after all, an infinity of possible stimuli needs to be mapped onto a much smaller repertoire of concepts. Wittgenstein (1958) who was very much concerned with the problem that any (mental) image still needs to be interpreted, seemed inclined to conclude that mental images are irrelevant to behavior. There is however an alternative. Since mental representations can not rely on external interpretation (to avoid infinite regression), the perceptual process itself should provide the interpretation, or in other words representations need to be active interpretative processes. Definitively there is nothing in virtue of which a merely causal process, such as, for example, a photographic projection, could provide an interpretation of the ambiguity. Thus we see that in order to refer (or represent) a special interpretation process is required. Important questions arise: What are the essential characteristics of such an interpretation process? Is there a specific process of interpretation taking place in the brain during the perception of objects, different from a photographic projection and which could account for the ability of mental/brain states to represent external objects?

Neurophysiological constraints

As we mentioned above, most of the input to cortical neurons is produced by feed-back (recurrent) connections. As illustrated in attractor networks (Amit, Sagi, & Usher, 1990; Sompolinsky & Kanter, 1986; Elman, 1990, see also von Stein, this volume) this results in context behavior: the new states of the network depend not *only* on the sensory input but also on the history of previous states of activation. Similarly, in the visual cortex, contextual modulation effects (due to stimuli outside the classical receptive field) have been well documented (Allman, Miezin, & McGuinness, 1985; Knierim & Van Essen, 1992) and explained on the basis of lateral interactions (Stemmler, Usher, & Niebur, 1995). Thus, while recurrent connectivity is an excellent scheme for

implementing contextual behavior, it runs against the (neo)behaviorist scheme that associates stimuli with consistent internal responses.

Another important problem often discussed in computational neuroscience is the format of neural representations: localistic vs distributed. Localistic representations suffer from many problems in relation to computational power (the combinatorial problem) and inflexible and non-adaptive behavior (McClelland, Rumelhart, & the PDP Research Group, 1986b; McClelland, Rumelhart, & Hinton, 1986a) While distributed representations fair much better in this regard, they need to face the problem of binding and segmentation which results when more than one object-stimulus need to be represented, at the same time (Singer & Gray, 1995). An attractive scheme for binding visual features into a coherent percept consists of synchronising the activity of their neural representations (Milner, 1974; von der Malsburg, 1981; Crick & Koch, 1990b; von der Malsburg, 1996). Recent neuropsychophysical (Singer & Gray, 1995) and psychophysical studies (Fahle, 1993; Leonards, Singer, & Fahle, 1996; Usher & Donnelly, 1998) provided support for that hypothesis. For this scheme to operate, however, distributed neural representations need to be able to synchronise or desynchronise in relation to global stimulus properties.

In the next section we present a scheme for neural representations that account for reference while allowing for contextual behavior, and according to which representations are active interpretative processes that are also able to engage in binding and segmentation.

A scheme for Neural representations

Categorisation and categorical perception

Perceptual errors that are non-systematic or which are due to systematic-contextual effects are difficult to account for within the present referential theories for representation. We believe that this is the case because, despite recent progress in *probabilistic causality* (Suppes, 1970; Eells, 1991) most attempts to explain reference rely on deterministic processes to mediate information transmission in causal chains (Dretske, 1981, 1983), which, in

turn, makes them unable to deal with ambiguity and probabilistic behavior inherent in most human (and animal) behavior (see, e.g., Suppes' reply to Dretske, 1983).

A consistent solution to the reference problem can be obtained, however, if we relax the requirements of deterministic causal processes to include probabilistic ones. The idea is simple. Not each time that I think to see a cat, there is a cat in front of me (and sometimes I think to see one without it being there), however when I think to see one, its more likely that a cat is there (than anything else).⁶ This means that because of context (or perceptual noise), we can never have perfect confidence that what we think to perceive is veridical in a single particular situation, however, we have a good basis for believing that (averaging over situational contexts) it is quite likely to be.

We believe that indeed, brain states can represent external objects, not due to some magical brain properties, but because of special interpretation-like causal but *probabilistic* processes taking place during sensory perception. Moreover, we believe that such a process is now well understood in cognitive psychology and neuroscience. We turn now to the process of *categorisation* and *categorical perception* which has been demonstrated both in human and in animal subjects (Harnad, 1987; LaMotte & Mountcastle, 1979; Wytttenbach, May, & Hoy, 1996), and which we believe to have the necessary characteristics for solving the intentionality problem.

There are several ways in which the process of categorical perception is fundamentally different from a mapping such as a shadow or a photograph. Unlike those passive mappings which transform a continuity of objects to a similar continuity of projections (e.g., infinitesimal rotations or distortions of the objects are mapped into the projection), the process of categorical perception involves a discontinuous decision-like process: continuous families of similar stimuli are mapped (probabilistically, depending on environmental or perceptual noise) onto discrete categorical concepts. Thus unlike a photographic projection, categorical perception is an active interpretation process in which small irrelevant transformations (distortions) are ignored

⁶This can be formalised by using the mathematical tool of mutual-information (Cover & Thomas, 1991), which also takes care of problems arising from differences in environmental frequencies (maybe cats are extremely rare compared to dogs).

and which involves perceptual boundaries⁷. Computational models for perception and memory (Anderson, Silverstein, Ritz, & Jones, 1977; Hopfield, 1984) have also demonstrated how systems of neural elements can generate categorical perception, generating computations which are error-corrective due to *attractor* dynamics. As a consequence, perceptual states do not need an external interpretation; they are part of the interpretation.

To understand how the process of categorisation solves the problems of perceptual mental representations, we should consider the *ensemble* of objects (characterised by some statistical distribution) that an organism with a cognitive system encounters and interacts with during its life. One can then propose that a brain categorisation state represents the most *typical* object that is likely to be mapped to (categorised by) that state (this is well defined by the statistical characteristics of the ensemble of objects and by the stochastic categorisation process⁸). The neural architecture needed to mediate processes of categorical perception develops with learning experience, so that the repertoire of categorisation states is continually refined. This scheme bears some similarity to the approach of Dretske (1981) who emphasised the role of the learning period for providing the reference of representational states; the statistical ensemble reflects the learning experience of the organism. However, unlike in that scheme (Dretske, 1981), mis-representations are naturally accounted for without the need to rely on the assumption that the learning stage is errorless, in order to distinguish true from incorrect reference. They are the instances when an object has been categorised (due to noise or to contextual influences) to a categorisation state which is atypical for it. Such occurrence should be rare (otherwise the categorisation state would be a typical one)⁹ The process of categorisation can thus be seen as

⁷Of course, some simple forms of “categorical perception” takes place even in simple physical systems such as a thermostat which maps all temperature conditions above a specified boundary to a “hot” category and the temperatures below the boundary to a “cold” category.

⁸Such a “typical object” can correspond to what cognitive psychologists call a *prototype*.

⁹It is interesting to note that this mechanism for causality satisfies the principle described by Fodor (1990) as asymmetric dependency. If object A is misrepresented by the representation of object B, the causal process involved is mediated (thus dependent) on the representation of B (its attractor pattern), while the opposite is not true (categorizing object A as A does not depend on the representation of object B).

an instance of statistical hypothesis testing: finding which object (out of a set of alternatives) is the more likely to have caused the perceptual state. Surely, such perceptual predictions might be wrong: the brain state might be generated as result of an optical illusion, a drug induced hallucination, etc. Even under such “atypical” conditions, however, the perceptual brain state refers to and represents the object which would generate the trace in typical situations relative to the past experience of the organism.

Neural implementation

We propose that the following scheme provides a scheme for the implementation of representation and information coding in the cortex.

Categories are represented by the identity of (distributed) neural populations (Georgopoulos, Kettner, & Schwartz, 1986; Miyashita & Chang, 1988). Temporal patterns (firing rate; Barlow (1972) or synchrony; Decharms and Merzenich (1996)) signal the confidence level of perceptual hypothesis. In addition, synchrony of neural populations encodes *relational* properties (i.e., binding) between the elemental properties (features) represented by those populations (Gray, König, Engel, & Singer, 1989; Crick & Koch, 1990a, 1990b; Singer & Gray, 1995) and possibly mediates the process of attentional selection (Niebur & Koch, 1994).

The way in which distributed cell populations mediate categorisation is explained by attractor networks (Anderson et al., 1977; Hopfield, 1984; Amit, 1989). In the following we will focus on another processes mediated by neural populations: dynamic binding.

Synchrony and binding

In our previous models (Usher & Niebur, 1995; Usher, Stemmler, & Niebur, 1996) we have also shown that synchronised neural activity corresponding to grouping of coherent visual information arises in simple models of the primary visual cortex due to recurrent lateral connections. For instance, Figure 1 illustrates an elementary form of binding in a model of visual cortex, implemented as a two-dimensional sheet of integrate-and-fire neurons with local and long-range connectivity. To demonstrate binding, we computed the cross-correlations between local field potentials (LFP) recorded (in the simulation) at two sites which

are stimulated either by the same or by orthogonal orientation gratings (within the same distance between each pair of sites). As shown in Figure 1 (bottom), the correlation between cells receiving the same input is significantly stronger than that between cells that receive different inputs. One should notice that this effect is stimulus dependent and can not be explained by the simple fact that connected cells synchronise (Mirolo & Strogatz, 1990). In fact, the degree of correlation between two given cells depends on their stimulation and is not only a function of the connectivity.





Figure 1: Activity in neuronal array. **Top:** Array of 100×100 cells where the gray-level of each pixel corresponds to the preferred orientation of the respective cell (lightest gray: horizontal, light gray: 45° , dark gray: 90° , black: 135°). White pixels represent spiking cells. **Bottom:** Cross-correlation functions of local field potentials. The dotted curve is between LFP recorded at sites stimulated with the same orientation grating and the full line is between LFP recorded at sites stimulated with orthogonal gratings.

Rate Fluctuations

We propose that the irregularity of neural discharge (Softky & Koch, 1993) is due to a large extent to self-generated fluctuations in the firing rate of neural populations. In another article (Usher & Niebur, 1999) we show that experimental data are consistent with the predictions of network models which generate their own rate fluctuations (Usher, Stemmler, Koch, & Olami, 1994; Usher, Stemmler, & Olami, 1995; Hansel & Sompolinsky, 1996). Such self generated fluctuations arise due to the feedback connections, in particular for center-surround lateral connectivity, enabling the networks to respond fast by rapidly amplifying fluctuations in the external input (Usher et al., 1995). Such fluctuations might underlie internal computations and reflect fluctuations in the certainty and the confidence level of perceptual/cognitive hypotheses. Optical imaging techniques have recently revealed evidence for such fluctuations (Arieli et al., 1996). In particular, Arieli et al. found that the variability of evoked visual responses is due to a large extent to the variability in the spontaneous activity, present previously to the sensory stimulation.

Conclusion

We have tried to demonstrate that, consistently with new findings in neurophysiology and progress in computational neuroscience, a scheme for referential neural representation can be developed. According to this scheme, mental representations are implemented by distributed neural populations which are causally linked with perceptual stimuli via the probabilistic process of categorical perception. This process requires an interpretation from a continuum of possible stimuli to a finite repertoire of concepts and is subject to contextual modulations. Moreover, we have shown that neural populations have the ability to synchronise, and to mediate binding between features of the same object. Rate fluctuations might reflect changes in confidence level or in the certainty of perceptual hypotheses, which occur even in the absence of a stimulus, within the spontaneous ongoing cortical activity (Arieli et al., 1996). This suggests an active view of the perceptual/cognitive system where even in the absence of inputs, various hypotheses are being continuously entertained at various degrees (and bound in various combi-

nations), and information from the past is integrated within present responses, supporting an active view of cognition.

Acknowledgements

We wish to thank Michael Herrmann for very stimulating discussions on Bayesian hypothesis testing. The work of EN is supported by a Sloan Fellowship and by the Markey Foundation.

References

- Allman, J., Miezin, F., & McGuinness, E. (1985). Direction- and velocity-specific responses from beyond the classical receptive field in the middle temporal visual area (MT). *Perception, 14*, 105–126.
- Amit, D. J. (1989). *Modelling Brain Function*. Cambridge University Press, Cambridge.
- Amit, D. J., Sagi, D., & Usher, M. (1990). Architecture of attractor neural Networks performing cognitive fast scanning. *Network, 2*, 189–216.
- Anderson, J. A., Silverstein, J. W., Ritz, S. A., & Jones, R. S. (1977). Distinctive features, categorical perception and probability learning: Some applications of a neural model. *Psychological Review, 84*, 413–451.
- Arieli, A., Sterkin, A., Grinvald, A., & Aertsen, A. (1996). Dynamics of ongoing activity: explanation of the large variability in evoked cortical responses. *Science, 273*, 1868–1871.
- Barlow, H. B. (1972). Single units and sensation: a neuron doctrine for perceptual psychology?. *Perception, 1*, 371–394.
- Brentano, R. (1973, orig. published 1874). *Psychology from an empirical standpoint* (translated by A. Pancurello and D. Terrell and L. McAlister). Humanities, New York.
- Collins, A. M., & Loftus, E. F. (1975). A spreading activation theory of semantic processing. *Psychological Review, 5*, 85–88.
- Collins, A. M., & Quillian, M. R. (1969). Retrieval time from semantic memory. *J. of Verbal Learning and Verbal Behavior, 8*, 240–247.
- Cover, T. M., & Thomas, J. A. (1991). *Elements of Information Theory*. Wiley, New York.
- Crane, T. (1995). *The mechanical mind*. Penguin, London, UK.
- Crick, F., & Koch, C. (1990a). Some reflections on visual awareness. *Cold Spring Harbor Symp. Quant. Biol., 55*, 953–962.
- Crick, F., & Koch, C. (1990b). Towards a neurobiological theory of consciousness. *Seminars in the Neurosciences, 2*, 263–275.
- Cummins, R. (1991). *Meaning and mental representation*. A Bradford Book. MIT Press, Cambridge, MA.
- Decharms, R. C., & Merzenich, M. M. (1996). Primary cortical representation of sounds by the coordination of action potential timing. *Nature, 381*, 610–613.
- Douglas, R. J., & Martin, K. A. C. (1991). A functional microcircuit for cat visual cortex. *J. Physiol., 440*, 735–769.
- Douglas, R. J., Martin, K. A. C., & Whitteridge, D. (1988). Selective responses of visual cortical cells do not depend on shunting inhibition. *Nature, 332*, 642–644.
- Dretske, F. (1981). *Knowledge and the flow of information*. Blackwell, Oxford.
- Dretske, F. (1983). Précis of Dretske *Knowledge and the Flow of Information*. *Behavioral and Brain Sciences, 6*(1), 55–63.
- Eells, E. (1991). *Probabilistic causality*. Cambridge University Press, Cambridge, UK.
- Elman, J. (1990). Finding Structure in Time. *Cognitive Science, 14*, 179–211.
- Fahle, M. (1993). Figure ground discrimination for temporal information. *Proceedings of the Royal Soc. London, 254*, 199–203.
- Ferster, D., & Spruston, N. (1995). Cracking the neuronal code. *Science, 270*, 756–757.
- Fodor, J. (1987). *Psychosemantics*. MIT Press, Cambridge, MA.
- Fodor, J. (1990). *A theory of content and other essays*. MIT Press, Cambridge, MA.

- Georgopoulos, A. P., Kettner, R. E., & Schwartz, A. B. (1986). Neuronal population coding of movement direction. *Science*, *233*, 1416–1419.
- Gray, C., König, P., Engel, A., & Singer, W. (1989). Oscillatory responses in cat visual cortex exhibit inter-columnar synchronization which reflects global stimulus properties. *Nature*, *338*, 334–337.
- Gregory, R. L. (1998). *Eye and Brain: The psychology of seeing* (5 edition). Oxford University Press, Oxford.
- Hansel, D., & Sompolinsky, H. (1996). Chaos and synchrony in a model of a hypercolumn in visual cortex. *Journal of Computational Neuroscience*, *3*(1), 7–34.
- Harnad, S. (1987). *Categorical perception: The groundwork of cognition*. Cambridge University Press, Cambridge, UK.
- Hopfield, J. J. (1984). Neurons with Graded Responses Have Collective Computational Properties Like Those of Two-State Neurons. *Proc. Nat. Acad. Sci., USA*, *81*.
- Knierim, J. J., & Van Essen, D. C. (1992). Neuronal responses to static texture patterns in area V1 of the alert macaque monkey. *J. Neurophysiology*, *67*(4), 961–980.
- LaMotte, R. H., & Mountcastle, V. B. M. (1979). Disorders in somethesis following lesions of parietal lobe. *J. Neurophysiology*, *42*(2), 400–419.
- Leonards, U., Singer, W., & Fahle, M. (1996). The Influence of Temporal Phase Differences on Texture Segmentation. *Vision Res.*, *36*(17), 2689–2697.
- McClelland, J., Rumelhart, & Hinton, G. (1986a). The appeal of parallel distributed processing. In Rumelhart, D., & McClelland, J. (Eds.), *Parallel Distributed Processing*, Vol. 1, chap. 1, pp. 3–44. MIT Press, Cambridge.
- McClelland, J., Rumelhart, D., & the PDP Research Group (1986b). *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*, Vol. 2. MIT Press, Cambridge.
- Milner, P. (1974). A model for visual shape recognition. *Psychol. Rev.*, *81*, 521–535.
- Mirollo, R., & Strogatz, S. (1990). Synchronization of pulse-coupled biological oscillators. *SIAM J. on Applied Mathematics*, *50*(6), 1645–1662.
- Miyashita, Y., & Chang, H. (1988). Neural correlate of pictorial short term memory in primate temporal cortex. *Nature*, *331*, 68–70.
- Niebur, E., & Koch, C. (1994). A model for the neuronal implementation of selective visual attention based on temporal correlation among neurons. *Journal of Computational Neuroscience*, *1*(1), 141–158.
- Putnam, H. (1981). *Reason, truth, and history*. Cambridge University Press, New York.
- Searle, J. R. (1983). *Intentionality: An essay in the philosophy of mind*. CUP, Cambridge.
- Shadlen, M. N., & Newsome, W. T. (1995). Is there a signal in the noise?. *Current Opinion in Neurobiology*, *5*, 248–250.
- Singer, W., & Gray, C. M. (1995). Visual feature integration and the temporal correlation hypothesis. *Annu. Rev. Neurosci.*, *18*, 555–586.
- Softky, W. (1995). Simple codes versus efficient codes. *Current Opinion in Neurobiology*, *5*, 239–247.
- Softky, W., & Koch, C. (1993). The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs. *J. Neurosci.*, *13*(1), 334–350.
- Sompolinsky, H., & Kanter, I. (1986). Temporal Association in Asymmetric Neural Networks. *Physical Review Letters*, *57*, 2861–2864.
- Stemmler, M., Usher, M., & Niebur, E. (1995). Lateral Interactions in Primary Visual Cortex: A Model Bridging Physiology and Psychophysics. *Science*, *269*, 1877–1880.
- Suppes, P. (1970). *A probabilistic theory of causality*. North Holland, Amsterdam.
- Usher, M., & Donnelly, N. (1998). Visual synchrony affects binding and segmentation in perception. *Nature*, *in press*.

- Usher, M., & Niebur, E. (1995). Stimulus-dependent synchronization in a model of orientation-tuned spiking neurons. In Bower, J. M. (Ed.), *Computation and Neural Systems*, Vol. 3, pp. 341–346. Kluwer, Norwell, MA.
- Usher, M., & Niebur, E. (1999). Firing Patterns And Neural Codes: Models, data and implications. *in preparation*.
- Usher, M., Stemmler, M., Koch, C., & Olami, Z. (1994). Network Amplification of Local Fluctuations Causes High Spike Rate Variability, Fractal Firing Patterns and Oscillatory Local Field Potentials. *Neural Computation*, 6(5), 795–836.
- Usher, M., Stemmler, M., & Niebur, E. (1996). The role of lateral connections in visual cortex: dynamics and information processing. In Sirosh, J., Mikkulainen, R., & Choe, Y. (Eds.), *Lateral Interactions in the Cortex*. The UTCS Neural Networks Research Group, Austin, Texas. Electronic book, ISBN 0-9647060-0-8, <http://www.cs.utexas.edu/users/nn/web-pubs/htmlbook96>.
- Usher, M., Stemmler, M., & Olami, Z. (1995). Dynamic pattern formation leads to 1/f noise in neural populations. *Physical Review Letters*, 74(2), 326–329.
- von der Malsburg, C. (1981). The correlation theory of brain function. Tech. rep. 81-2, Max-Planck-Institute for Biophysical Chemistry, D-3400 Goettingen, Germany.
- von der Malsburg, C. (1996). Binding in models of perception and brain function. *Current Biology*, 5, 520–526.
- Wittgenstein, L. (1958). *Philosophical Investigations* (Second edition). Blackwell, Oxford.
- Wytttenbach, R. A., May, M. L., & Hoy, R. R. (1996). Categorical perception of sound frequency by crickets. *Science*, 273, 1542–1544.