Neuron, Vol. 24, 7–9, September, 1999, Copyright ©1999 by Cell Press

The Binding Problem

Review Introduction

Adina L. Roskies*
Neuron

Since its original formulation as a theoretical problem (von der Malsburg, 1981), "the binding problem" has captured the attention of researchers across many disciplines, including psychology, neuroscience, computational modeling, and even philosophy. Despite the issue's prominence in these fields, what "binding" means is rarely made explicit. In this paper, I will briefly survey the many notions of binding and will introduce some issues that will be explored more fully in the reviews that follow.

What Is Binding?

The canonical example of binding is the one suggested by Rosenblatt (1961; see also von der Malsburg, 1999 [this issue of *Neuron*]), in which one sort of visual feature, such as an object's shape, must be correctly associated with another feature, such as its location, to provide a unified representation of that object. Such explicit association, or "binding," becomes especially important when more than one visual object is present, in order to avoid incorrect combinations of features belonging to different objects, otherwise known as "illusory conjunctions" (Triesman and Schmidt, 1982). Considerable psychological evidence exists for the occurrence of illusory conjunctions (see Wolfe and Cave, 1999 [this issue]), suggesting that in certain cases, binding is indeed a problem for the brain. In addition, evidence from neuroanatomy and neurophysiology indicates that processing streams in the visual system are segregated, so that feature dimensions such as color, motion, location, and object identity are processed in separate brain regions. It should be noted, however, that while earlier characterizations of anatomical processing streams emphasized their segregation, more recent studies increasingly show a large degree of cross-talk between brain regions and more evidence for sensitivity to dimensions other than the preferred one for many brain regions (see Ghose and Maunsell, 1999 [this issue]). Thus, although neuroanatomical information will be central to understanding how the brain processes stimuli and forms representations, our current knowledge of neuroanatomy is sufficient to constrain neither the problem of binding nor its solution.

Binding Problems, Binding Solutions

The singular term "problem" suggests that binding is a unitary problem. In fact, the binding problem is a class of problems, and some of the confusion in discussions of binding may stem from the fact that different phenomena are being referred to by a single name. Besides visual binding, which includes binding information across visual space, binding information across types of features, and binding neural signals across cortical space, binding occurs in other modalities. For instance, auditory binding may be needed to discriminate the sound of a single voice in a crowd; binding across time is required for interpreting object motion; and cross-modal binding is required to associate the sound of a ball striking a bat with the visual percept of it, so that both are effortlessly perceived as being aspects of a single event. I like to refer to these sorts of problems as perceptual binding problems, since they involve unifying aspects of percepts. In addition, there are cognitive binding problems: they include relating a concept to a percept, such as linking the visual representation of an apple to all the semantic knowledge stored about it (it is edible, how it tastes, used in pies, etc.); cross-modal identification, such as being able to identify an item that has previously only been seen by how it feels; and memory reconstruction, the linking of previously encoded information to form a structured and unified representation. While this perceptual/cognitive distinction is somewhat artificial, it serves to highlight the fact that binding occurs in many different kinds of brain processes. The reviews that follow tend to focus on the visual binding problem, although not exclusively. It behooves us to be as explicit as possible about the nature of the particular problem under investigation, for while it is likely that most different forms of binding problems are solved via common mechanisms, that is only a parsimonious assumption, and ultimately it is an empirical issue.

On the other side of the coin, it is worthwhile to remember that something as complex as binding, writ large, may not have a single mechanistic solution. The potential mechanisms for binding suggested in these reviews are not mutually exclusive: there is ample room for attentional, temporal, and combinatorial coding mechanisms to work together in the brain to process information, and there may well be other as yet undiscovered mechanisms at work as well.

Is Binding a Problem?

There are two types of indications that binding is a problem. The first is an indication that the brain has difficulty in binding. This can occur in normal brains when there are temporal or capacity limitations, which lead to errors such as illusory conjunctions (see Wolfe and Cave, 1999; Treisman, 1999 [this issue]). It can also occur in damaged brains, when deficits appear that make it clear that binding is critical for normal cognitive operation. For example, in Balint's syndrome, bilateral parietal damage causes simultagnosia, the inability to perceive more than one object at a time. Other experiments show unexpected dissociations between pathways. For instance, Goodale et al. (1991, 1994) report evidence for a double dissociation between a vision-for-perception and a vision-for-action pathway: patients with lesions restricted to occipitoparietal cortex can discriminate objects according to shape, but cannot properly grasp them in accordance with the visual percept, while others with lesions in ventral cortex fail at the discrimination task,

*E-mail: aroskies@cell.com.

but can respond perfectly well to that stimulus with an appropriate grasping action (Goodale et al., 1991, 1994). In this case, it is not entirely clear whether binding normally takes place between these pathways, but the fact that similar dissociations are not seen in undamaged brains suggests that information is shared between pathways, and this may be accomplished through binding. Thus, both experiments and case studies reveal that the brain does engage in some operations responsible for appropriately combining neural signals.

The second indication that binding is a problem is a conceptual one: even if the brain usually does not appear to have a problem in correctly binding signals, we as scientists still lack an understanding of how information variously distributed in patterns of neural firing results in coherent representations. Thus, binding is a problem in that it requires an explanation.

Mechanisms of Binding

Several of the papers in this issue of *Neuron* treat, in detail, possible mechanisms of binding. Reynolds and Desimone (1999) explain how attentional mechanisms could be indispensible for binding, and indeed much evidence from psychology and neurology suggests that attention and binding are intimately related (Wolfe and Cave, 1999). Both Gray (1999) and Singer (1999b) discuss how temporal synchrony (or temporal correlation) could be used to bind neural signals and describe the evidence for and against temporal binding. They bring to the fore an impressive amount of evidence from diverse neuroscientific approaches, and provide a perspective on temporal aspects of neural coding from the cellular to the systems levels. From the data presented, I think it is clear that neural oscillations and synchronous signals are present in the brain and that neurons have the machinery to exploit these signals (for instance, in changing synaptic strengths). However, there is only suggestive but not yet incontrovertible evidence that these signals are used for binding, or that they play any critical role in brain function. These proponents of temporal correlation are extremely measured in their discussion, and are careful to emphasize that available evidence is consistent with, but does not prove, that temporal correlation is used for binding. In the other camp, Shadlen and Movshon (1999) offer an incisive critique of the temporal binding hypothesis, arguing that temporal signals do not have the capability of encoding binding relations. Several of the reviews suggest experiments that could help cement or destroy the case for temporal synchrony in perceptual binding, although executing these experiments properly will be a great challenge.

Ghose and Maunsell (1999) and Riesenhuber and Poggio (1999a), in contrast to the others, argue that there is really no binding problem, in the sense that the classical hierarchical scheme of neural coding, going back to Hubel and Weisel, can encode the combinatorial information required for perception without leading to a combinatorial explosion. While Ghose and Maunsell claim that specialized groups of cells are at least as powerful as temporal correlations could be for supporting rich representations, they nonetheless suggest that an explanation for unitary representations is lacking and that some form of binding must exist. Riesenhuber and Poggio do not address this larger issue, but they outline a model that stands as a plausibility proof for combinatorial coding being sufficient for object recognition. Their model uses a moderate number of broadly tuned neurons to perform an object recognition task, and it is successful even in the presence of other objects. Although their model is consistent with neurophysiological data, it remains to be seen whether an extended model could deal with the diversity of stimuli and complexity of tasks that brains handle easily every day. In addition, the model's metric for object recognition is that a tuned unit responds more to the object it is tuned for than to distractors. It will be interesting to see how robust the performance is when the number of units and the number of distractors grows.

Von der Malsburg (1999), in his review, makes two particularly important points that add perspective to the discussion. The first is that the flexibility of brain function sets it apart from network models, and the second is that learning or plasticity must be a fundamental component of whatever the mechanisms are that bind neural signals. While the goal of a model is often to simulate a single task that the brain can perform, the "goal" of the brain is to be able to handle whatever complex stimuli or situations it might encounter. Rosenblatt's (1961) solution to the binding problem demonstrates that solving any given problem is often rather straightforward, and complex mechanisms such as temporal binding can appear to be unnecessary. However, difficulties such as combinatorial explosions (but see Ghose and Maunsell, 1999 [this issue]) and ontogenetic factors arise when trying to model a more general-purpose device, and the necessity of binding becomes apparent. How general is Riesenhuber and Poggio's model? Can it also be the framework for other visual tasks besides recognition? How flexible is this type of network? Until issues of scaling and task specificity are more thoroughly explored, it is difficult to say how applicable such a model may be for understanding brain function.

Is Binding Enough?

Both Ghose and Maunsell and Shadlen and Movshon point out what is perhaps the fundamental problem with the binding hypothesis: even supposing that temporal coding is the vehicle for signaling which neural populations should be bound together, the theory does not adequately address how those combinations are computed. In a sense, the binding problem is pushed back one level, for it must be solved at least partially for the necessary temporal correlations to be established. Are neurons with appropriate receptive field properties and anatomical connectivity a sufficient basis for the generation of correlated signals? What is the role of top-down connections in establishing proper synchrony? These questions are difficult to answer and difficult to model with biological realism. A similar problem rests with the output side of the temporal binding hypothesis—if correlations are the signal for binding, how are those signals read out? The readout problem is one of the most puzzling and fundamental problems for systems neuroscience in general: how is the firing of populations of neurons interpreted and transformed by other neurons to result in decision, action, perception, etc? This problem plagues most, if not all, models of brain function, for in modeling it is the modeler that attributes semantics to nodes in a network. One way to potentially avoid that problem would be to close the loop between the model and the world, so that the world impinges upon the network, which then acts upon and affects the world, thus generating semantics through action.

This brings me to mention (by virtue of my editorial license) what is perhaps the most mystifying binding problem of all: the problem of consciousness. How does something as simple and mechanistic as neural firing add up to subjectivity, raw feelings, a self? Are the mechanisms that allow us to attribute the correct color and shape to an object the same ones that lead to the unity of phenomenal experience? Will the solution of the binding problem be the solution to the mystery of consciousness? I will not belabor the point, since answers will be long in coming, but although none of the scientists who authored the reviews that follow discuss binding with respect to consciousness, I will wager that the a good part of the interest, excitement, and contentiousness that surrounds the binding debate is attributable to the magnitude of the issues with which it is connected. The following pieces provide a comprehensive review of the status of the binding problem at the dawn of the new millenium. It will be extremely interesting to repeat this exercise in a decade or two, to chart the progress that we as scientists make on one of the most puzzling and fascinating issues that the brain and cognitive sciences have ever faced.

References

A comprehensive reference list for all reviews can be found on pages 111–125.