

The Binding Problem(s)

8/25/2010 9:38 AM

Jerome Feldman

Abstract

The neural binding problem (NBP) encompasses three distinct situations: feature binding, variable binding, and the subjective unity of perception. Feature binding involves associating the correct (visual) features with objects and is fairly well understood. Variable binding arises in natural language and other abstract thought; there are ingenious proposed models, but no experimental confirmation. The subjective unity of perception is the deepest variant of the NBP, but unfortunately, contemporary science has nothing to say about it.

In its most general form, the binding problem concerns how items that are encoded by distinct brain circuits can be combined for perception and action. Any coherent distributed system needs a way of assimilating information so, at a basic level, the binding problem is unavoidable. We start by considering the abstract computational problem and coordinated action in social systems as well as the traditional neural binding problem (NBP).

Any large parallel system will have a lot of information that cannot be fully accessible at every node and must be abstracted. The brain, with its millions of retinal cells, is one example, but the problem is inherent. Any such system should ideally make decisions/actions based on all available information, but this is combinatorially impossible – the system architecture needs to privilege certain combinations. The brain has the additional constraint that almost all connections are local. The brain's organizing principle is topographic feature maps¹ and in the visual system these maps are spatial².

The purpose of combining information is to make good decisions and actions. Consider the analogy of a large human organization, such as a company or government agency. A prototypical company executes discrete actions including establishing facilities, acquiring materials, developing and marketing products, buying politicians, etc. Some government agencies also do things. The capabilities for all these activities are distributed (as in the brain) without any individual or small group having complete understanding and yet the organization takes unified actions. Looking ahead, this shows that coherent behavior does not require the unified perception that we subjectively experience.

A real-world analog of the binding problem is “connecting the dots” in intelligence operations. Famously, on Christmas of 2009, a young Nigerian was able to board

multiple flights and arrive over Detroit with pounds of explosive on his body, despite several warning signs. Even public accounts of this intelligence catastrophe present a striking illustration of the binding problem, identifying fourteen separate information integration failures.

The neural binding problem (NBP) encompasses three distinct situations: feature binding, variable binding, and the subjective unity of perception, defined and discussed below. These are “problems” because we know that the brain has many distinct specialized circuits, and don’t know how these myriad computations are combined for perception, thinking, and action.

It is important to recognize that the brain is a neural system that evolved to run a physical body in a social environment. It is constantly trying to find a best fit between the agent’s goals and noisy perceptual input and is subject to all manner of illusions. Current research has largely abandoned the notion of an isolated NBP and studies binding as part of overall brain function. Following tradition, we will focus mainly on the visual system.

The Unity of Perception

We will start with the deepest and most interesting variant of the NBP, the subjective unity of perception. This is closely related to the problem known as the Illusion of a stable visual world³. Traditionally, the NBP concerns instantaneous perception and does not consider integration over saccades. But in both cases the hard problem is explaining why we experience the world the way we do. There is now overwhelming biological and behavioral evidence that there is no high resolution, full field, visual representation in the brain, but that is what we subjectively experience. As is well known, current science has nothing to say about subjective (phenomenal) experience and this discrepancy between science and experience is also called the “explanatory gap”.

Feature Binding

Fortunately, quite a lot is known about feature binding, the simplest form of the NBP. There has been much more work on feature binding than on variable binding, which will be discussed later. The basic question is ancient – why don’t we confuse e.g., a red circle and a blue square with a blue circle and a red square. There is an extensive continuing literature of feature binding. Triesman⁴ is an excellent survey of the early literature and Velik⁵ is a more recent review with a good historical perspective.

While linking features to the correct object and location is a *requirement* for effective vision it is not normally a *problem*. The visual system is spatiotopically organized and most detailed vision is done in foveal fixations which are coordinated in space and time. Also, visual features are routinely bound by gestalt principles, linked to known object

images, etc. There is a question of how these coherent feature bundles are *remembered*, and we will discuss this below.

In fact, a more basic challenge in vision might be called the *unbinding* problem. An individual photoreceptor cell has no way to distinguish a change in illumination from a reflectance change or self-motion from target motion, but the agent relies on such distinctions. The vast expansion of visual cells from about one million in the optic nerve to billions in visual cortex is generally understood to carry out the transform from conflated proximal signals to estimates of the features of their distal sources⁶. Computational theories of this unbinding go back to Zipser and Andersen⁷ and are important in current work.

Another salient fact is that the visual system can perform complex recognition rapidly enough to preclude anything but a strict feed-forward computation. There are now detailed computational models⁸ that learn to solve difficult vision tasks and are consistent with much that is known about the hierarchical nature of the human visual system. The ventral (“what”) pathway contains neurons of increasing stimulus complexity and concomitantly larger receptive fields and the models do as well.

Attention and Feature Binding

In fact, much of the binding of visual features (shape, size, color, texture, motion, etc.) is done well only in foveal vision⁹. There are about three fixations per second and, during a fixation, there is usually a single item of focal interest and so the binding of features is easy. All of the foveated features are local in time and space and thus bound together. In addition, we have known for decades that effective attention can also be covert, without saccades, and this will be discussed below.

Essentially all the experimental results on illusions in feature binding arise from overloading the system in one way or another. Some example manipulations include brief presentations, masking, and binocular rivalry. Stressful cases can disrupt the normal feature binding mechanisms. There is some inconsistency in the terminology: “binding” is sometimes used to refer only to the problematic, attention-demanding cases. We will discuss some basic results on this *stressed* feature binding and then return to the role of attention.

One of the most striking examples of a problem in stressed neural feature binding is one of the earliest results. Given a brief presentation of several randomly oriented letters ‘P’ and ‘Q’, people will see an illusory ‘R’ about 10% of the time. This is assumed to be caused by combining the ‘\’ tail of the ‘Q’ with a ‘P’ to activate the perception of an R. The display contains good evidence for all the features of an ‘R’ so the result is not too surprising. Another basic set of results concerns “pop-out” phenomena. When a target

(e.g. red horizontal line) shares only one feature with background distractors, it is easily detected, but when it shares multiple features with some distractors, detection requires serial attention. Reynolds and Desimone⁹ survey the early results on attention in feature binding. Attention is also proposed as the key to combining the computations of the ventral (what) and dorsal (where) streams of the visual system. Chikkerur et al. have built a detailed biologically grounded Bayesian model, using both spatial and feature attention, and compared it with a wide range of experimental data¹⁰.

There is a good deal of ongoing research involving the NBP, utilizing a wide range of experimental techniques. Seymour et al.¹¹ used functional magnetic resonance imaging (fMRI) to study which human brain areas are most active when color and motion are perceived separately and which areas respond most to conjoined stimuli. They presented circular displays with red dots rotating one direction and green dots in the opposite one. A sophisticated pattern classifier was able to distinguish which features were coupled from the fMRI signature. Interestingly, some voxels as early as primary visual cortex V1 had detectably different responses to the different pairings. Whitney¹² provides a nice summary of these and related results and their implications for the NBP.

Obviously enough, we must combine all the relevant visual features of an object in order to recognize it. Partly because of the subjective unity of perception, it is natural to imagine and model this process as entailing some unified representation of all the features. But multiple smaller combinations would also do the job and there are a number of reasons to believe that this is more likely in the brain. Humphreys¹³ reviews the clinical literature suggesting multiple binding processes.

As discussed above, the neurons in early and intermediate vision are sensitive to multiple, but not all, stimulus dimensions¹. Morita et al.¹⁴ describe a number of binocular rivalry binding experiments and models suggesting that pair-based feature coding is important in vision. Their basic display varied three features across the two eyes: shape (flower, snowflake), color (red, green) and rotation (clockwise, counter-clockwise). When all three contrasting features were presented simultaneously, there was strong rivalry, leading to alternating perceptions. When only a single contrasting attribute at a time was used, subjects reported indistinct or missing objects much more often. Crucially, paired attributes led to behavior like the simultaneous three attribute case.

Further experiments explored illusory conjunctions when three 2-attribute images were displayed rapidly (94 ms. apart). If the images were the same for both eyes and there was no rivalry, most subjects saw three consecutive objects and the illusion involving all three attributes conjoined was seldom perceived. In the rivalry condition, subjects saw the full 3-attribute image about half the time and were not aware that it was illusory. The paper also includes a related short-term visual memory experiment and a

discussion of simple computational models involving mutual inhibition at the level of paired-attribute units. More generally, theories of feature binding inherently entail some model of visual memory.

Short Term Feature Memory

Feature binding would be of no use if it were evanescent; there must be some kind of short term memory of the feature bundle. Traditionally, the feature NBP refers only to short time periods. There is a literature on more general questions of binding and memory and Zimmer et al.¹⁵ is a good place to start. I will say more about binding in models of *episodic* memory in the section on variable binding.

As discussed above, most detailed feature binding happens in foveal vision. Since there are about three saccades per second, combining and retaining this information is a considerable challenge. For one thing, visual short term memory is now known to be much less stable than it subjectively seems. A wide range of results on “change blindness” shows that unattended items in short term memory are malleable in many ways. Our exquisitely sensitive change detection circuitry is inhibited during saccades and this is exploited in tests of short term memory. The Simons and Rensink survey¹⁶ covers a range of recent results on change blindness with possible implications for feature binding.

As we mentioned, the visual system is spatiotopically organized and spatial coherence provides the core of feature binding – things appearing in the same place tend to group. This was called “spatial tagging” of feature bundles in the early literature⁴. It is not so simple because both our eyes and objects move and there are several distinct kinds of spatial map – retinotopic, egocentric, allocentric, etc. There are additional multi-sensory maps that deal with binding across modalities and with connections to motor circuits. Still, to first order, short term memory for feature binding is spatial. Additional evidence for this position comes from fMRI experiments like those of Shafritz et al.¹⁷. They found that parietal cortex, central to spatial processing, was highly activated when stressed binding could be aided by spatial cues, but not when the cues were temporal.

There is also a body of work directly examining short term memory for feature binding. For example, Karlsen et al.¹⁸ studied when feature binding was automatic and when it seemed to require attention. Their displays had three or four simple colored objects, like a circle, a triangle or a cross. In the base case the objects were colored and subjects had no trouble recognizing if a target probe was the same as one in the display. The main manipulation was to separate the color from the figure, either by a short time interval or a by small vertical displacement on the screen. Each of these changes caused a significant, but modest decrement in accuracy. In addition, they examined whether a concurrent task of counting backwards would interfere in any or all of the

conditions. To their surprise, the interference was greater in the case of unified presentation and the paper discusses possible explanations for this result.

In another experiment, Bouvier and Treisman¹⁹ showed that top-down feedback appears to be necessary for stressed feature binding. In this experiment there were six “+” signs each of which had one colored bar, either horizontal or vertical. The task was to report the color and orientation of the colored bar in the target figure, which was indicated by four small dots surrounding the “+”. The key manipulation was sometimes having the four dots persist for 300 ms after the target “+” disappeared. This “trailing mask” is known to disrupt feedback, but not forward processing. As predicted, a trailing mask led to a significant reduction in orientation accuracy when subjects had to bind the color and orientation of one bar of the “+”, but not for a single bar. Interestingly, this effect was minimal when the location of the target was known in advance, enabling covert attention to the target area.

Temporal Phase Binding

Through the decades before the spatial character of visual memory was fully understood, several computational models of feature binding were proposed^{20, 21}. The most interesting and well studied proposed binding mechanism is *temporal phase synchrony*. Timing considerations, like spatial organization, are fundamental to neural processing at all levels²². Things that occur at or near the same time are treated differently in perception, action, memory, and learning. Temporal phase synchrony proposals go beyond this truism and suggest a powerful additional mechanism.

One of the basic features of the brain’s architecture is that everything is connected and potentially active at the same time. An elegant idea, dating at least to the 1960s⁵ involves dividing local firing patterns into separate phases, like time-domain multiplexing in engineering. This would, in principle, allow several distinct sets of non-interfering bindings.

The phase binding approach breaks the cycle of neural firing into discrete time slices. When an attribute node fires in-phase with an object node, this coincidence represents a binding between them. The best-known model of this sort is SHRUTI²⁰, and its mechanisms have been carefully examined from several perspectives. Figure 1 below shows an example of temporal phase binding, where time moves along the X-axis. We will consider this example in more detail later - for now just look at the bottom six rows. Notice that the triangles (denoting spike trains) in row 1 remain aligned with those in row 5 and similarly for rows 3 and 6. In this example there are only these two phases and each captures a binding: (Tom:? with owns:x) and (Book:? with owns:y). The trapezoids in line 7 depict the envelope of each phase cycle. Similar mechanisms could, in principle, be used to bind visual features like motion and color to shape.

Biological considerations

As in any neural modeling context, there are two distinct criteria for models and theories of the NBP: computational adequacy and neural plausibility. Temporal phase binding has been studied intensively in both dimensions. As was mentioned, no one questions that time in general and synchrony in particular is central to neural computation. But the particular mechanism of phase synchronization remains contentious.

There have been detailed modeling studies and simulations supporting the idea that stable phase synchronization by systems with neurally plausible properties. There have also been experimental findings consistent with phase binding, but these have been hard to replicate. Shadlen and Movshon²³ present a detailed analysis that questions the plausibility of the idea and supports a more structural and spatial model. Temporal phase coherence is no longer considered a major contender in feature binding, in part because it would be much too slow to account for experimental data. It is much more relevant in variable binding where most other techniques don't apply.

Variable binding

Neural realization of variable binding is completely unsolved, but is not unsolvable^{20, 21}. All animals need feature binding, but variable binding only arises in language and other symbolic thought. As a simple example of variable binding, consider the sentence: "Mary bought a book and gave it to John". We can draw several inferences from this statement, using rules that can be easily stated in logic, such as:

$\text{owns}(z,y) \text{ and } \text{gives}(z,x,y) \Rightarrow \text{owns}(x,y)$ OR $\text{buys}(x,y) \Rightarrow \text{owns}(x,y)$.

It is straightforward to implement such rules in math or programming, but we don't know how the brain does it. Because variable binding is a characteristic of language and abstract reasoning, there are no animal models or experiments available

In conventional computing, we assume that different program modules all have access to the values of (global) variables and can modify their behavior appropriately. Any theory of neural computation needs some mechanism for achieving this kind of global effect. There are three variables (z,x,y) in the rules above and they can be bound to a very wide range of possible fillers – no fixed neural network could capture all the possibilities. And, of course, such inferences chain and can get quite complex. While feature binding is a problem only in stressed situations, neural binding of variables is a challenge in all cases. Variable binding and the related function called unification are ubiquitous in any theory of language understanding. In our example, "Mary bought a book and gave it to John"., 'Mary' fills (is bound to) the agent role of both 'bought' and 'gave', 'John' fills the recipient role and 'book' fills the theme or indirect object role of 'gave'. Every sentence involves this kind of variable binding and

there is no experimental evidence on how the brain does this. The linguist Ray Jackendoff²⁴ has suggested that the variable binding problem is the key to any neural theory of language.

A recent article by van der Velde and de Kamps in the Behavioral and Brain Sciences²⁵ and the accompanying commentary explore a wide range of connectionist approaches to the binding problem. The most basic model is brute force enumeration of all possible variable bindings, sometimes with coarse-coded conjunctive binding to mitigate its exponential complexity. More recently, van der Velde and de Kamps²⁵ employ such a crossbar network in their Neural Blackboard model. In this design, rather than synchrony or passing around some sort of signature, there are connections between computational nodes that are ordinarily disabled, but may be enabled and when enabled allow signals to travel between the two nodes for a period of time. Thus, it attempts to solve the binding problem by making temporary links between nodes.

None of these methods work for the general case where new entities and relations can be dynamically added, as is common in language. The essential difference is that there are a potentially unbounded number of items that might be bound to a variable, so none of the pair-coding or crossbar techniques described above will work. For example, if I tell you that my granddaughter Sonnet is brilliant, you have a new person to consider as a possible filler for variable roles and also a number of new facts for use in inference.

A second approach has been to use sign (signature) propagation. In sign propagation, each variable in an expression has its own node (a group of neurons working together). This node can represent and transmit a particular signature corresponding to a concept, so the signature is essentially treated as a name for the concept²¹. The main difficulty is that then there must be one signature for every representable object – so each signature must carry about 20 bits of information and a new signature must be created for each new item encountered. There is no biologically plausible suggestion for how the brain might do this.

A third, and the most widespread approach, is that of phase synchronization, also known as temporal synchrony –as described in the previous section. In fact, early exploration of phase synchronization was the origin of the term binding in neuroscience²⁶. As Malsburg (personal communication) writes: "When I coined the term binding in the neuroscience context, the source of the expression was exactly that -- variable binding, as in computer science"

This approach breaks the cycle of neural firing into discrete time slices. When a variable node fires in-phase with a concept node, this coincidence represents a binding between

them. The best-known model of this sort is SHRUTI^{20, 27}, and its mechanisms have been carefully examined. Its biological basis was discussed in the previous section - here we will explore its representational capabilities.

Figure 2 depicts a fragment of a Shruti²⁰ inference network, which uses temporal phase binding. The network on the left represents the simple predicate calculus rule:

$\text{buys}(x,y) \Rightarrow \text{owns}(x,y)$.

As you would expect, the circular nodes labeled x and y , represent the variables in our rule. The upward pointing house shapes labeled $?$ are used to propagate queries like: “Does Tom own a book?”. The downward pointing house shapes convey positive or negative answers for a query if the information is available in the network.

Let’s consider a Boolean query on whether Tom owns a book. The network on the right of Figure 2 shows that this version of Shruti can support universal and existential quantifiers and also a simple ontological hierarchy with Tom as a person and therefore a possible agent. There would also be a network of facts, perhaps one that Tom bought the book “Ulysses”. Our logical rules are (as usual) universal and the query is existential – is there a book that Tom owns?

This would be posed in Shruti by first assigning separate clock phases to the pairings (Tom:? x with owns: x) and (Book:? y with owns: y) as depicted at the bottom of Figure 1. Then the $?$ node on the bottom left of Figure 2 would be activated. Figure 2 shows one way (purchase) that an agent could come to own something, but there are others and they would also be linked to the ‘owns’ relation. Each such causal rule involves a mediator circuit as shown on the middle left of Figure 2; the mediator assures that only facts that have all the specified bindings can participate in inference. Now, spreading activation on the upward $?$ node path effectively searches for a fact in memory having the proper bindings (phases) . If there is such a fact e.g., “Tom bought Ulysses”, it will activate the $+$ node of the antecedent clause on the upper left of Figure 2. This process is depicted in the top five lines of Figure 1; after two cycles, a buy relation involving Tom and a book is queried and (after two more cycles) is activated. Because of the downward connections in Figure 2, this in turn will activate the $+$ node of the consequent “owns” relation – yielding a positive answer to the query. The system can also (temporally) bind the variable y to Ulysses, yielding a sharper answer – yes, Tom owns Ulysses.

Various Shruti implementations²⁷ have extended these capabilities to handle queries with multiple variables, quantifiers, probabilistic relations, etc. The structure of the networks for these impressive inferential capabilities does not depend crucially on phase binding and can be largely applied to other proposed variable binding mechanisms.

As discussed in the previous section, decades of research have yet to find convincing evidence for temporal phase binding and people continue to explore possible alternatives. One recent effort²⁸ attempts to combine many of the ideas of Shruti with a basic variable binding mechanism closer to the signature method. Like Shruti they begin with the fact that people can only deal with a small number (~7) of bindings at a time. In this case, a signature passing system could get by with ~3 bits, which is plausible. The basic idea of replacing eight time slices with a 3 bit signature is straightforward and was mentioned in a 1993 paper²⁰.

To link the short signatures to the (still) large number of possible concepts, Barrett et al.²⁸ uses a central structure that controls binding. This also enables some operations that SHRUTI cannot perform. First, it permits the network to keep track of specific bindings, where otherwise they would be lost as a time slice or signature spreads through the network. The central binding structure also allows for more complex abilities, such as the unification of signatures that have been determined to represent the same object. Furthermore, a central binder allows conflict-free signature allocation, which cannot be performed without global information. The Barrett et al. paper²⁸ also considers a number of auxiliary tasks that would be needed in a neural system for variable binding and inference. There is currently no evidence for or against binding with short signatures and no experiments have been proposed.

Episodic memory models

In addition to short-term memory, there are at least three distinct modes of long-term memory: semantic memory (like Figure 2), procedural memory (like dancing) and episodic memory (like your first date). There continues to be active computational and neural modeling of all these functions, but only episodic memory overlaps with the binding problem. The reason is obvious – a memory of a particular situation or episode entails binding together many particulars of time, place, players, results, etc.

There is an extensive history of computational models of episodic memory and of the hippocampal complex which is known to be heavily involved. Shastri²⁹ has a comprehensive review of these efforts and also the most detailed proposal for linking computational models to experimental findings. The core of Shastri's proposal involves temporal phase binding, quite like that described in Figures 1 and 2 above. The model extends variable binding to add circuits for memory retrieval and for binding-error detection and detection and integration. All of this functionality is modeled and simulated by elaborate circuitry postulated to reside in various divisions of the hippocampal complex.

An alternative formulation by Cer et al.³⁰ describes a complex computational memory model with three separate binding mechanisms for the hippocampus, posterior cortex

and prefrontal cortex, none of which involves temporal phase binding. There is ongoing work on a wide range of neuro-computational memory models, but only a few that explicitly consider the binding problem¹⁵.

Conclusions

The three versions of the binding problem persist in their separation. The question of subjective experience continues to be intractable. Feature binding is not a problem under normal conditions although much is still to be learned. Feature binding under stress remains a fruitful source of perception and memory experiments.

Neural realization of variable binding is completely unsolved, but is not unsolvable. Because variable binding is a characteristic of language and abstract reasoning, there are no animal models or experiments available. All of the proposed computational theories are quite complex and none have experimental support.

It appears that the time has passed when significant advances in cognitive science can be achieved by considering isolated “problems” and “solutions”. The issue of how the brain combines information from different circuits is important, but only as one capability of a system that has many performance requirements.

Related Articles

Cogsci-079: Language and Brain

Cogsci-083: Language, Connectionist and Symbolic Representations of

Cogsci-168: Perception, Philosophical issues about

References

1. Kaas J. Topographic Maps are Fundamental to Sensory Processing Brain Research Bulletin, 1997, 44:107–112
2. Lennie P. Single units and visual cortical organization. Perception, 1998, 27: 889 - 935
3. Martinez-Conde S., Krauzlis R., Miller, J, Morron C, Williams, D., & Kowler, E.. Eye movements and the perception of a clear and stable visual world. *Journal of Vision*, 2008, 8(14):i, 1, doi:10.1167/8.14.i.
4. Treisman A. Solutions to the binding problem: Progress through controversy and convergence. Neuron, 1999, 24:105-125

5. Velik R. From single neuron-firing to consciousness--towards the true solution of the binding problem, *Neurosci Biobehav Rev.*, 2010, 34(7):993-1001
6. Barlow HB. Why have multiple cortical areas? *Vision Research* 1986, 26: 81-90
7. Zipser D, Andersen, RA. A back-propagation programmed network that simulates response properties of a subset of posterior parietal neurons. *Nature* **331**, 679 - 684 (25 February 1988); doi:10.1038/331679a0
8. Serre T, Oliva A, Poggio T. A feedforward architecture accounts for rapid categorization. *Proc Natl Acad Sci USA*, 2007, 104(15):6424–6429
9. Reynolds JH, Desimone R, The role of neural mechanisms of attention in solving the binding problem. *Neuron*,1999, 24(1):19-29, 111-25.
10. Chikkerur SS et al. What and where: A Bayesian inference theory of attention. *Vision Research*, 2010, doi:10.1016/j.visres.2010.05.13
11. Seymour K, Clifford CW, Logothetis NK, Bartels A, The coding of colour, motion and their conjunction in human visual cortex. *Curr. Biol.*, 2009, 19:177–183.
12. Whitney D, Neuroscience: Toward Unbinding the Binding Problem *Current Biology* 1990, 19(6):251-253
13. Humphreys GW, Conscious visual representations built from multiple binding processes: Evidence from neuropsychology. *Progress in Brain Research*, 2003, 142:243-255.
14. Morita M, Morokami S, Morita H Attribute Pair-Based Visual Recognition and Memory. *PLoS ONE*, 2010, 5(3): e9571. doi:10.1371/journal.pone.0009571
15. Zimmer HD., Mecklinger A., Lindenberger U. (Eds.), (2006) Handbook of binding and memory: Perspectives from cognitive neuroscience. Oxford U. Press
16. Simons, DJ and Rensink RA. Change blindness: past, present, and future *TRENDS in Cognitive Sciences* 2005,9(1) :16-20
17. Shafritz KM, Gore JC, Marois R, The role of the parietal cortex in visual feature binding. *Proc Natl Acad Sci U S A*, 2002, 99(16):10917-22.
18. Karlsen PJ, Allen RJ, Baddeley AD, and Hitch GJ Binding across space and time in visual working memory *Memory & Cognition*, 2010, 38 (3), 292-303
19. Bouvier, S. and Treisman. A. (2010) Visual Feature Binding Requires Reentry *Psychological Science* published online 8 January 2010

20. Shastri L, Ajjanagadde V. From simple associations to systematic reasoning Behavioral and Brain Sciences 1993, 16:417--494,.
21. Browne A, Sun R. Connectionist variable binding. Springer Verlag, Heidelberg, 2000
22. Feldman JA (2010) Ecological expected utility and the mythical neural code. *Cogn Neurodyn*. 2010 March; 4(1): 25–35.
23. Shadlen MN, Movshon JA. Synchrony unbound: a critical evaluation of the temporal binding hypothesis. *Neuron* 1999, 24:67–77
24. Jackendoff R. Foundations of Language. Oxford University Press, Oxford, 2002.
25. van der Velde F, de Kamps M. Neural blackboard architectures of combinatorial structures in cognition. Behavioral and Brain Sciences 2006, 29:37-70
26. von der Malsburg C. The correlation theory of brain function, MPI Biophysical Chemistry ,1981, Report 81-2
27. Wendelken C, Shastri L, Multiple instantiation and rule mediation in SHRUTI Connection Science, 2004, 16:211–217
28. Barrett L, Feldman JA, Mac Dermed L. A (somewhat) new solution to the binding problem. *Neural Computation*, 20: 2361-2378
29. Shastri L. Episodic memory and cortico-hippocampal interactions, Trends in Cognitive Sciences, 2002, 6: 162-168. 2002
30. Cer DM, O'Reilly RC. Neural Mechanisms of Binding in the Hippocampus and Neocortex: Insights from Computational Models. In Zimmer HD., Mecklinger A., Lindenberger U. (Eds.), (2006) Handbook of binding and memory: Perspectives from cognitive neuroscience. Oxford U. Press