Synchronous vs. Conjunctive Binding: A False Dichotomy?

Robert F. Hadley (hadley@sfu.ca)
School of Computing Science, Simon Fraser University
Burnaby, B.C. V5A 1S6, Canada

Abstract

It is argued that claims about the superiority of synchronous binding methods are premature. A hypothetical explanation for synchronous firing is offered, which portrays such firing as a symptom of a deeper binding process that employs conjunctive nodes. It is shown that only a feasibly small number of conjunctive binding nodes would be required in the deeper process.

Introduction

A long standing issue in connectionist research has concerned the manner in which the human brain contrives to *bind* values to variables, fillers to roles, and (sensory) features to features. Closely related to this issue is the quest for an elegant and efficient solution to "variable binding" in artificial connectionist networks. In the present paper, I shall be concerned ultimately with biological neural networks, and for this reason the discussion will focus upon the two *binding methods* whose biological bases have been most directly confirmed. In brief, these two methods are: (a) *conjunctive coding*, where the binding of two elements is signaled by the cyclical firing of a particular node (realized by one or more neurons) that receives input from other nodes that represent the elements which are to be bound together (or "conjoined"). (b) *synchronous firing*, where distinct nodes (or distinct sets of neurons) are represented as *being bound together* when the distinct nodes fire in approximately synchronous cycles.[1]

The concepts of conjunctive coding and binding via synchronous firing will be elaborated in following sections, but for the present some observations may be helpful. The use of *conjunctive coding* as a method for binding values to variables and fillers to roles can be observed in several early connectionist models (cf. Rumelhart, McClelland, et al, 1986) and has been especially valuable in network models of language parsing and/or acquisition (see Cottrell, 1985, Stevenson, 1994, Hadley and Cardei, 1999, Hadley and Hayward, 1997). Moreover, the binding of sensory features via explicit, conjunctive binding neurons has been empirically verified in studies of the visual cortex of mammals and other cortical regions (Singer, 2003). Such binding of elements (whether they be features, concepts, variables, or roles) is known to be both reliable and fast. However, binding via conjunctive coding has often been criticized as being combinatorially explosive in the number of nodes (or neurons) required to effect all the bindings possible in a given domain. Unfortunately, this criticism has at times been made in a sweeping form, without regard to the fact that the number of conjunctive nodes required may merely be the simple product of the number of fillers and roles involved (the latter being feasibly sparse in non-perceptual domains).

In contrast with conjunctive coding, synchronous firing, considered as a binding method, has in recent years attracted considerable notice, both because of its biological basis (cf. Singer, 1999, 2003) and because it (purportedly) avoids the proliferation of nodes attributed to conjunctive coding. Evidence for the role of synchronously firing neurons in sensory feature binding is substantial (Singer, 1999), and synchronous firing has been used as a basis for variable binding in some models of inference, including SHRUTI (Shastri and Ajjanagadde, 1993) and LISA (Hummel and Holyoak, 1997).

Despite the considerable positive acclaim that synchrony has attracted as a solution to the binding problem, its touted advantages over conjunctive coding might, for all we actually know, be illusory. I say this because synchronous neural firing cannot be regarded as a fundamental, unanalysable phenomenon. It is well known that the firing of an individual neuron is caused by underlying processes, both chemical and electrical. For this reason, among others, it would be scientifically incoherent to regard the *systematic* firing of neurons in synchrony as a fundamental, inexplicable phenomenon. We cannot both claim that the firing of an individual neuron has deeper causes and that the synchronous firing of a set of neurons must be regarded as absolutely basic and inexplicable. Yet, the causes of synchronous firing in non-perceptual realms (which include language comprehension, reasoning, and other forms of higher cognition) are not presently known. The few causal indications that do exist are compatible with the following supposition: apart from spurious coincidences, neurons fire in synchrony only when they are each connected to explicit conjunctive binding nodes (dub them CBN) whose cyclical firing engenders synchrony in the firing cycles of neurons that the CBNs are immediately connected to. The simplest form of this supposition would require one conjunctive binding node for each pair of synchronously firing nodes.

I have no wish to defend the simple form of this supposition, but, as will soon emerge, I will propose a variant conjecture which requires only a feasibly small number of conjunctive nodes. Regardless of that, we should bear in mind that until the causes of synchronous firing are understood, we are not in a position to judge the number of neurons involved in the genesis of synchrony firing. In consequence, we are not entitled to flaunt the putative greater efficiency of synchronous binding schemes.

In section 2, the *genesis problem* for synchronous firing is examined. There, I present a hypothetical model which might, in principle, explain how the relevant kind of synchronous firing is ever engendered. This model posits the mediation of a feasibly small number of conjunctive coding nodes which lie at the core of a network which ultimately links with neurons that are caused to fire synchronously. Although this proposed model is unlikely to be accurate in its details, it is hopefully *indicative of the direction* in which a true explanation of the causes of synchronous firing is to

---

[1] Other proposed solutions to the binding problem include Lange and Dyer (1989); and CONSYDERR, by Sun (1992). Each of these solutions possesses merit and features of interest.

be found, *at least in the domain of filler-role binding*. If something akin to this hypothesized model is correct, then the usual criticisms of conjunctive binding would not be germane to the phenomena of value-variable binding and filler-role binding, both of which occur in reasoning, language comprehension, and other cognitive domains. (These last two forms of binding are strongly analogous, and hereafter I shall, for brevity, simply refer to them as "filler-role" binding.)

Genesis of Synchronous Binding

In this section I present a hypothetical model for how two separate nodes, each of which might be realized by a distinct set of several neurons, could be caused to fire in synchrony. The model posits an underlying *core network*, in which CBNs (conjunctive binding nodes) play a key role in the process that engenders synchronous firing. The scope of this model is restricted to instances of filler-role binding. In particular, it is assumed that a single "filler node" (selected from a finite set of candidate fillers) should have the capacity to bind with a single "role node" (likewise selected from a finite set of candidate roles) on a given occasion. It is also assumed that, within the *filler-role domain* of applicability, the most active filler node should bind with the most active role node, provided the two nodes are both "currently unbound" (that is, they are both free to bind). [2] Examples of connectionist systems which employ this strategy (of binding an unbound filler node to an unbound role node whenever both are the most active nodes possessing appropriate connectivity) include the language learning and parsing system presented in Hadley and Cardei (1999), and the language parsing system of Stevenson (1994).

Here, I will be assuming that both *role nodes* and *filler nodes* are particular types of *concept representations*. For example, it seems reasonable to suppose that our understanding of the *role* of "being the agent of some action" involves having the concept of "being an agent". Also, filler nodes can straightforwardly correspond to concepts. E.g., the concept of "dogs" clearly plays the agent role in the proposition "dogs chase cats". Having said this, further elaboration is in order. In what follows, it should be understood that a "concept node" is *not* being identified with a single neuron. Most, and possibly all, of our concepts involve the representation of multiple *features* (e.g., "frog" has features corresponding to shape, colour, etc.). Such features, in turn, are widely believed to be representable only when specific sets of neurons become active. For this reason, I will assume that the active representation of any concept node will likewise require the activation of a significant number of neurons. Thus, each *concept node* is implemented by a "neuron-set", and *nodes* are abstractions of such sets.

Moreover, for all that is known to the contrary, it is plausible that within each such neuron-set, there exists a small subset (of perhaps 10 or fewer neurons) that form a reverbatory *clique* (or group) such that:

(a) every neuron within the clique forms excitatory connections with all others in that clique,

(b) every neuron, N, in the total "concept neuron-set", is connected to several members of that clique in a way that permits N to be jointly excited by the activation levels of those members.

In what follows, I take it as a working hypothesis that for each set of neurons, whose joint activations reliably qualify as representing a given concept, C (whether C is a role or a filler), there does exist such a clique. Also, due to conditions (a) and (b), above, it follows that the activation of most (say, 70% or more) members of a concept node's clique will typically lead, via spreading activation, to the activation of all elements of the concept's *total neuron-set*. In light of this, the following shorthand is adopted: when I describe "two nodes as being connected (by a link)", I am actually saying that within each of the two cliques associated with those *abstract nodes*, most (70% or more) of the neurons within one of the cliques make afferent (input) connections on most of the neurons in the other clique, and vice-versa. Given the multiplicity of connections between the neurons in the respective cliques, and also within each clique, it will be possible for each of the two "nodes" to spread its activation to the other node.

Now, before delving into further details, I would remind readers that the model described presently is not intended as a biologically accurate account of how synchronous firing is caused. Rather, it is best viewed as a suggestive metaphor which illustrates one possible way in which CBNs might provide underpinnings for synchronous firing. The metaphor displays how such synchrony could be viewed as a symptom of a binding's having occurred, rather than as the mechanism of role-filler binding. Moreover, the model demonstrates that, in principle at least, a feasibly small number of CBNs could provide the required underpinnings.

Initially, for diagrammatic simplicity, the model is presented on a small scale, and requires only about 30 CBNs. Yet, even this scaled down version would suffice for a pool of 5000 filler (concept) nodes and a pool of 1000 role nodes. It will soon emerge, however, that even when the model is scaled up to accommodate, say, 12 times that number of filler and role nodes, only about 6400 CBNs would be required.

A noteworthy feature is that, although nodes within the model's diagram (see Figure 1) are shown in fairly close proximity and spatial symmetry, in reality there would be no need for nodes to be in close proximity, or to be separated by regular distances. That is, even a perfect biological analogue to the model could incorporate neurons that are widely distributed, and distances between functionally similar nodes could be rather irregular. E.g., the nodes shown in light grey (or *blue* in colour printings) would not need to be separated by equal distances, although functionally they would need to preserve a connectivity pattern analogous to the one displayed.

Also worth noting is that there is some evidence that mechanisms that actually cause synchronous firing, in cases where the firing is not directly provoked by sensory input, can involve widely distributed cortical connections. As Singer (1999) observes, there are "indications that [the causes] are mediated at least in part by cortico-cortical connections" amongst cells that reside not only "in the same area but also cells distributed across different areas and even across the two hemispheres."

---

[2] A given node is considered to be the "most active" filler node (respectively, role node) provided it has the highest firing rate among all filler (respectively, role) nodes connected to a given *core network*.

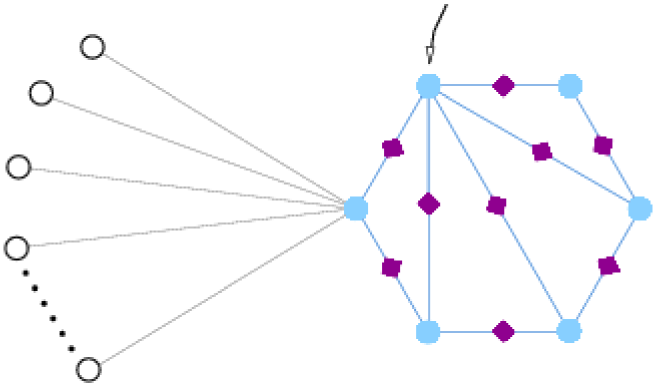

Figure 1: Source nodes (white circles at far left) provide input to the hex-nodes (shown in light grey). The diamond-shaped, dark nodes are hex-binders.

Moreover, synchrony has been observed in the firing of neuron sets located as distantly as the two hemispheres. Given this, it is likely that mediation between the two neuron sets will involve intervening neurons. There is, then, *prima facie* plausibility, at least, to the notion that synchronous firing would be engendered by intervening CBNs. We now consider one way in which that might occur.

Architectural Aspects of the Model

Figure 1 displays, schematically, the general structure of the model in question. The differing colours of the nodes correspond to the differing functions they possess. For example, nodes displayed in *white* are the *source nodes*, that is, they represent those fillers and roles whose initial firings trigger a series of complex interactions that eventually result in a single filler node and a single role node coming to fire synchronously. Only a very small subset of the totality of white source nodes is displayed in Fig. 1. Note that all source nodes so displayed are linked to a single node, shown in light grey. We will assume that attached to all other grey nodes are comparable sets of white source nodes. The white nodes are, thus, separable into clusters, depending upon which grey node they “feed into”. Let us assume that attached to each grey node is a cluster of, on average, about 1000 white, source nodes.

In contrast with the source nodes, nodes displayed in light grey and in dark belong to a *core network* which is responsible for establishing a uniform firing rate that is eventually conveyed to a single filler and a single role node. Nodes displayed in dark are CBNs. Presently, more shall be said about both grey and dark nodes. All links displayed in Fig. 1, as well as other links implicitly assumed to be present, are reciprocal. (Excitation between nodes can flow in both directions; reciprocal pairs of axons are assumed.) Within the *core network*, six grey nodes are displayed. These are called *hex-nodes*, because, in this scaled-down model, the grey nodes form the vertices of an implicit hexagon. Because there are six hex-nodes, and each is connected to a distinct cluster of white nodes, there are six clusters of source nodes (five clusters are not displayed). We shall assume that exactly one of these clusters is comprised solely of *role* nodes, while the remaining clusters are each comprised entirely of filler (or concept) nodes.

Now consider the grey hex-node that is identified by a small arrow. Note that this node is attached to several dark diamond-shaped nodes, which are each attached to other grey nodes. As mentioned, each of the dark diamond nodes represents a CBN. Their task is to bind together a pair of grey hex-nodes.

Although not fully displayed (for reasons of clarity), each grey node is assumed to be connected to every other grey node, via an intervening dark CBN. Thus, each grey hex-node possesses as many attached dark nodes as the grey node indicated by the small arrow. That is, all grey nodes are fully connected to one another via intervening, dark, CBNs. Moreover, although not displayed in Fig. 1, all dark CBNs possess direct, reciprocal inhibitory links to one another. Thus, the CBNs all belong to a single, winner-take-all (WTA) network. The dark CBN that receives the greatest initial activation (at a key time) will eventually win a competition with all other dark CBNs.

Similarly, let us assume that all source (white) nodes that belong to any given cluster will possess direct, reciprocal, inhibitory connections with one another. That is, all source nodes in any single cluster belong to a single WTA competitive network, but the relevant inhibitive links are not shown in Fig. 1. Altogether, then, there are seven WTA networks implicit in the system displayed by Fig. 1. One of these WTA networks involves just the dark CBNs within the core. The remaining six WTA networks each involve just the source nodes within a cluster (one WTA net per cluster). Assume that all inhibitory links supporting WTA competitions are equally weighted.

Before delving into further processing aspects, let us note that each node within the core network (i.e., each hex-node and each CBN) may best be viewed as an abstraction of a neuronal clique. Internally (as with the concept nodes), each clique would contain mutually excitatory neurons. Readers may have observed that cliques as here conceived bear a striking resemblance to certain of Hebb’s *cell assemblies* (Hebb, 1949). Due to the high degree of intra-connectedness within cliques, we may reasonably suppose that all neurons within a clique attain nearly the same firing rate. Consequently, the *firing rate of any abstract node* will here be identified with the average firing rate within the corresponding clique. (Prior remarks about connectivity *between* cliques also apply here.)

Processing Aspects of the Model

We now consider how activation levels (which correspond to firing rates) are spread within the overall connectionist system, and how this propagation eventually establishes synchronous firing for a pair of filler and role nodes. For clarity, the processes are presented below as a series of steps, although in the hypothetical “reality” being modeled, distinct stages may be difficult to separate. Also, it is assumed here that all *active* nodes are subject to some degree of decay in their firing rates, but that such decay is *not* so rapid as to immediately nullify a node’s functional interactions.

1. At the outset, we assume that within each cluster of white, source nodes, there are some nodes that possess at least moderately active firing rates. Also, due either to processes involved in the comprehension of an externally presented *sen-*

tence, or in the formation of some individual *thought*, a particular source (filler) node in one cluster has acquired a higher activation level (firing rate) than all other filler nodes, and similarly, a particular role node in a *different* cluster attains a higher activation level than all other role nodes. The firing rates of these two highly active nodes are *not* assumed to be equal. I do assume, however, that these two most active nodes should become "bound" to each other, and that a sign of this binding is to be synchronous firing. (Recall that there exist connectionist parsing systems which illustrate how and why the filler and role nodes that *should bind* are the most active nodes of their respective types that *are available* to bind.)

2. Once nodes in each of the source clusters have become effectively active, each of those nodes spreads activation both towards their respective grey hex-node and towards the other source nodes within the WTA cluster that the given source node belongs to. (Recall that all source nodes within a given cluster possess reciprocal inhibitory links with each other.)

3. As this spreading activation reaches the grey hex-nodes, WTA competitions are initiated in each of the six clusters of source nodes. Eventually, a single winning source node is selected in each of the six clusters, and other source nodes in those clusters fall below effective activation levels. However, the WTA competitions need not be completed before activation has reached grey and dark nodes in the core network.

4. As grey hex-nodes receive activation (measured in firing rates) from the source nodes in their respective "feeding clusters", each such hex-node is caused to assume the firing rate of the most active source node in the respective cluster. This occurs because (we may assume) all links entering a grey hex-node make *disjunctive connections* with that node; thus, the received activation is not additive. (Bear in mind also that a hex-node could be realized by a clique of neurons.)

5. Note that there are exactly two distinct grey hex-nodes (call them X1 and X2) that have, respectively, assumed the firing rates of the *most active role* node and the *most active filler* node (from among the source nodes). This is so because all role nodes are in a separate cluster from all filler nodes. Just one of the hex-nodes receives activation from the role nodes.

6. Once each of the grey hex-nodes has (respectively) acquired the firing rate of its most active source node, the hex-node conveys its new firing rate to each of the dark CBNs to which it is immediately attached (call these dark nodes "hex-binders").

7. A WTA competition ensues among the dark hex-binders. Exactly one hex-binder (dub it HB) is directly connected to the two most active hex-nodes. (That is, HB is directly connected to X1 and X2.) For this reason, HB will become the most highly activated hex-binder, and thus will win the WTA competition among all the hex-binders. Upon winning this competition, HB will acquire some high, effective firing rate, R. The remaining hex-binders, having lost the competition, will fall to low, non-effective activation levels.

8. Because all hex-binders other than HB are non-effectively active, they are unable to re-excite hex-nodes to which they are directly connected. In contrast, HB, being highly active, will now re-stimulate its attendant hex-nodes, X1 and X2. In particular, HB is now able to convey its firing rate, R, to X1 and X2.

9. As a consequence of step 8, X1 and X2 now assume essentially equal firing rates, and are now in a position to convey this same rate, R, to their respective source nodes.

10. The reader will recall that while activation has been spreading within the core network, WTA competitions have occurred in each of the six source clusters. Since these competitions began at an early stage, we may assume that, prior to (or not later than) the time that a winner has been selected among the hex-binders in the core, winner nodes have also been selected in each of the six source clusters. Among these six winners will be the originally most active role node (RN) and filler node (FN). Recall that RN and FN are, respectively, directly connected to X1 and X2.

11. Hex-nodes X1 and X2 have effectively the same firing rate, R. Moroever, these are the only two hex-nodes directly connected (respectively) to the two source nodes, RN and FN. Also, among the grey hex-nodes, only X1 and X2 have recently been stimulated by a winning hex-binder (there was only one winning, dark hex-binder). Due to this recent re-stimulation, X1 and X2 are assumed to fire towards their respective source clusters. Whether any of the other hex-nodes do likewise is irrelevant, since only X1 and X2 have any particular reason to have the same firing rate, R.

12. Therefore, among the six active *source* nodes (which have each undergone some activation decay *since becoming winners*), only RN and FN will now receive re-stimulation from hex-nodes that can be safely assumed to have an approximately identical firing rate. Thus, among the source nodes, only RN and FN will now reliably acquire an identical firing rate.

In summary, RN and FN have now been caused to fire in (at least, approximate) synchrony. Their firing rate may not be identical to that of X1 and X2, but that will not matter. What matters is that RN and FN have now assumed approximately equal firing rates. If their rate is essentially the same as that of X1, X2 and HB, this does no harm, because the latter three nodes do not belong to the set of role and filler nodes.

Scaling Up the Model

The model described thus far clearly requires only a small number of CBNs. For, the CBNs occur only between each pair of hex-nodes, and there are just six hex-nodes. (Thus, less than 6^2 CBNs are present.) However, because I have assumed that only about 1000 source nodes are connected to each hex-node, the total number of filler and role nodes permitted would be around 6000. Nevertheless, the architecture could easily be scaled up to accommodate, say, a total of 80,000 role and filler nodes, without involving a vast number of CBNs. For example, assume that instead of just six hex-nodes, the core network contains 80 nodes whose connectivity and functionality is analogous to that of the hex-nodes. Thus, the perimeter of the core network would now contain 80 grey nodes, and there would be 80 separate source-node clusters. Each of the 80 grey nodes must, of course, be connected to the remaining 79 grey nodes via an intervening dark CBN (hex-binder), but this necessitates less than 80^2 CBNs. So, less than 6400 CBNs would be required. Each CBN might be implemented by, say, 10 neurons within a clique, and we should factor in the 80 hex-nodes. Nevertheless, the total "neuron count" for the core network would then be about

65,000, which is a modest number. Assuming, once again, that each of the 80 source clusters contains roughly 1000 filler or role nodes (not both kinds in the same cluster), we see that about 80,000 filler and role nodes could be accommodated by means of a highly feasible number of CBNs.

Summary and Conclusions

In the foregoing, I have discussed the comparative merits of conjunctive coding and synchronous firing as methods of binding and have argued:

(i) that we are not actually in a position to judge the number of nodes required by *synchronous binding*, because we do not yet know the numerical range of neurons involved in the underlying processes that engender this type of binding.

(ii) that the genesis of synchronous firing in non-sensory domains may, for all we know, involve a substratum of CBNs.

In addition, I have, in section 2, presented a connectionist model, grounded upon conjunctive coding, which illustrates how synchronous bindings could be engendered in the realm of roles and fillers. The model clearly requires just a comparatively small number of conjunctive nodes, and (contrary to my diagrams' proportions) is compatible with a widely distributed configuration of nodes, separated by irregular distances. As previously acknowledged, the model is not proposed as biologically realistic. Rather, it is presented as a hypothetical analogue whose details may suggest the direction in which the true explanation of synchronous binding is to be found. In any case, the hypothetical model illustrates an important possibility, i.e., that binding via synchronous firing is merely a symptom of a deeper form of binding, one that intrinsically involves the operation of conjunctive binding nodes.

References

Cottrell, G. 1985, Connectionist parsing. *Proceedings of the Seventh Annual Conference of the Cognitive Science Society*, 201-211, (Mahwah, NJ: Lawrence Erlbaum Associates, Publishers).

Hadley, R.F. and Cardei, V.C., 1999, Language acquisition from sparse input with no error feedback. *Neural Networks*, **12**, 217-235.

Hadley, R.F. and Hayward, M.B., 1997, Strong semantic systematicity from Hebbian connectionist learning. *Minds and Machines*, **7** , 1-37.

Hebb, D., 1949, *Organisation of Behaviour*, London.

Hummel, J.E., and Holyoak, K.J., 1997, Distributed representations of structure: A theory of analogical access and mapping. *Psychological Review*, **104**, 427-466.

Lange, T. and Dyer, M., 1989, High-Level inferencing in a connectionist network. *Connection Science*, **1**, 181-217.

Rumelhart, D.E., McClelland, J.L., and the PDP Research Group (Eds.), 1986, *Parallel Distributed Processing; Explorations in the Microstructure of Cognition, Volume I* (Cambridge, MA: MIT Press).

Shastri, L., and Ajjanagadde, A., 1993, From simple associations to systematic reasoning: A connectionist encoding of rules, variables, and dynamic bindings using temporal synchrony. *Behavioral and Brain Sciences*, **16**, 417-494.

Singer, W., 2003, Synchronization, Binding and Expectancy. In M.A. Arbib (Ed.), *The Handbook of Brain Theory and Neural Networks*, 1136-1143, (Cambridge, MA: MIT Press).

Singer, W., 1999, Neuronal synchrony: A versatile code for the definition of relations? *Neuron*, **24**, 49-65.

Stevenson, S., 1994, A unified model of preference and recovery mechanisms in human parsing. *Proceedings of the Sixteenth Annual Conference of the Cognitive Science Society*, 824-829, (Mahwah, NJ: Lawrence Erlbaum Associates, Publishers).

Sun, R., 1992, On variable binding in connectionist networks. *Connection Science*, **4**, 93-124.