right kind of input nodes of the neural control circuits in order to bind word assemblies to structure assemblies and to bind different structure assemblies via their subassemblies. For instance "cat chases mouse" and "mouse chases cat" both have the same surface N-V-N structure and are distinguishable by the blackboard model only if agent and theme are assigned inversely. The underlying event structure, however, cannot simply be read off the word order as the authors seem to suggest. This is clear from notorious examples like "The man frightens the child" versus "The child fears the man." Both sentences describe the same event, but neither the position in the sentence nor the syntactic category signals the thematic role of "man" and "child."

Moreover, concurrent activation of assemblies for "mouse" and "chases" does not resolve whether "mouse" is the actor or the theme of "chases." Hence simultaneous activity as a binding mechanism, that is, "and-binding," works only if the assignment of thematic roles is already accomplished by a preprocessor. vdV&dK do not sufficiently comment on how this task is achieved. The proposed model therefore does not solve the binding problem but merely shifts the burden of constructing appropriate neural representations of sentence structure to unexplained preprocessing of the linguistic input. Consider, for instance, the sentence "The fact that the boy comprehends surprises the teacher." Unlike the very similar examples in the target article, this sentence does have two readings for the same N-that-N-V-V-N sequence. How the model "parses" this sentence depends entirely on the preprocessor's choices on whether "the fact" should be assigned the theme role of "comprehends" or not.

This shortcoming is particularly problematic for real-world ambiguous sentences. If vdV&dK's model already runs into trouble for simple three-word sentences, what would it make of real-world sentences for which there can be an interaction of several types of ambiguities, such as attachment ambiguities (like PP attachments and reduced clause attachments), functional ambiguities (like subject, object and indirect object assignments), and scope ambiguities (like different scope of adjectives). Any processing or parsing model for natural language should be able to tell us how various (neural) representations can be assigned to ambiguous sentences, and many parsing models even go further and also tell us how to select the preferred representation (e.g., Bod 1998; Charniak 1997);. vdV&dK's model does neither; as a parsing model for real-world sentences the approach is thus inadequate.

Should we, then, interpret the model as just a neurally inspired model of representation? The problem with such an interpretation is that without specifying the processes that do something with the representations – learning, interpretation, production, inference – it is unclear what the requirements on representations are. At the very least, representations should distinguish between all sentences that have different phrase structure. Unfortunately, the authors do not show that their formalism can make these distinctions for more than a handful of simple examples, nor do they analyze the expressive power of the formalism proposed and compare it with established results in theoretical linguistics (e.g., Chomsky 1957; Joshi et al. 1991). The real question that we believe remains unanswered is how networks of neurons, of the type that are found in the human brain, can construct appropriate representations and put them to use in actual language processing.

ACKNOWLEDGEMENT

Willem Zuidema is funded by the Netherlands Organisation for Scientific Research (Exacte Wetenschappen), project number 612.066.405.

How neural is the neural blackboard architecture?

Yoonsuck Choe

Department of Computer Science, Texas A&M University, College Station, TX 77843-3112. ***choe@tamu.edu*** ***http://faculty.cs.tamu.edu/choe***

Abstract: The target article does not provide insight into how the proposed neural blackboard architecture can be mapped to known neural structures in the brain. There are theories suggesting that the thalamus may be a good candidate. However, the experimental evidence suggests that the cortex may be involved (if in fact the blackboard is implemented in the brain). Issues arising from such a mapping will be discussed.

In the target article, van der Velde &de Kamps (vdV&dK) present a neural blackboard architecture for the representation of combinatorial structures in cognition. However, they do not provide much insight into how their proposed architecture (Fig. 3 in target article) can be mapped onto known neuroanatomy and neurophysiology.

In this commentary, I will assess the idea that the thalamus is a key structure to which the neural blackboard architecture can be mapped. In fact, there are existing theories in this line: In addition to the work of Newman and colleagues (1997), which is cited in the target article, Harth and colleagues (1987) and Mumford (1991; 1995) also propose that the thalamus may be a good candidate for such an internal blackboard or a sketchpad. The rational is that the thalamus receives dense feedback from all parts of the cortex and hence is an ideal place for the results of cortical processing to be written and combined. Hence the thalamus on its own may take on an integrative role under these theories. (Also see Choe 2004 for a possible role of the thalamocortical circuit in processing analogy.)

However, neuroscientists have argued that the thalamus does not play such an integrative role in which signals from multiple sources in the cortex converge to a small, local region of the thalamus (Sherman & Guillery 2001, p. 102). On the other hand, there is also evidence that there may be intrathalamic interactions that can allow some form of integration through the thalamic reticular nucleus containing inhibitory neurons (Crabtree et al. 1998). Hence it is unclear whether a blackboard, if it is in fact implemented in the brain, can be found solely within the thalamus or if other parts of the brain such as the cortex can also serve a role.

Here I will assess the compatibility of the gating circuit in the target article's Figure 3 with the thalamocortical circuit. The basic finding is that the model of vdV&dK needs both the cortex and the thalamus, as well as the thalamic reticular nucleus (TRN) as a neural substrate. Figure 1(a) shows a schematic diagram of the thalamocortical circuit based on known experimental findings (see Sherman & Guillery 2001 for a review). Figure 1(b) shows how the gating circuit in the target article's Figure 3 can be nicely mapped onto the thalamocortical circuit, where the major representations, such as part-of-speech and thematic role, and the control all come from the cortex, and the disinhibitory gating happens at the thalamic reticular nucleus (TRN) and the thalamic relays. A connection new in (b) compared to (a) is the connection from X_{out} to Y, which can be seen as an ascending branch of the axon from T_2 to C_2 (see, e.g., Macchi et al. 1984). The mapping in Figure 1(b) suggests that the blackboard architecture may not be confined to the thalamus as Newman and colleagues, Harth and colleagues, and Mumford and colleagues suggest. Rather, the cortex would serve as the major blackboard, and the thalamus would only play the role of gating.

Several issues become apparent when the architecture is put into a concrete context as shown in Figure 1(b). One key assumption in the neural blackboard is that duplication of structure assemblies is easy compared to that of word assemblies. But what does it mean for an assembly to be duplicated? There are two possibilities: (1) the anatomical structure is duplicated on the fly, which seems highly unlikely (but Figure 7b in the target article indicates this may be the case); or (2) the pattern of activation that represents a structure assembly is duplicated, which seems too biologically underconstrained and abstract (i.e., how can various binding and control reach the dynamically changing patterns?).

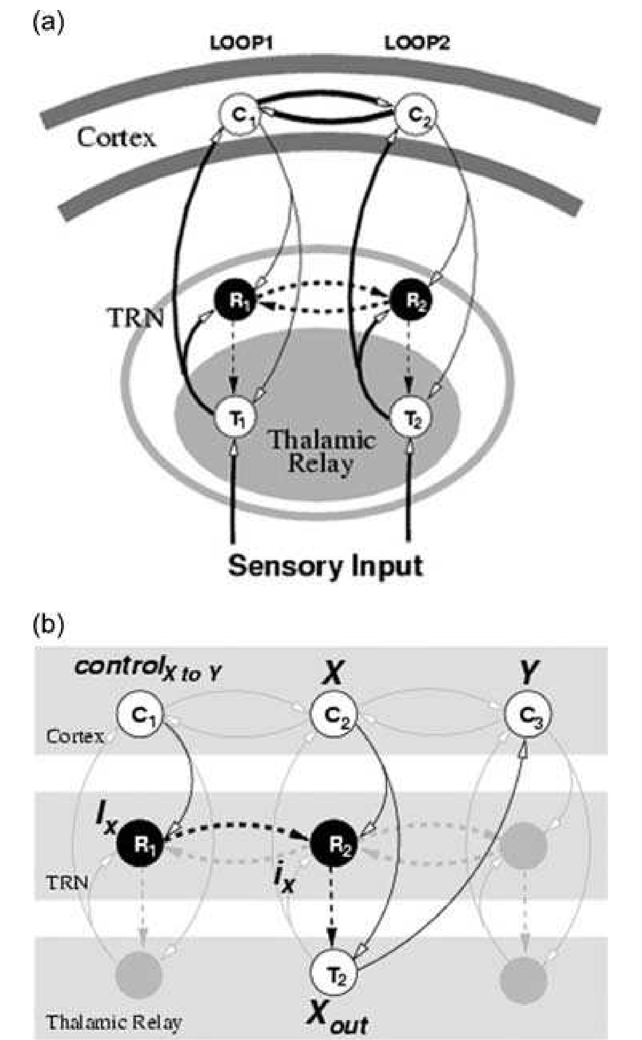

Figure 1 (Choe). Thalamocortical circuit and the gating circuit mapped onto it. (a) A schematic diagram of the thalamocortical circuit is shown (adapted from Choe 2004). White discs are excitatory neurons and black ones are inhibitory neurons, whereas solid edges are excitatory connections and dashed ones inhibitory. (b) The gating circuit in the neural blackboard architecture is mapped onto the thalamocortical circuit. The neurons and connections that may not participate in the gating circuit are shaded out in dark gray. Besides that, the only difference between the circuit here and that in (a) is the ascending thalamocortical projection from T_2 to C_3, i.e., X_{out} to Y (see text for a discussion).

Another issue is about how different word assemblies can be bound to a particular structure assembly construct such as N_1 in the target article's Figure 2 (e.g., cat -x- N_1). For this to work, potentially all cortical areas representing a word assembly (taking the place of X in Fig. 1[b]) need to have a physical connection to the relevant gating circuit in the thalamus (i_x and X_{out}) that project to N_1 (Y): a combinatoric explosion of connections. A related problem is about how control from the central pattern generator ($control_{X\ to\ Y}$) gets attached to the gating circuit (I_X). For the central pattern generator precisely to control the flow of activation in all the gating circuits, it needs to project to a large number of participating TRN neurons, which may, again, become too resource-intensive.

In sum, (1) it is possible that the gating circuit can be mapped to a concrete neural circuit in the thalamocortical network; but (2) with this concrete view of the gating circuit, the flexibility or dynamic nature of the proposed architecture comes under question, thus undermining the claimed productivity in the blackboard architecture. The blackboard architecture may be sufficient when it remains abstract, but once it becomes grounded in neural terms, its combinatoric power may suffer. Hence the question: How neural is the neural blackboard architecture?

How anchors allow reusing categories in neural composition of sentences

William J. Clancey

NASA Ames Research Center, Intelligent Systems Division, Moffett Field, CA 94035. ***William.J.Clancey@NASA.gov*** ***http://bill.clancey.name***

Abstract: van der Velde's &de Kamps's neural blackboard architecture is similar to "activation trace diagrams" (Clancey 1999), which represent how categories are temporally related as neural activations in parallel-hierarchical compositions. Examination of other comprehension examples suggests that a given syntactic categorization (structure assembly) can be incorporated in different ways within an open composition by *different kinds of anchoring relations* (delay assemblies). Anchors are categorizations, too, so they cannot be reused until their containing construction is completed (bindings are resolved).

Conceptual Coordination (Clancey 1999, hereafter CC) attempted to bridge connectionism and symbolic theories of cognition by analyzing slips, autism (Clancey 2003), learning in well-known architectures (e.g., MOPS, EPAM, SOAR), dream phenomenology (Clancey 2000), and natural-language comprehension difficulties. The neural blackboard architecture of van der Velde & de Kamps (vdV&dK) (pp. 15ff) has properties that directly map to "activation trace diagrams" (CC, p. 6), including: (1) representing a conceptualization of a sentence on a blackboard; (2) not copying structures (i.e., one physical instantiation in a given sentence diagram); (3) preserving temporal relations of neural activation; (4) composing word assemblies by structural relations (e.g., a noun phrase, NP); (5) creating word structures (e.g., "the cat - NP") by temporary binding/tagging (CC, p. 273); (6) binding words to multiple structures in different roles (e.g., both agent and theme); and (7) holding a partially completed structure active (a "delay" activity in vdV&dK; "anchor" in CC, p. 246).

Figure 1 represents conceptualization of a double-centered embedded sentence from vdV&dK using the notation from CC. During the composition process, the first noun phrase (NP1) is held active by an anchor, contained within subsequent phrases, and then resolved in an agent role by a right branch. A verb phrase is represented (in vdV&dK's terms) as an agent and verb on a line segment (e.g., "boy-likes"), with a theme (e.g., object, "dog") indicated by an arrow below. This sentence is analogous to "The cat that the bird that the mouse chased scared ran away" (CC, fig. 10.14, p. 256). vdV&dK (pp. 47ff) indicate that these sentences are difficult to comprehend because there are two NPs (NP2 and NP3) that could bind to two different theme assemblies (e.g., "likes" and "bites").

Based on analysis of other convoluted examples (provided by Lewis 1996), the binding problems involved in sentence comprehension can be restated and generalized. The following principles postulate restrictions on how word (semantic) and structural (syntactic) categorizations are employed in a given construction (from CC, pp. 261–62):

(P1) A semantic categorization can play multiple syntactic roles but not the same role twice in an active construction (e.g., "cat" cannot be an agent in two verb phrases).

(P2) A syntactic categorization may "index" multiple semantic categorizations but may be active only once at a given time (e.g., only one active NP; the others are in "completed" verb phrases).