
Combinatorial structures and processing in Neural Blackboard Architectures

Frank van der Velde
CTIT, CPE-BMS
University of Twente, Enschede
IOP, Leiden
The Netherlands
f.vandervelde@utwente.nl

Marc de Kamps
Institute for Artificial Intelligence
and Biological Systems
University of Leeds
Leeds, United Kingdom
M.deKamps@leeds.ac.uk

Abstract

We discuss and illustrate Neural Blackboard Architectures (NBAs) as the basis for variable binding and combinatorial processing the brain. We focus on the NBA for sentence structure. NBAs are based on the notion that conceptual representations are in situ, hence cannot be copied or transported. Novel combinatorial structures can be formed with these representations by embedding them in NBAs. We discuss and illustrate the main characteristics of this form of combinatorial processing. We also illustrate the NBA for sentence structures by simulating neural activity as found in recently reported intracranial brain observations. Furthermore, we will show how the NBA can account for ambiguity resolution and garden path effects in sentence processing.

1 Introduction

We outline and illustrate the implementation and processing of combinatorial structures in Neural Blackboard Architectures (or NBAs). First, we will discuss the ideas on which NBAs are based. Then, we will illustrate sentence representation and processing in the NBA for sentence structure [1]. Next, we will simulate neural activation patterns that arise in the NBA during sentence processing and relate it to recently reported intracranial brain observations. Finally, we will illustrate and discuss how the NBA handles linguistic phenomena such as unproblematic ambiguities and garden path sentences [2].

2 In situ concept representations

The development of Neural Blackboard Architectures is based on the notion that representations in the brain, such as representations of concepts, are ‘in situ’, in line with the neural assemblies proposed by Hebb [3]. Figure 1 (based on [4]) illustrates this with the concept *cat*. As Hebb argued, a neural assembly develops over time as neurons that process information about a concept become interconnected. These neurons could be located in different parts of the brain, depending on the nature of the underlying concept. So, in case of a concept like *cat*, they will consist of neurons in the visual and auditory cortex, involved in the perception of cats. They could also consist of neurons involved in actions related to cats (e.g., pronouncing the word *cat*).

Concept representations with neural assemblies could be distributed, but parts of the assembly could also consist of more local representations. They could consist of neurons involved in processing information about the concepts, but also of neurons involved in actions related to concepts. This entails that both incoming and outgoing connections determine the nature of a concept representation [5]. Furthermore, unlike the original assemblies as described by Hebb, concepts as intended here

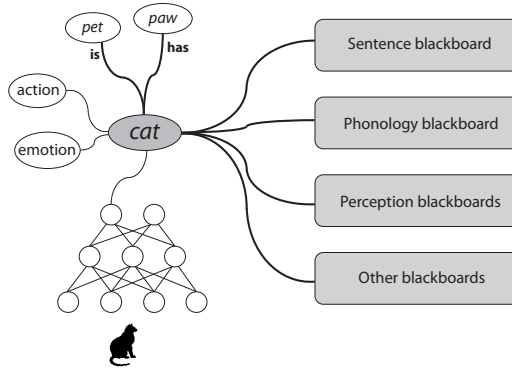


Figure 1: In situ concept representation (based on [4]).

are not only associative neural structures. They also incorporate relations, as illustrated with the relations *is pet* and *has paw* in Figure 1.

The assembly or web-like structure of a concept representation entails that concepts representations are ‘in situ’ [4]. That is, wherever a concept is activated it always consists of the activation of the assembly of that concept or a part of it. In this view, it is not possible to make a copy of a concept representation and store it elsewhere in the brain to, for example, create a combinatorial structure with it or use it in a reasoning process. In this sense, in situ concept representations are fundamentally different from symbols in symbol architectures [6]. Due to the in situ nature of concepts representations, they are also grounded in perception or action (depending on the nature of the concept).

The in situ nature of neuronal concept representations imposes constraints on the way they can be combined in higher-level cognitive combinatorial structures, such as sentences. In situ concept representation cannot be copied or transported to construct combinatorial structures. Instead, as illustrated in Figure 1, they are embedded in ‘neural blackboards’ to form combinatorial structures. We assume that there will be specific neural blackboards for specific forms of combinatorial structures, and concept representations will be embedded in a number of them depending on the nature of that concept and the role it plays in specific combinatorial structures. For example, the concept *cat* will be embedded in a phonology blackboard, forming the word *cat* based on familiar phonemes. It will also be embedded in a sentence blackboard, needed to form sentence structures. It could also be embedded in a perception blackboard underlying the combinatorial nature of visual cognition [7]. Other blackboards, e.g., needed for sequential processing or reasoning, could be included as well.

The neural blackboards allow the construction of (potentially novel) combinatorial structures based on (familiar) in situ concept representations, using forms of variable binding (depending on the specific blackboard). They aim to satisfy the following characteristics:

- The concept representations remain in situ (and thus grounded) even when they are a part of (novel) combinatorial structures.
- The combinatorial structures or parts thereof are content addressable, even with novel combinatorial structures (of familiar constituents).
- The variable binding of concepts in blackboards creates the connection paths needed to produce behavior (linking sensory and motor activation), even with (novel) combinatorial structures.

We will discuss and illustrate these characteristics using the sentence blackboard.

3 Sentence neural blackboard

Figure 2 illustrates the structure of the sentence *Sonnet sees (the) dog* in the Neural Blackboard Architecture for sentence structure [1]. Here, we assume that the words *sees* and *dog* are familiar

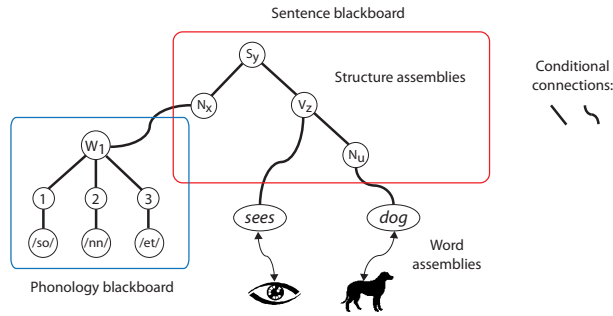


Figure 2: Sentence representation in a sentence neural blackboard.

constituents but *Sonnet* is a new name, not heard before. The representations of *sees* and *dog* consist of neural assembly structures ('word assemblies'), in line with the concept *cat* in Figure 1. For simplicity, we represent them with ovals, but the full connection structure is always implied.

The sentence blackboard allows the creation of a (temporal) connection structure that binds the word assemblies in line with the relations expressed by the sentence. To this end, the NBA consists of 'structure assemblies' such as N_x and N_u in Figure 2 (representing nouns) and S_y (sentence) and V_z (verb). The word assemblies can bind to structure assemblies of their kind, and structure assemblies can bind to each other in line with the sentence structure. The binding process results from specific 'conditional connections' (outlined below).

As noted in [8], humans have the ability to include new words and names in (even novel) sentences structures, exemplified here with the name *Sonnet*. However, the ability to include novel words in sentences is restricted to new words based on familiar phonemes. In the NBA this ability results from a phonology blackboard in which new names can be formed based on familiar phonemes [9]. The phonology blackboard interacts with the sentence blackboard, which allows the incorporation of the new name in a sentence structure [10]. Here, *Sonnet* activates a set of phonemes (illustrated with the pseudo phonemes *so*, *nn*, and *et*), which bind to a structure assembly in the phonology blackboard (W_1). In turn, this structure assembly binds to the sentence structure in the sentence blackboard.

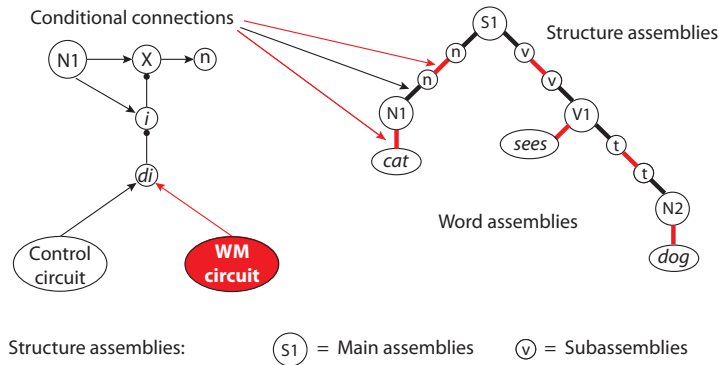


Figure 3: Binding in the sentence neural blackboard.

The binding process in the NBA is illustrated in more detail in Figure 3 with the sentence *cat sees dog*. The conditional connections, here the red and black thick lines, consist of gating circuits. In the NBA, each word assembly (e.g., of a noun) is connected to a set of structure assemblies of the same kind (all N_i assemblies in the case of a noun) with gating circuits. (The number of these connections can be reduced by representing words in the phonology blackboard first, see [9].) In turn, each structure assembly consists of a 'main assembly', such as N_1 , and (a set of) subassemblies, such as n or t . The connection between a main assembly and a subassembly consists of a gating circuit as well.

Structure assemblies of different kinds, such as V1 and N2, are connected by their subassemblies of the same kind. Here, their t (theme) subassemblies, which represents the fact that a verb can have a theme (object). This connection also consists of a gating circuit.

The gating circuits operate by disinhibition (*di*), illustrated in Figure 3. So, when N1 is active, it activates a neuron (or neuron population) *X* and an inhibitory neuron (or population) *i*. The latter inhibits *X*, which blocks the flow of activation. But when *i* itself is inhibited (by neuron or population *di*) activation can flow from N1 (via *X*) to *n*.

Gating circuits can be disinhibited (or ‘activated’) in one of two different ways. In case of gating circuits between main assemblies and subassemblies the activation results from an external control circuit that activates the *di* population. This is how syntactical operations affect binding in the blackboard. A control circuit could have recognized that *sees dog* represent a verb and a theme. It then activates all *di* populations in the gating circuits between all V_i and N_j assemblies and their t assemblies. As a result, the active V_i and N_j will activate their t subassembly.

Gating circuits between subassemblies and between word and main assemblies are activated by (specific) ‘working memory’ (WM) populations. A WM population remains active for a while, after initial activation, by reverberating activation in the population [11]. An active WM population binds the assemblies to which it is connected in the manner outlined below.

3.1 Variable binding in a connection path

Figure 4 illustrates how binding is achieved in the NBA. The parts (a), (b) and (c) illustrate the same binding process with increasing detail. In (a), the binding between the t subassemblies of V1 (or V1-t) and N2 (N2-t) in Figure 3 is repeated. Figure 4b illustrates that this binding is based on a ‘connection matrix’, which consists of columns and rows of ‘connection nodes’ illustrated in Figure 4c.

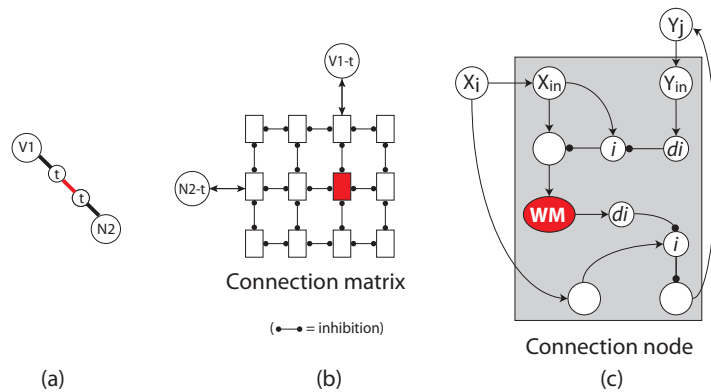


Figure 4: Binding by activity in a Connection Matrix (based on [1]).

Each specific V_i -t and N_j -t pair of subassemblies is interconnected in a specific connection node, located in a connection matrix dedicated to binding V_i -t and N_j -t subassemblies. In general, when two assemblies X_i and Y_j (e.g., V1-t and N2-t) are concurrently active in the processing of a sentence, they activate a WM population in their connection node by means of a gating circuit, as illustrated in Figure 4c. In turn, the active WM population disinhibits a gating circuit by which activation can flow from X_i to Y_j , and another such circuit, not show in (c), by which activation can flow from Y_j to X_i . As long as their WM population is active, X_i and Y_j are ‘bound’ because activation will flow from one to the other whenever one of them is (initially) activated.

The NBA allows any noun to bind to any verb in any thematic role using dedicated connection matrices. Also, the NBA has structure assemblies that can bind to other structure assemblies, such as S1 in Figure 3, or clause structure assemblies ([1]). In this way, hierarchical sentence structures can be represented, such as relative or complement clauses.

A given word assembly can bind to different structure assemblies at the same time, allowing the creation of sentence structures in which words are repeated (e.g., *cat sees cat*). Yet, word assemblies remain in situ in this way, so words in sentence structures are always accessible and grounded.

The role of the NBA is thus to establish a ‘connection path’ for any sentence structure that can be formed in a language. Although it has been questioned whether this would be possible [8], an available connection path is in fact a necessary condition for any model of (neuronal) variable binding. That is, whenever two neuronal representations are bound in a functional way, there is some kind of (direct or indirect) connection path between those representations. This path is needed to produce (meaningful) behavior based on the functional binding at hand. For example, to answer questions about the relation expressed by the binding [10]. The ability of the brain to establish such connection paths derives from its ‘small-word’ like connection structure [12]. The NBA indeed has such a structure [1].

In the NBA the connection matrices are given. But in a pilot simulation [13] we investigated the possible development of such a connection structure. The simulations indicated that a structure like a connection matrix could arise during development, based on initially random activation and connection patterns. But, of course, each node in the connection matrix would also have to have the circuits as illustrated in Figure 4c for the binding process to occur.

3.2 Content addressable sentence structures

The role of the NBA to (temporarily) create a connection path needed for the production of behavior is illustrated in Figure 5. This figure illustrates how the NBA can answer (binding) questions (for a simulation, see [1]).




Figure 5: Content addressable sentence structures and competition in the NBA in answering (binding) questions. Grey nodes are active.

The structure of the sentence *Bill knows John*, stored in the NBA, can be used to answer a (binding) question like “What does Bill know?” (or *Bill knows?*). The question activates the assemblies for *Bill* and *knows* and provides information that *Bill* is the subject of *knows*. This results in the activation of the sentence structure *Bill knows*. The question also asks for the theme of the relation *Bill knows*, which can be used to activate the gating circuits for the theme (t) subassemblies. This will initiate a flow of activation from V1 to N2 and thus to *John* in Figure 5, producing the answer to the question.

The NBA can correctly produce the answer even when, for example, *Bill* is a constituent of other sentences as well, as in *Bill sees Mary*. The representation of *Bill* is the same in situ word assembly in both sentences (which directly derives from the notion of in situ representation). As noted above, they are distinguished in the NBA by the binding with different noun main assemblies (here, N1 vs N3). The question *Bill knows?* activates *Bill*, which in turn activates both N1 and N3 to which it is bound. Because the question entails that *Bill* is the subject, this results in the activation of both S1 and S2.

However, the question also identifies *knows* as the verb and not *sees*. This results in the activation of V1 which inhibits the activation of V2, because main assemblies of the same kind inhibit each other in the NBA [1]. In this way, asking for the theme produces the correct answer *John* (bound to *knows*) instead of the incorrect answer *Mary* (bound to *sees*).

Figure 5 illustrates the role of a connection path in producing behavior. The question provides sensory information resulting in the activation of *Bill knows*. The answer (*John*) is produced by a motor program that would be activated by the in situ word assembly of *John*. If there were not a connection path that would link the sensory information (*Bill knows*) to the motor activation (*John*) it is difficult to see how this behavior could be produced. Yet, this kind of behavior is possible for any arbitrary binding of words in (potentially novel) sentence structures or even new words in sentence structures, as in Figure 2. In each of these cases, the behavior of answering a question (probing for relation information or binding) depends on connecting sensory information to motor activation, in a manner that satisfies the binding relations expressed by the relation at hand. The NBA produces this connection path due to its small-world like connection structure and its ability to bind arbitrary relations by the binding process outlined above.

Figure 5 also illustrates the role of content addressability in the NBA. Due to the in situ nature of word (concept) representation, the question *Bill knows?* directly activates the word assemblies *Bill* and *knows*. Because they are bound to the sentence structure in the NBA (i.e., because word representations remain in situ in any sentence structure) the activation of *Bill* and *knows* directly activates the partial sentence structure *Bill knows*.

It seems that the structure of a question itself indicates that sentence structures are content addressable. After all, the question *Bill knows?* presents the words *Bill* and *knows*, indicates that *Bill* is the subject of *knows* and indicates that the response should be the theme. All of this information is precisely the information needed to directly activate the partial sentence structure of *Bill knows* (instead of other partial structures such as *Bill sees*) and to initiate the flow of activation to produce the answer *John*. It is difficult to see how this process could be more direct and thus faster than in this way. Given the evolutionary pressure on fast and yet meaningful behavior, it seems that the NBA fits these constraints in an optimal way.

4 Sentence processing in the NBA

Figure 6 illustrates how sentence processing can proceed in the NBA [14]. Here, sentence processing results from the interaction between a ‘control network’ and the activation in the neural blackboard. Figure 6 illustrates the processing of *cat sees dog* presented in Figure 3.

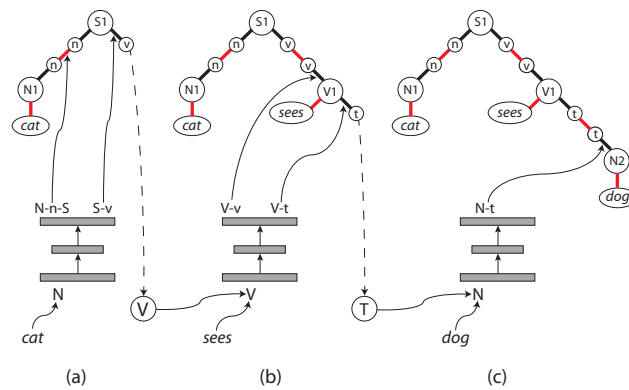


Figure 6: Sentence processing in the NBA (based on [14]).

Sentence processing is incremental in the NBA. The input to the control network consists of word information (e.g. Noun, Verb) and predictions based on the activation in the blackboard that arise during processing. Based on its input, the control network has learned to activate the gating circuits in the NBA needed to create the sentence structure. In (a) *cat* is recognized as a noun and the network activates the binding between S1 and N1 with their n subassemblies. It also activates the v subassembly of S1, which in turn activates a ‘prediction node’ V, predicting the occurrence of a verb of the main clause. In (b) the control network, based on the combination of the prediction node V and the occurrence of the verb *sees*, activates the v subassembly of V1, which results in the binding of S and V. It also activates the t subassembly of V1 and a prediction node of a theme (T),

anticipating an object of the verb. In (c), the combination of the T node and the noun *dog* results in the binding of V1 and N2 with their t subassemblies.

In [14] we trained a control network with a set of training sentences with different syntactic forms (e.g., subject, theme, relative clause, complement clause). The network could then control a substantially larger set of sentences in which these syntactic forms were combined in different ways. For example, the control network could produce the correct binding operations in sentences with arbitrary deep (center) embeddings. However, the binding process in the NBA starts to break down with these sentences after one or two embeddings, due to competitions with the connection matrices involved [1, 14]. Hence, the NBA could account for both competence and performance aspects of language processing.

4.1 Simulating neuronal activation during sentence processing

Brain activation observed with fMRI indicates that neuronal activation increases for longer and more complex sentences [15]. Observations with intracranial electrodes provided more detailed information [16]. In particular, when (sub) phrases in a sentence are formed (or ‘merged’ in linguistic terms) neuronal activity first increases but then briefly decreases.

Figure 7 shows the activity in the NBA when the sentence *Ten sad student of Bill Gates will move* from [16] is processed (taking *Bill Gates* as a single noun). Figure 7 (left) shows the basic structure of the sentence in the NBA. Figure 7 (right) shows the sum of all activity (Total) in the NBA, and the sum activities of the main assemblies (MA), the subassemblies (SA), and the WM populations in the connection matrices. All assemblies and gating circuits are modeled with Wilson Cowan population dynamics [17]. MA and SA assemblies have reverberating activity as well (until they are inhibited). Activation values are normalized to fit in one plot.

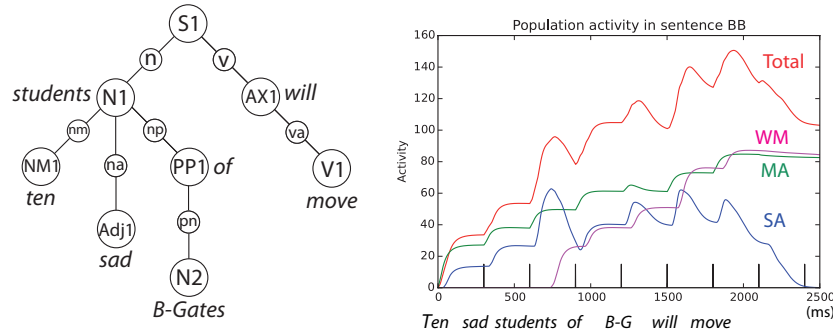


Figure 7: NBA activity in processing *Ten sad student of Bill-Gates will move*.

The sum of all activity resembles the patterns observed in [16]. For example, after *Ten sad students* (forming a phrase) activation increases but then drops, as with other phrases in the sentence. This is in particular due to the subassemblies, which are inhibited by feedback from their connection matrix when a binding has succeeded. But overall activation increases with sentence length and complexity [15].

Figure 8 illustrates the neural activity in the NBA when it resolves an unproblematic ambiguity (UPA) presented in [2]. Here, it is UPA13a (*Bird found in room died*) vs UPA13b (*Bird found in room debris*). The NBA resolves ambiguities by a competition between conflicting WM populations [18]. The NBA structure of both sentences and the conflicts (red dashed lines) are illustrated in the figure (left).

The assumption is that UPA13a is the basis interpretation. So, *found* is the verb of a relative clause bound to *bird*. Simulation shows that this incremental sentence processing proceeds without problems. So, the processing of UPA13b starts in the same way. The figure shows the individual WM activations obtained with UPA13b, which reveal the bindings achieved (the sum of this activity would result in a plot as in Figure 7). When *bird* (N1) is presented it binds to S1, indicated by the black line S1-N1 representing the WM binding population. After the initial activation by the word, this population remains active at a sustained level, so the binding remains intact during (and after)

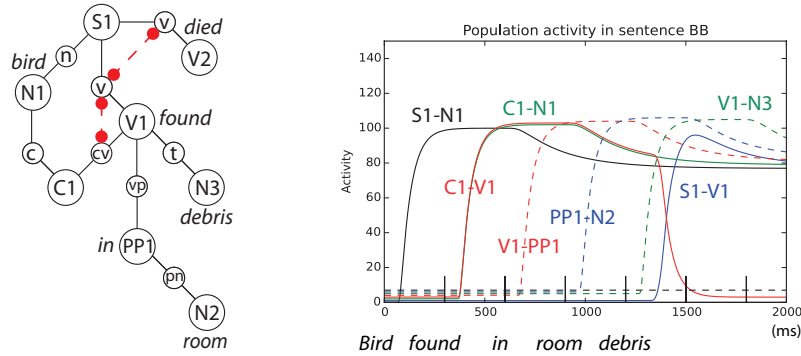


Figure 8: NBA activity in processing *Bird found in room debris*.

the sentence processing. The combination *bird found* is interpreted as *found* (V1) being the verb of a relative clause bound to N1 (as in UPA13a). This binding succeeds, as indicated by the WM populations C1-N1 (green line) and C1-V1 (red line). The bindings of *found in the room* proceed without problems, indicated by the WM populations V1-PP1 (red dash) and PP1-N2 (blue dash).

However, the last word *debris* provides a different interpretation of *found*. Instead of the verb of a relative clause it is the main verb of the sentence. This interpretation initiates the activation of the WM population V1-N3 (green dash), which represents the binding of *debris* as theme of *found*. It also initiates the activation of the WM population S1-V1 (blue line) which represents the binding of *found* as the main verb. This binding is in conflict with the binding C1-V1 (red line). The competition between the WM populations results in the decline of the activation of C1-V1, which resolves the binding conflict. Hence, ambiguity resolution in the NBA results from dynamic competitions in the architecture [18].

Alternatively, UPA13b could be the basis interpretation. Then *found* is the main verb directly. In that case *died* causes a conflict in UPA13a. However, simulations show that this conflict cannot be resolved because V1 blocks the binding of V2 (*died*) to S1. Also, V1 cannot bind anymore as verb to a relative clause of *bird* (N1), because the activations of N1 and V1 have been or are inhibited by other main assemblies (N2 and V2 respectively). Hence, this ambiguity cannot be resolved by the NBA. However, this ambiguity is similar to the classic garden path sentence *The horse raced past the barn fell* [19], which indicates that the NBA can account for both ambiguity resolution and garden path effects [18].

5 Conclusions

Neural Blackboard Architectures (NBAs) can account for variable binding and (novel) combinatorial processing in neural terms. They also satisfy a number of characteristics of brain representation and processing. Conceptual representations are always in situ, also when they are a part of a combinatorial structure. Hence, combinatorial structures are content addressable. This allows their direct activation based on input information. NBAs also provide the connection paths that are needed to account for the production of behavior. The sentence NBA can account for observed patterns of neuronal activation during sentence processing. It can also handle linguistics phenomena such as ambiguity resolution and can account for garden path sentences by the dynamic competitions in the architecture.

Acknowledgments

The work of Frank van der Velde was funded by the project ConCreTe. The project ConCreTe acknowledges the financial support of the Future and Emerging Technologies (FET) programme within the Seventh Framework Programme for Research of the European Commission, under FET grant number 611733. The research of Marc de Kamps has received funding from the European Union Seventh Framework Programme (FP7/2007-2013) under grant agreement no. 604102 (HBP) (Ref: Article II.30. of the Grant Agreement).

References

- [1] van der Velde, F. & de Kamps, M. (2006). Neural blackboard architectures of combinatorial structures in cognition. *Behavioral and Brain Sciences*, 29, 37-70.
- [2] Lewis, R. L. (1993). *An Architecturally-based Theory of Human Sentence Comprehension*. Thesis Carnegie Mellon University, Pittsburgh, PA.
- [3] Hebb, D. O. (1949). *The organisation of behaviour*. New York: Wiley.
- [4] van der Velde, F. & de Kamps, M. (2011). Compositional connectionist structures based on in situ grounded representations. *Connection Science*, 23, 97-107.
- [5] van der Velde, F. (2015). Communication, concepts and grounding. *Neural networks*, 62, 112 - 117.
- [6] Newell, A. (1990). *Unified theories of cognition*. Cambridge, MA: Harvard University Press.
- [7] Van der Velde, F. & de Kamps, M. (2001). From knowing what to knowing where: Modeling object-based attention with feedback disinhibition of activation. *Journal of Cognitive Neuroscience*, 13(4), 479-491.
- [8] Feldman, J. (2013). The neural binding problem(s). *Cognitive Neurodynamics*, 7, 1-11.
- [9] Van der Velde, F. & de Kamps, M. (2006). From neural dynamics to true combinatorial structures (reply to commentaries). *Behavioral and Brain Sciences*, 29, 88-108.
- [10] van der Velde, F. & de Kamps, M. (2015). The necessity of connection structures in neural models of variable binding. *Cognitive neurodynamics*, 9, 359-370. DOI 10.1007/s11571-015-9331-7
- [11] Amit, D. (1995). The Hebbian paradigm reintegrated: Local reverberations as internal representations. *Behavioral and Brain Sciences*, 18, 617-657.
- [12] Shanahan, M. (2010). *Embodiment and the inner life*. Oxford: Oxford University Press.
- [13] van der Velde, F. & de Kamps, M. (2011). Development of a connection matrix for productive grounded cognition. *IEEE International Conference on Development and Learning (ICDL)* (Vol 2), 1-6. DOI: 10.1109/DEVLRN.2011.6037343
- [14] van der Velde, F. & de Kamps, M. (2010). Learning of control in a neural architecture of grounded language processing. *Cognitive Systems Research*, 11, 93-107.
- [15] Pallier, C., Devauchelle, A-D. & Dehaene, S. (2011). Cortical representation of the constituent structure of sentences. *PNAS*, 108, 2522-27.
- [16] Nelson, M. J., El Karoui, I., Rangarajan, V., Pallier, C., Parvizi, J., Cohen, L., Naccache, L. & Dehaene, S. (2014). Constituent structure representations revealed with intracranial data. *Society for Neuroscience Annual Meeting*. Washington, DC, USA.
- [17] Wilson, H. R. & Cowan, J. D (1972). Excitatory and inhibitory interactions in localized populations of model neurons. *Biophysical Journal*, 12, 1-24.
- [18] van der Velde, F. & de Kamps, M. (2015). Ambiguity resolution in a Neural Blackboard Architecture for sentence structure. In Tarek R. Besold & Kai-Uwe Kühnberger (eds.), *Proceedings of the KI 2015 Workshop on Neural-Cognitive Integration*. Dresden, Germany.
- [19] Bever, T. G. (1970). The cognitive basis for linguistic structures. In Hayes, J. R., (ed.) *Cognition and the Development of Language*. New York: Wiley.