European Review, Vol. 16, No. 4, 483–495 © 2008 Academia Europæa
doi:10.1017/S1062798708000409 Printed in the United Kingdom

Fuster's Cherries and Linguistic Trees

HELMUT SCHNELLE

Rochstrasse 4, D-10178 Berlin, Germany. E-mail: helmutschnelle@t-online.de

Bridging the gap between linguistic structure representations and neurocognitive representations is a difficult challenge. This article presents an outline of how a formally specified system of constituent structure grammar could be translated into a distributed hierarchy network of associated modules. A set of syntactic constituent rule units could be reinterpreted as complex neuronal modules, which interactively generate momentary binding associations. The iterating binding activity in the network corresponds to the syntactic structure representation of a given sentence.

'Cherries turn red as they get ripe'. The neuroscientist J. Fuster chose this sentence to exemplify hypothetical interactions between networks of posterior and frontal cortices in the construction of a simple sentence by a highly schematic diagram in Figure 1.[1] The linguist is far from being satisfied by this representation. The core of morpho-syntactic construction should rather be presented by a schematic tree structure like that of Figure 2.[2] It shows how the syntactic properties of parts of the sentence are related and categorized, where it is obvious to each linguist that the schema represents at best some core aspects that need to be complemented by further linguistic explanations. Similarly the neuroscientist Fuster explains in his book how other parts of the cortex interact with those presented in the diagram. These interactions should be taken into account. Given the different representational frameworks it is difficult for the analysts of neurocognition and of linguistics to come to a mutual understanding. The neuroscientist Fuster's cherries grow on trees in a complex of electronic orchard networks, whereas the linguistic grammarian concentrates on generatively ordered arrangements of tree constructs. I now want to show how both could learn from each other.

Let me start again: language is easy. Everybody speaks one and knows how to use it. In order to speak a language, having learned it as a child is sufficient.

https://doi.org/10.1017/S1062798708000409 Published online by Cambridge University Press

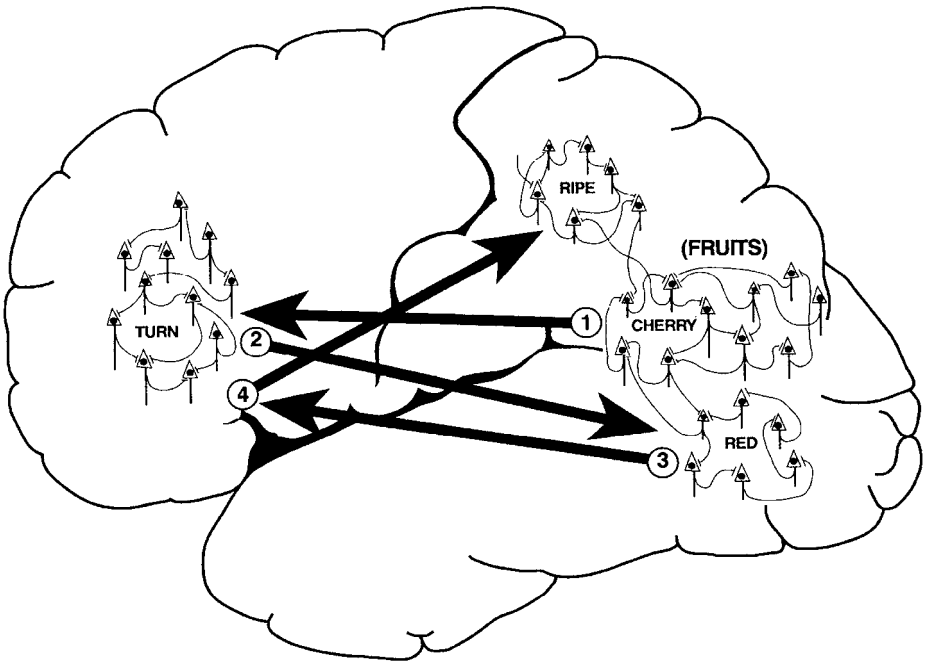


Figure 1. Cherries turn red as they get ripe. Reproduced by kind permission from J. M. Fuster (2003).[3]

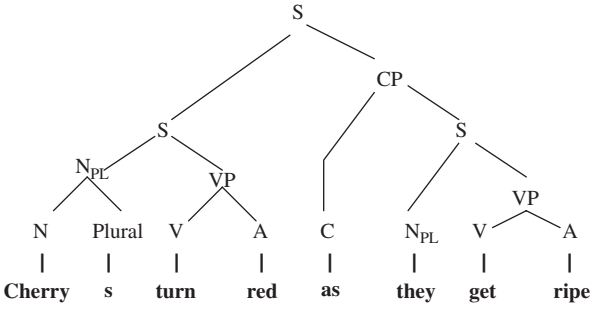


Figure 2. Constituent tree representing the syntactic structure of the example sentence

On the other hand, language is complicated – it is difficult to explain why a language is as it is, in its sounds, its words, and its sentences. The difficulty is obvious to everyone who wants to learn a foreign language.

Difficult things need explanation. Explaining the knowledge of language is the task of linguistics. After several thousand years of studies, the explanations of

https://doi.org/10.1017/S1062798708000409 Published online by Cambridge University Press

structures of grammars and lexica, which obviously underlie the speakers' and hearers' knowledge, are rather advanced today. The normal perspective is that linguistic explanation contributes to characterizing a central component of the human mind.

However, in the present understanding of most linguists, it must also be acknowledged that explaining the mental knowledge of a language is not enough. If the speakers' or hearers' knowledge were not organized in their brains, they could not apply it, as various types of brain lesions prove. As a consequence, the science of language should ultimately explain both mental knowledge of language and the brain's organization of language. It should also explain how mind and brain correspond in the domain of language.

On the other hand, the discussion of syntactic, morphological and lexical phenomena does not explain in any way at which locations in the brain the proper temporal order and phrasal grouping of morphemes and words are organized in the acts of perception and articulation. In other parts of his book,[1] Fuster gives some general explanations how the word activities represented in Figure 1 are integrated in interactive processing influenced by the activity of frontal and prefrontal brain. These are explanations about which the knowledge concentrated linguists remain silent. Linguists are satisfied with those aspects of mental knowledge that represent empirically justified grammatical and lexical structures in terms of notational configurations.

However, as acknowledged, mental knowledge analysis and neuronal brain analysis should not remain separated. In the present research situations of the two disciplines, we cannot do more than propose and study models that seem to be plausible for bringing structures and organizations of the linguistic mind and the neurocognitive brain into correspondence. Even this is a difficult task, in particular when we want to correlate techniques of formal symbolic structure representations with systematic descriptions of connection-based organization in cognitive neuroscience!

We must expect that perspective correspondence is more difficult to model than the similar task of relating software and hardware processes in computer science. The 'software' and 'data structures' that appropriately constitute the functional mind must be far more tightly bound up with the nature of the neuronal network 'hardware'. This certainly holds for the relation between mental linguistic structure and biological brain structure based neuronal organization.

Most linguists are sceptic. Typical is Chomsky's argument that the gap between formalist linguistic theories and neuronal descriptions of language organization cannot be bridged:[4] Psychologists have long realized that a description of what an organism does and a description of what it knows can be very different things. A formal study of grammar models is a mental approach to clarify the knowledge achieved by the mature speaker of a language. It is not concerned with the

https://doi.org/10.1017/S1062798708000409 Published online by Cambridge University Press

description of what a speaker or hearer does in his acts of speaking or understanding, nor in how brains organize these acts. However, it must be assumed that the neural organization somehow represents knowledge. The basic question remains: how is the formalized mental knowledge available in the neural brain. The sceptical stance seems to be justified: one can scarcely hope to develop a sensible theory of the actual organization of language before the basis of a serious and far-reaching formal account of what a language user knows has been completed.

I do not share this view. Instead I believe that the best approach to meet complex challenges consists of gradually developing and studying stages of complexity in well-founded working models. This was, by the way, also Chomsky's solution when in the 1960s he developed mathematical linguistics, and this must be the solution in studying empirically-based working models of neuronal organization. At the time, Chomsky selected a simple type of formal model: constituent structure grammar.[5,6] It was clearly understood as an initial framework for his studies, and is indeed a good choice for my studies also.[7] It seems to be useful in exercises helping us to develop expertise in mind–brain correspondence modelling before we are able to solve more complex and comprehensive solutions.

In fact, Chomsky's task was also translational. He wanted to relate mere collections of intuitive linguistic knowledge or empirically observed facts to strictly ordered and systematically explained system descriptions, applying mathematical structure models. What he was looking for was a mathematically formalized grammar, which would define a mathematical theory of syntax providing the core component for a more comprehensive mathematical linguistics that would satisfy the challenges of the discipline. He knew that his pre-mathematical intuitions of a linguistically more comprehensive and more adequate grammar, his understanding of a transformational grammar, were too difficult to be solved in a single swoop. As a mathematical form of grammar structure, he therefore concentrated on the formal understanding of constituent structure grammars, a particularly simple class of the grammar framework. Although he explicitly acknowledged the shortcomings of this type of grammar he wrote a complete chapter about the formal properties of grammars in which he emphasized that there are very good reasons why the formal investigation of the theory of constituent structure grammars should be intensively pursued.[8] Indeed, it does succeed in expressing certain core aspects of grammatical structure.

Empirical foundation of neuronal network constructions: cortical hierarchies and cortical modules

As there is scepticism concerning the availability of sufficiently powerful knowledge of neurocognitive modelling, I shall refer to two empirically established

https://doi.org/10.1017/S1062798708000409 Published online by Cambridge University Press

structure frameworks for neurocognitive modelling, namely distributed connectivity hierarchies and the cortical modules.[9] The modules are operative units in which several hundred neurons are combined in clusters, each extending approximately over half a millimetre. Moreover, during the growth process of an individual, each module develops an internal circuitry. The specific circuitry results from experience and is adapted to the specific task in which the module must carry out its interaction with other modules in the network connecting distant modules. Since the distant modules interact, they form a possible module interaction association. The distributed connectivity hierarchy of module associations constitutes the complete system of module external connectivity. Since the internal module circuitry organizes the required specificity for each location, the association actively interacts, task-specifically, in generating the appropriate functions of the distributed connection hierarchy.[7]

I now want to give an outline of how a formally specified system of constituent structure grammar could be reinterpreted as a distributed hierarchy network of associated modules and how each syntactic constituent rule unit could be reinterpreted as a complex neuronal module. Specific combinations of rules for generating a syntactic sentence representation would be translated into a module association of rules. Such modules are usually composed of functionally homogeneous cellular aggregates made up of coalesced mini-columns.[10]

The discussion of the principles of linguistic structure and neuronal organization will now be given preference. I shall start by introducing some essential characteristics of hierarchies and modules. In this way I want to refute the assumption that our knowledge about neurocognitive frameworks might not be sufficient for model development.

It is the notion of hierarchy complex that suggests the idea of correspondence of linguistic tree systems and neuronal hierarchies of perception and executive action. How should we spell out the details of constitution arrangement (a) of descriptively denoted formal syntax categories, as for instance in the form of the schematic tree of Figure 2, or, (b) of organizationally interactively operating neuronal modules arranged in a hierarchy? Their formats of representation are certainly different. Trees are static configurations passively composed of rule applications. An association of potentially interacting modules, the typical complex in a neurocognitive hierarchy – which we shall also call cognit^3 – is a process in which the interaction of the mutually fitting distributed module activities results in a self-organized binding group as a momentary structure representation pattern. Trees are passive constitution patterns; cognits, that is potentially interactive associations, are self-organizing activity configuration binders. Part of the translation challenge is the mapping of the linguists' rule-determined generator system of passive configurations into a neurocognitive generator system consisting of a collection of self-determining modules whose

https://doi.org/10.1017/S1062798708000409 Published online by Cambridge University Press

collective activation in the appropriate situation and context develops into a fixed attractor activity fitting to the situation.

In both cases there are convergent branching constitutions. Should we succeed in translation, a syntactic constituent structure tree could perhaps correspond to a cognit of a neuronal constituent structure network, and a translation of the system of the constituent structure rules would lead to a model of a neuronal syntax network that would generate each of the possible cognits contained in the network, whenever situation and context are given. Linguists often believe that dynamic systems of this type could never be as powerful as formalist descriptions. This view has been refuted by showing that, in principle, interactive unit systems are as powerful as Turing machines.[11]

Neuronal hierarchies

Let us begin with the discussion of the characteristic properties of neuronal hierarchies observed in the perception and executive networks of the brain. After many years of research, the neuroscientists Felleman and Van Essen could show by neuroanatomical methods that visual perception organization in macaques and cats forms a distributed hierarchical system of interacting area modules.[12,13] Seventy-two organization areas were connected by 758 connections in macaques and by 1100 in cats. The hierarchy represents many stages of processing that are increasingly distant from primary sensory areas. The areas in the sensory cortex react to primitive properties of visual objects, whereas the higher areas represent properties resulting from association-determined combinations of simpler ones. The former ideas that these 'combination' processes operated in one direction only – bottom-up or top-down – have been refuted. It has been discovered that most connections in the hierarchy are reciprocal. It seems that only by iterated flows of processing in both directions can perception be represented with clarity. The required time span of iteration is about 100 ms, a time span that a person experiences as a moment, an instantaneous event.

There are certain differences with systems of linguistic tree hierarchies. In view of our translation goal it is important to register them precisely. Felleman and Van Essen emphasize that their lines are, in general, to be understood as reciprocal. Each connection of module x with module y is an axonal connection from x to y that is parallel to an axonal connection from y to x. Studying the networks presented by Felleman and Van Essen there is convergent branching – the connections from several lower nodes in the hierarchy come together in higher nodes. These connections converge from bottom to top or in the feed-forward direction. But there are also connections that diverge, thus appearing to converge in the top to bottom, or feedback direction. Moreover, there are collateral connections between nodes on the same level. Another point to be mentioned is that there is

https://doi.org/10.1017/S1062798708000409 Published online by Cambridge University Press

no indication in the graphic representation of the distributed neuronal hierarchy that the same area name recurs at several different nodes. Thus, there is no label recurrence, no repetition of label tokens. Moreover, the nodes in neuronal hierarchies are not exclusively combinational, as in formal constituent structure rule systems; some of them are instead alternative selecting choice nodes, nodes that are well known from coding systems.[14,15]

Neural hierarchies differ in practically all of these features from linguistic constituent structure trees. The lines in the tree graphs, as in Figure 2, are not understood as reciprocal; they only converge from bottom to top and do not represent diverging branches or collaterals. The possibility of symbol recurrence, i.e. the unlimited availability of symbol tokens, is of fundamental importance for constituent structure grammars, since the notion of strings is basic for their definition. This is formal linguistics heritage from algebra or algorithm theory. Thus, the lack of recurrence in neuronal hierarchy networks might prima facie seem to prevent formal linguistics in general. I shall return to discuss this point. Finally, constituent structure grammars do not use choice points explicitly, although they are implicitly present.

Over many years I have developed working models that translate a hierarchical system of rule-based linguistic structures into a neuronal organization hierarchy.[7] It finally turned out that even the principles of mental structure representation applied in formal syntax can be transformed into more appropriate cognit representation if the model introduces specifics of neuronal hierarchies, namely reciprocal connections, plurality of connection orientation types (convergence, divergence, and collaterals), choice points, and the elimination of token occurrences and string-based formalisms.

However, as already mentioned, introducing the last aspect makes it particularly difficult for formalists to understand a brain organization model of language. They might be perplexed when their object domain is rather represented by dynamic state-space descriptions as discussed by Jackendoff.[16] The idea that a specific structure could indeed be implied by operation circuitry internal to the associated modules and by their interaction circuitry is strange to linguists, who are used to conceiving structure descriptions in terms of configurations of symbol strings or of symbol trees generated as constituent branching or by a combination and merging of constituent units. They learned from algebra and the theory of algorithms that calculation is based on string operation. In my view, however, the change of perspective is absolutely necessary. Consequently, I shall try to explain the essentials of formal syntax in the simple case of a constituent structure grammar. Our intended change of perspective from a static rule-generated syntax structure to interactively active module networks, shall be explained by means of some examples that illustrate the principles and essentials.

https://doi.org/10.1017/S1062798708000409 Published online by Cambridge University Press

The essentials of constituent structure syntax

A context-free constituent structure grammar is a special case of a formally specified system that is based on notions of algebra and the theory of algorithms. In linguistics, the formalisms of this kind are considered as notational schemata of mentally specifiable structure insights. As already said, this type of grammar or, more specifically, of syntax, serves as an exercise frame, the 'trivium' before studying more complex and linguistically more adequate grammar representations. However, we must also say that this grammar type is not a mere construct. In a sense, basic varieties of languages are of this type. 'Basic varieties' are the kinds of pidgin languages that second-language speakers normally develop on their own when they are not explicitly taught.[17]

The general principles of formal grammar organization are as follows. For each sufficiently simple sentence the system's rule applications generate the sentence's syntactic structure. In a sub-class of constituent structure grammars, each generated representation is a formal tree known from mathematical graph theory, a schematic configuration with branching nodes labelled by syntactic category terms. Figure 2 above has presented an example. For our present purposes of illustration, I borrow Chomsky's early example.

(1) 'The man hit the ball.'

A primitive sentence indeed.[18] Its syntactic structure is presented in Figure 3.

Which is the appropriate understanding of a constituent structure rule system for English sentences? As usual, the rule system is presented as a collection of single rules each of which is presented by a configuration of formal symbols. Its character as a system derives from principles of combining rule applications for generating the coherent tree. For our example sentence, the fragment of syntax consists of merely three rules as in the grammar fragment (G1).

(G1) (1) S → NP - VP, (2) VP → V - NP, (3) NP → Det - N

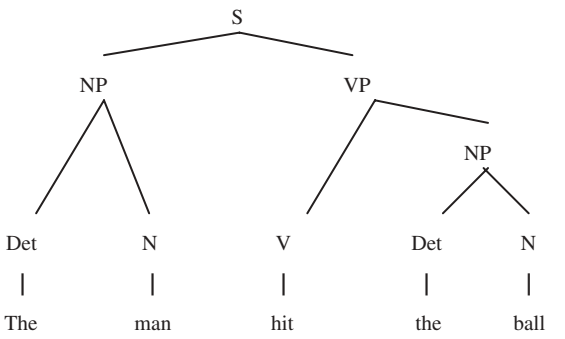


Figure 3. Chomsky's earliest constituent structure representation (1957)[18]

https://doi.org/10.1017/S1062798708000409 Published online by Cambridge University Press

The application of these rules consists of replacing symbols and symbol sequences, and generating the following syntactic derivation (D1) ending in a sequence of parts of speech categories. (Selecting, for each part of speech term, a lexical word may result in the sentence representation of the last line.)

(D1)	Start:	S
	Rule (1); S replacement:	NP - VP
	Rule (3); NP replacement:	Det - N - VP
	Rule (2); VP replacement:	Det - N - V - NP
	Rule (3); NP replacement:	Det - N - V - Det - N
	Sequence of word selections:	The man hit the ball

The reader will easily verify the similarity of this branching derivation with the configuration assigned to the cherries sentence presented above in Figure 2. In both of them the character of hierarchical arrangement is transparent: the part of speech sequence at the bottom (represented by terminal symbols), the non-terminal phrasal terms at intermediate positions and the convergence on the category S (sentence) on top. It is this structural feature that is important: The rule applications do not represent the speakers' or hearers' acts but are only formal relations in the generated tree schema.

Let me briefly introduce another schema for representing rules. The used representations are equivalent but more transparent than the classical form of the rule. The new variety will be helpful below. In it, a context-free constituent structure rule is represented by so-called treelets. Figure 4 represents the treelets of the grammar fragment (G1).

If rules are thus represented by treelets, the generation of the tree is not a replacement of a symbol by a string, but a 'clipping' of the treelets, since each treelet is already a tree fragment. The result is obviously the same as above, namely Figure 3. Obviously, the third treelet is clipped twice, first at the NP contained in the S-treelet, and second at the NP contained in the VP-treelet.

Translating trees, rules and rule complexes into neuronal modules

If all treelets that generate the syntactic structure of the sentence are clipped we might be tempted to simply assign to each node a local neuronal assembly, a module. Figure 5 would be the result. Unfortunately, Figure 5 is wrong. Our translation principle required definitely that each symbol (and not symbol token!)

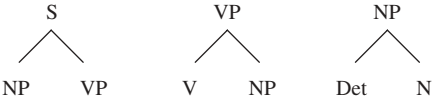


Figure 4. Three treelets to be clipped into the structure part of Figure 3

https://doi.org/10.1017/S1062798708000409 Published online by Cambridge University Press

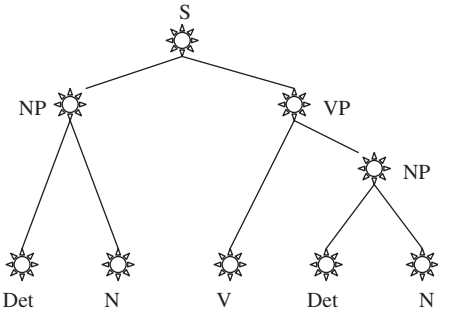


Figure 5. Clipped tree transformed in module arrangement

be translated into exactly one local neuronal assembly, or module. This principle excludes the repetition or recurrence of symbols; it does not allow several tokens of the same symbol! What should we do? We take the S-treelet and connect it to the NP-treelet, and then we take the VP-treelet and connect it – also(!) – to the NP-treelet. Instead of two occurrences of NP we now have two inputs at NP. But now the connection system is no longer a tree. It is as if we brought the two NP tokens of the ordinary clipping together and have them overlap at one NP-type position, the one to which we assign the neuronal module. The consequence is clear: the module that results from the two NP nodes must now have an internal circuitry that manages the different cases being either activated by a signal from S or by a signal from VP. Or if the tree should be interpreted bottom-up, the module should have some internal decision competence for appropriately sending the signals to S or to VP.

It seems that the situation becomes even more confusing. If the NP module receives top-down signals simultaneously from S and VP, what should be done in the next step? Generate 'the boy' and 'the girl'? But how? The order that was implicit in the clipped treelet arrangement, is lost since the two NP-treelet tokens overlap. The organization is much less transparent than might have initially appeared. The solution is the same as the one I had already in mind in my early proposals.[19]

In technical terms, the solution results in a development leading to a fixed attractor activity in a signal process represented in a state space. I have mentioned this possibility already. I shall not explain the solution in all of its details but the reader should at least get an idea about the dynamic process patterns that result from our translation into hierarchy distributed modules. Here are the essential ideas to be applied to the network with overlapping NP; these ideas rely on connectivity-determined signal flow. First, the static configuration of a tree is translated into a dynamic process that is generated and controlled by the internal circuitry of the modules. Second all lines are doubled into reciprocal pairs. Reciprocal connection was already a property of neuronal hierarchy, which we discussed above. We now have a configuration like Figure 6.

https://doi.org/10.1017/S1062798708000409 Published online by Cambridge University Press

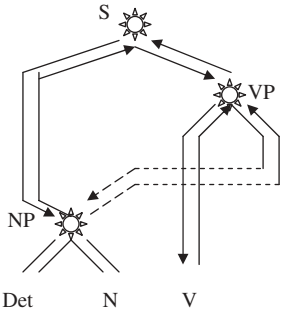


Figure 6. Module organized network with reciprocal connection network and non-recurrent arrangement of syntactic types

Let me now sketch the bottom-up situation. The reader should follow the description, by referring to Figure 6. The beginning of a sentence is heard and causes an expectation signal flowing from the S-module to the NP-module, and from there to the Det Module. In the NP module, the internal circuitry stores for a short period the fact that the input signal came from S. The Det module registers by neuronal analysis that the initial sound pattern 'the' is indeed a Det. This causes a signal sent along the bottom-up connectivity line (recall the reciprocal pair of connections) to the NP-module. Since this module corresponds to the NP rule its circuitry registers the completion of Det and concludes that an initiating expectation signal should be sent to the N-module. This registers that the sound 'boy' indeed satisfies the expectation and consequently sends a completion signal back to the NP-module. This registers completion, and asks about the state of registered direction store. A combination with the storage signal sends the completion signal forward to S. The module S has an internal circuitry corresponding to the conditions of the rule or treelet. It determines that after completed NP, VP should be sent an initiating signal. It causes VP to signal expectation to the V module. V registers that 'hit' is indeed a V, a fact that causes it to send the completion signal to VP. Corresponding to the rule, VP circuitry now sends an expectation signal to…? Well to NP! Thus, NP gets for a second time an expectation signal, but this time from VP. NP's circuitry short-term stores, having received a signal from the VP input connection, sends an expectation signal again to Det, which registers the presence of a Det, sends the completion signal to NP, which causes, corresponding to the rule, an expectation signal to NP, that consults the short-term store, and sends a completion signal to VP that in turn sends a completion signal to S. Since the connections in the modules are short, one can assume that any act mentioned takes about 25 ms, the sum about 400 ms, less than the normal speech length for the sentence.

Now, should the input sounds not have been very distinct, a repetition would immediately be initiated, thus leading to more 'distinctivity' by approximating a fixed attractor activity, a kind of periodic attractor.

https://doi.org/10.1017/S1062798708000409 Published online by Cambridge University Press

This was just a first sketch of how the static constituent structure tree is translated into a neuronal network model that makes use of typical properties known from neuronal networks: no token repetition, reciprocal connectivity and convergence. Discussing the fact that more complex rule system translations will lead to divergence, co-laterality and choice points must be postponed to my forthcoming book due to be published by Cambridge University Press. This will also present a more comprehensive account of the theoretical frameworks and the details of translating formalisms into dynamic networks and show how static linguistic trees may be planted into a cherry orchard of temporal attractor processes.

References and Notes

1. J. M. Fuster (2003) *Cortex and Mind-Unifying Cognition Figure* (New York: Oxford University Press), p. 209.
2. This figure resulted from discussions with my colleague M. Bierwisch.
3. J. M. Fuster (2003) *Cortex and Mind-Unifying Cognition Figure* (New York: Oxford University Press), pp. 14, 96.
4. N. Chomsky (1963) Formal properties of grammars. In: R. D. Luce, R. R. Bush and E. Galanter (eds) *Handbook of Mathematical Psychology, Vol. II*, Chapter 12, p. 326.
5. N. Chomsky (1963) Introduction to formal analysis of natural languages. In: R. D. Luce, R. R. Bush and E. Galanter (eds) *Handbook of Mathematical Psychology, Vol. II*, Chapter 11, pp. 290–325.
6. For more details see below, the section about the essentials of constituent structure syntax.
7. H. Schnelle (1996) *Die Natur der Sprache- Die Dynamik der Prozesse des Sprechens und Verstehens* (Berlin, NewYork: W de Gruyter), pp. 294–298. H. Schnelle (1981) Elements of theoretical net-linguistics. *Theoretical Linguistics*, **8**, 67–100.
8. N. Chomsky (1963) Formal properties of grammars. In: R. D. Luce, R. R. Bush and E. Galanter (eds) *Handbook of Mathematical Psychology, Vol. II*, Chapter 12, p. 325.
9. V. B. Mountcastle (1998) *Perceptual Neuroscience – The Cerebral Cortex* (Cambridge MA: Harvard University Press) Chapters 7 and 10, esp. p. 268, Figure 10.6.
10. J. M. Fuster (2003) *Cortex and Mind-Unifying Cognition Figure* (New York: Oxford University Press), p. 64.
11. R. Penrose (1989) *The Emperor's New Mind – Concerning Computers, Minds and the Laws of Physics* (New York: Oxford University Press), p. 171.
12. V. B. Mountcastle (1998) *Perceptual Neuroscience – The Cerebral Cortex* (Cambridge MA: Harvard University Press), p. 268, Figure 10.6.
13. D. J. Felleman and D. C. Van Essen (1991) Distributed hierarchical processing in the primate cerebral cortex. *Cerebral Cortex*, **1**, 1–47.

https://doi.org/10.1017/S1062798708000409 Published online by Cambridge University Press

14. N. Chomsky and G. A. Miller (1963) Introduction to the formal analysis of natural language. In: R. D. Luce, R. R. Bush and E. Galanter (eds) *Handbook of Mathematical Psychology, Vol. II*, Chapter 11, p. 278.
15. B. J. Baars and N. M. Gage (2007) *Cognition Brain, and Consciousness* (London: Elsevier, Academic Press), p. 319.
16. R. Jackendoff (2002) *Foundations of Language* (New York: Oxford University Press), p. 24ff., p. 56.
17. W. Klein and C. Purdue (1997) The basic variety: or couldn't language be much simpler. *Second Language Research*, **13**, 301–347.
18. N. Chomsky (1957) *Syntactic Structures* (s'Gravenhage: Mouton).
19. H. Schnelle (1964) Programmieren linguistischer Automaten. In: K. Steinbuch and S. W. Wagner (eds) *Neuere Ergebnissse der Kybernetik* (München: Oldenbourg), pp. 109–136.

About the Author

Helmut Schnelle is emeritus professor of linguistics. After many years of studies in mathematical and computational linguistics his main interests now lie with bridging the gap between language structure and neurocognition dynamics, and related problems of the philosophy of mind and brain.

https://doi.org/10.1017/S1062798708000409 Published online by Cambridge University Press