P. F. DOMINEY

A MODEL OF LEARNING SYNTACTIC COMPREHENSION FOR NATURAL AND ARTIFICIAL GRAMMARS

Abstract. One important function of syntactic analysis of speech is the assignment of thematic roles to noun phrases. Thematic roles of noun-phrases in canonical sentences are assigned in "default" orderings (e.g. Agent Object Recipient in English for the canonical sentence "John gave the ball to Mary."). This ordering is transformed in non-canonical sentences (e.g. "The ball was given to Mary by John."), whose thematic role assignment is guided, in part, by function items (e.g. prepositions "to" and "by"). Agrammatic patients are impaired in the syntactic comprehension of non-canonical sentences. It is not clear whether this is due to a failure to process function items, a failure to assign thematic roles, or both.
We have recently studied artificial grammar learning in a recurrent network model in which serial surface structure, and abstract transformational rules are represented by separate systems. We now examine the behavior of this dual system model in syntactic comprehension of canonical and non-canonical sentences. Function items are represented by the "surface" system and guide the application of transformational rules that are represented by the "abstract" system. After learning the structural regularities of the target language, the model predicts that failure in either of these systems, i.e. in the representation of function items or in the representation of syntactic knowledge will impair syntactic comprehension. The model also predicts that the expression of these impairments should not be restricted to performance with natural grammars but should also be manifest in artificial grammar conditions. With respect to this second prediction, we have recently confirmed that agrammatic patients indeed fail to process non-canonically ordered sentences both in natural and artificial grammars. In the artificial grammar sequences, this failure is purely related to application of the transformation, as no function items are involved in the task. The relevance of these results to current linguistic theory will be discussed.

1. INTRODUCTION

Thematic role assignment is a central function of syntactic analysis in sentence comprehension. The research objective in this chapter is to attempt to outline how this capability could be implemented in a neurophysiologically feasible neural network model that learns the structural regularities of the target language. We have recently developed a neural network model that simulates the 0-8 month infant's sensitivity to serial (Saffran, Aslin, & Newport, 1996), temporal/rhythmic (Nazzi, Bertoncini, & Mehler, 1998) and abstract structure (Marcus, Vijayan, Bandi Rao, & Vishton, 1999) of language (Dominey & Ramus, 2000). A key aspect of this model is that the infant's sensitivity to serial and temporal structure in language can be simulated by a "temporal recurrent network" (TRN) that relies on recurrent connections to encode context, and an associative memory that binds context representations to behavioral responses, while sensitivity to abstract rule-like structure requires the additional representation capabilities.

E. Witruk, A. D. Friederici, & T. Lachmann (Eds.), Basic Functions of Language, Reading, and Reading Disability, 61—77.

Such additional representational capabilities were initially developed to simulate human performance in the learning and transfer of abstract structure in sensorimotor sequence learning (Dominey, Lelekov, Ventre-Dominey, & Jeannerod, 1998). The two sequences *ABCBAC* and *DEFEDF* can be said to have different surface structures (i.e. serial orders), since they are made up of different elements. The share, however, the same abstract structure *123213*. Humans can learn this kind of abstract structure and demonstrate transfer of this knowledge to new sequence that are constructed from the same abstract structure (isomorphic sequences) e.g. *RZPZRP*, *HDIDHI* etc.

We demonstrated that while our TRN model can learn surface structure, it fails to learn abstract structure because it is incapable of representing the structural relations between repeating elements. In order to do so, we were required to add a continuously updated short term memory of the 7+/- 2 previous elements. Each new element in a sequence is compared with the contents of the STM, thus allowing detection and representation of abstract structure. For example, the abstract structure of sequence *ABCBAC* is represented as *u,u,u,n-2,n-4,n-3*, where "u" indicates "unpredictable, or no repetition", and "n-2" indicates "repetition of the element 2 places behind," etc. The thus updated model for abstract structure was able to simulate adult performance in serial reaction time tasks that required abstract structure learning (Dominey et al., 1998). Interestingly, the same model also simulates babies' sensitivity to the difference between auditory sentences based on the abstract structures *AAB* and *ABB*, as demonstrated by Marcus et al. (1999).

The current research tests the hypothesis that by allowing the models for surface and abstract structure to interact, we can develop a capability to simulate adult performance in thematic role assignment with dissociable processing of closed class items by the surface model, and open class items by the abstract model, as described below.

2. THEMATIC ROLE ASSIGNMENT

A central function in syntactic analysis or syntactic comprehension is the assignment of thematic roles to noun phrases in sentences. In a simplified manner, we can consider that in languages like French and English, there is a default or canonical order in which thematic roles are assigned (e.g. "Agent Object Recipient" in English for the canonical sentence "John gave the ball to Mary.") However, in non-canonical sentences (e.g. "The ball was given to Mary by John."), this ordering is transformed, and thematic role assignment is guided, in part, by function items (e.g. closed class words including prepositions "to" and "by", grammatical morphemes, etc.).

The ability to assign thematic roles has been quantified in different clinical tests used to asses agrammatism in aphasic patients suffering from lesions of the left cortical hemisphere. A well know version developed by Caplan, Baker, and Dehaut (1985) consists of 9 sentence types of varying syntactic complexity, five canonical and four non-canonical, and will serve as our target problem (see Table 1). Five sentences of each type are used for a total of 45 sentences. Sentences are real aloud to the patients in a pseudo-random order, and after each sentence, the subject should indicate by pointing at photographs "who did what to whom", indicating in canonical order the agent, object and recipient. Interestingly, a rather significant subgroup of these patients with left-hemisphere lesions demonstrate a deficit in thematic role assignment that is highly selective for non-canonical sentences. We note that the sentences are constructed so that no semantic interpretation can contribute to the role assignment, which must proceed entirely as guided by syntactic function items.

*Table 1. Nine Sentence Types as Specified in Caplan et al. (1985). Noncanonical Word Ordering Indicated by *.*

Two place verb sentences	Active (A)	The elephant hit the monkey.
	Passive (P)	The elephant was hit by the monkey. (*)
	Cleft-Subject (CS)	It was the elephant that hit the monkey
	Cleft-Object (CO)	It was the elephant that the monkey hit. (*)
Three-place verb sentences	Dative (D)	The elephant gave the monkey to the rabbit.
	Dative pass (DP)	The elephant was given to the monkey by the rabbit. (*)
Sentences with two verbs	Conjoined (C)	The elephant hit the monkey and hugged the rabbit.
	Subj-Obj Rel (SO)	The elephant that the monkey hit hugged the rabbit. (*)
	Obj-Subj Rel (OS)	The elephant hit the monkey that hugged the rabbit.

In order to realize such a task, a system should first be capable of distinguishing function words (or morphemes) from content words. Numerous behavioral and event-related brain potential studies indicate that indeed, adults process function and content words in a dissociated manner. The system must also be able to store the content words in a working memory, and then to access this memory in a non-standard order (i.e. different from the input order) guided by the function items. This capability to access the content words in the canonical order, guided by the function words, provides the basis for the assignment of thematic roles.

The dual-process model that we developed to simulate human performance in learning surface and abstract structure appears ideally suited for this task. Function items are represented by the "surface" system and guide the application of transformational rules that are represented by the "abstract" system. We can demonstrate that after training on a supervised version of the Caplan task in which the correct thematic role assignment is provided, the model can then perform the standard unsupervised task correctly, including the generalization to new sentences. The model can be rendered agrammatic either by disruption of the representation of the function items, or by disruption of the selection of nouns from the working memory based on the representation of the function words. Thus, the model predicts that failure in either of these systems, i.e. in the representation of function items or in the representation of syntactic knowledge will impair syntactic comprehension.

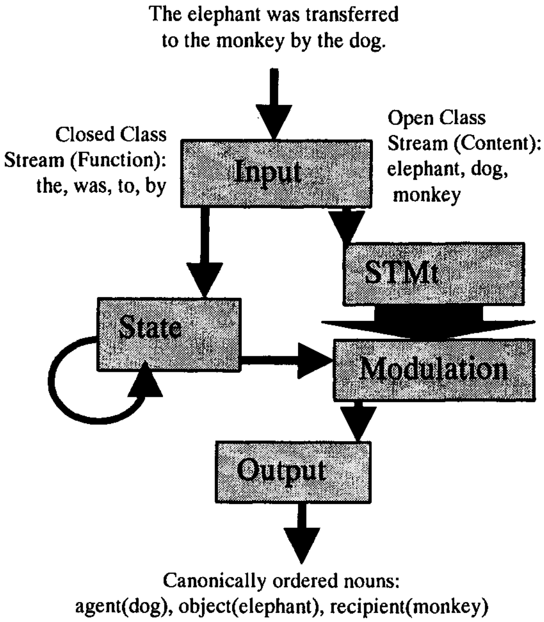

Figure 1. Syntactic comprehension model architecture.

3. OVERVIEW OF MODELING APPROACH

The model architecture is presented in Figure 1, and technical details are presented in the Appendix. In Figure 1, each of the labeled elements corresponds to a 5x5 array of leaky integrator neurons. A central premise of the model is that open and closed class items are processed by dissociable streams as suggested by studies of psycholinguistics, aphasia and language related ERPs (e.g. Friederici, 1985; Neville et al. 1993; Osterhout, 1997; Pulvermüller, 1995; Brown, Hagoort, & ter Keurs, 1999). This is reflected in the dual processing streams for closed and open class words. As the input sentence is read in, the recurrent network State encodes the ordered history of closed class words, and the short term memory for transformation (STMt) stored the open class words in the order of their arrival. Based on the syntactic context defined by the sequence of closed class words, State units activate Modulation units to access the contents of the STMt to retrieve the open class elements in their canonical order.

The model is constructed from ensembles of neuron-like units, each of which simulates the membrane potential of a neuron (or population of neurons) and generates an analog signal of firing/discharge rate. The model acquires the capacity to perform the syntactic comprehension task by learning that results in the modification of connection strengths between different neurons. The desired behavior is to reproduce human-like performance in Caplan's syntactic comprehension task in which sentences are presented, one word at a time to the model which then responds by emitting the nouns in their canonical order when presented with the remaining unassigned nouns.

The 9 sentence types from Caplan et al.'s (1985) syntactic comprehension task are re-coded so that each word corresponds to one element in the 25 element input array of the model. The words are categorized into closed and open classes. Closed class words enter a processing stream directed towards the recurrent network (State) which thus maintains an ongoing context of the sequence of closed class words presented for the sentence currently being processed. The open class words enter a processing stream by which they are stored in a working or short term memory. This continues until the final word of the current sentence is processed.

3.1 Training

During the training phase, a given sentence is presented, one word at a time in a temporal sequence as just described. Once the sentence is presented, the model is then provided with the open class nouns, one at a time, in their canonical order. In the Caplan protocol, this is equivalent to the experimenter herself pointing to the nouns in their canonical order after reading the sentence, i.e. showing the subject the correct responses. For canonical input sentences, the open class words are thus simply re-presented in the same order as in the initial input. For non-canonical input sentences, the nouns are transformed into their canonical order. During this learning phase, the model learns to associate the different patterns of closed class words (encoded in the neural activity of the State layer), with the appropriate (re)ordering of open class words.

For an example sentence, "The elephant was transferred to the monkey by the dog," after the sentence is presented, the model is then provided with the first canonically ordered thematic role, i.e. the agent which in this case is "dog." This matches with the contents of the 3rd element of the STM. A learning signal strengthens connections between the current set of active neurons in State (whose activity is due to the sequence of function words that was provided) and a modulatory neuron that will modulate that contents of the third element of the STM into the output. The result of such learning is that when the same type of sentence is presented, the active State units will activate the appropriate Modulation unit, directing the contents of the 3rd STM element to the output, thus providing the agent response.

As the supervised training example proceeds, the next canonically ordered element, the object "elephant" is provided. This matches with the first element of the STM, and the new State activity (which has been modified due to the match recognition) becomes associated with the modulation of the contents of the 1st STM element into the output, and so on for the recipient. The result of this training, for all 9 sentence types, is that activity in the closed class stream (encoded in the evolving activity of State during sentence presentation) allows the selection of elements from the STM in the appropriate learned canonical order. While clearly oversimplified from a linguistic perspective, the resulting representations are sufficient to allow the

http://www.springer.com/978-1-4020-7027-3

Basic Functions of Language, Reading and Reading
Disability
Witruk, E.; Friederici, A.D.; Lachmann, Th. (Eds.)
2002, IX, 377 p., Hardcover
ISBN: 978-1-4020-7027-3