

A Neurolinguistic Model of Grammatical Construction Processing

Peter Ford Dominey¹, Michel Hoen¹, and Toshio Inui²

Abstract

■ One of the functions of everyday human language is to communicate meaning. Thus, when one hears or reads the sentence, “John gave a book to Mary,” some aspect of an event concerning the transfer of possession of a book from John to Mary is (hopefully) transmitted. One theoretical approach to language referred to as construction grammar emphasizes this link between sentence structure and meaning in the form of grammatical constructions. The objective of the current research is to (1) outline a functional description of grammatical

construction processing based on principles of psycholinguistics, (2) develop a model of how these functions can be implemented in human neurophysiology, and then (3) demonstrate the feasibility of the resulting model in processing languages of typologically diverse natures, that is, English, French, and Japanese. In this context, particular interest will be directed toward the processing of novel compositional structure of relative phrases. The simulation results are discussed in the context of recent neurophysiological studies of language processing. ■

INTRODUCTION

One of the long-term quests of cognitive neuroscience has been to link functional aspects of language processing to its underlying neurophysiology, that is, to understand how neural mechanisms allow the mapping of the surface structure of a sentence onto a conceptual representation of its meaning. The successful pursuit of this objective will likely prove to be a multidisciplinary endeavor that requires cooperation between theoretical, developmental, and neurological approaches to the study of language, as well as contributions from computational modeling that can eventually validate proposed hypotheses. The current research takes this multidisciplinary approach within the theoretical framework associated with grammatical constructions. The essential distinction within this context is that language is considered to consist of a structured inventory of mappings between the surface forms of utterances and meanings, referred to as grammatical constructions (see Goldberg, 1995). These mappings vary along a continuum of complexity. At one end are single words and fixed “holophrases” such as “Gimme that” that are processed as unparsed “holistic” items (see Tomasello, 2003). At the other extreme are complex abstract argument constructions that allow the use of sentences like this one. In between are the workhorses of everyday language, abstract argument constructions that allow the expression of spatio-temporal events that are basic to human experience including active transitive (e.g., John took the car) and

ditransitive (e.g., Mary gave my mom a new recipe) constructions (Goldberg, 1995).

In this context, the “usage-based” perspective holds that the infant begins language acquisition by learning very simple constructions in a progressive development of processing complexity, with a substantial amount of early ground that can be covered with relatively modest computational resources (Clark, 2003; Tomasello, 2003). This is in contrast with the “continuity hypothesis” issued from the generative grammar philosophy, which holds that the full range of syntactic complexity is available in the form of a universal grammar and is used at the outset of language learning (see Tomasello, 2000, and comments on the continuity hypothesis debate). In this generative context, the universal grammar is there from the outset, and the challenge is to understand how it got there. In the usage-based construction approach, initial processing is simple and becomes increasingly complex, and the challenge is to explain the mechanisms that allow full productivity and compositionality (composing new constructions from existing ones). This issue is partially addressed in the current research.

Theoretical Framework for Grammatical Constructions

If grammatical constructions are mappings from sentence structure to meaning, then the system must be capable of (1) identifying the type of grammatical construction for a given sentence and (2) using the identified construction and its corresponding mapping to

¹CNRS UMR 5015, France, ²Kyoto University, Japan

extract the meaning of the sentence. Interestingly, this corresponds to Townsend and Bever's (2001) two steps of syntactic processing, that is, (i) parsing of the sentence into phrasal constituents and words (requiring access to lexical categories and word order analysis) and (ii) subsequent analysis of phrasal structure and long-distance dependencies by means of syntactic rules.

Figure 1A illustrates an example of an active transitive sentence and its mapping onto a representation of meaning. The generalized representation of the corresponding active transitive construction is depicted in Figure 1B. The "slots" depicted by the NPs and V can be instantiated by different nouns and verbs in order to generate an open set of new sentences. For each sentence corresponding to this construction type, the mapping of sentence to meaning is provided by the construction. Figure 1C illustrates a more complex construction that contains an embedded relative phrase. In this context, a central issue in construction grammar will concern how the potential diversity of constructions are identified.

Part of the response to this lies in the specific ways in which function and content words are structurally organized in distinct sentence types. Function words (also referred to as closed class words because of their limited number in any given language), including determiners, prepositions, and auxiliary verbs, play a role in defining the grammatical structure of a sentence (i.e., in specifying who did what to whom) although they carry little semantic content. Content words (also referred to as open class words because of their essentially unlimited number) play a more central role in contributing pieces of meaning that are inserted into the grammatical structure of the sentence. Thus, returning to our examples in Figure 1, the thematic roles for the content words are determined by their relative position in the

sentences with respect to the other content words and with respect to the function words. Although this is the case in English, Bates, McNew, MacWhinney, Devescovi, and Smith (1982) and MacWhinney (1982) have made the case more generally, stating that across human languages, the grammatical structure of sentences is specified by a combination of cues including word order, grammatical function words (and or grammatical markers attached to the word roots), and prosodic structure.

The general idea here is that these constructions are templates into which a variety of open class elements (nouns, verbs, etc.) can be inserted in order to express novel meanings. Part of the definition of a construction is the mapping between slots or variables in the template and the corresponding semantic roles in the meaning, as illustrated in Figure 1. In this context, a substantial part of the language faculty corresponds to a structured set of such sentence-to-meaning mappings, and these mappings are stored and retrieved based on the patterns of structural markers (i.e., word order and function word patterns) unique to each grammatical construction type. Several major issues can be raised with respect to this characterization. The issues that we address in the current research are as follows: (1) Can this theoretical characterization be mapped onto human functional neuroanatomy in a meaningful manner and (2) can the resulting system be demonstrated to account for a restricted subset of human language phenomena in a meaningful manner?

Implementing Grammatical Constructions in a Neurocomputational Model

Interestingly, this characterization of grammatical constructions can be reformulated into a type of sequence learning problem, if we consider a sentence as a sequence of words. Determining the construction type for a given sentence consists in analyzing or recoding the sentence as a sequence of open class and closed class elements, and then performing sequence recognition on this recoded sequence. Dominey (1995) and Dominey, Arbib, and Joseph (1995) have demonstrated how such sequence recognition can be performed by a recurrent prefrontal cortical network (a "temporal recurrent network" [TRN]) that encodes sequential structure (see also Dominey 1998a, b). Then corticostriatal connections allow the association of different categories of sequences represented in the recurrent network with different behavioral responses.

Once the construction type has thus been identified, the corresponding mapping of open class elements onto their semantic roles must be retrieved and performed. This mapping corresponds to what we have called "abstract structure" processing (Dominey, Lelekov, Ventre-Dominey, & Jeannerod, 1998). In this context, the reordering in the form-meaning mapping in Figure 1B can be characterized as the abstract structure ABCD-

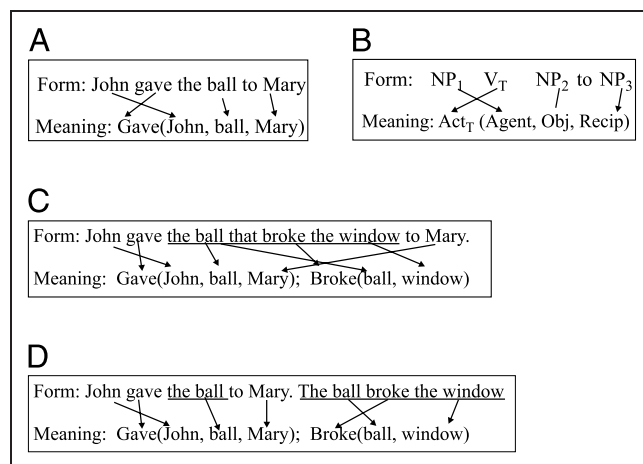


Figure 1. Grammatical construction overview. (A) Specific example of a sentence-to-meaning mapping. (B) Generalized representation of the construction. (C) Sentence-to-meaning mapping for sentence with relativized phrase. (D) Compositional sentence-to-meaning mapping for sentence with relative phrase in which the relative phrase has been extracted.

BACD where A-D represent variable slots. In order to accommodate such abstract structures, rather than representing sequences of distinct elements, we modified the recurrent network model to represent sequences of variables corresponding to prefrontal working memory elements (Dominey et al., 1998).

Concretely, from a developmental perspective, we demonstrated that the resulting abstract recurrent network (ARN) could simulate human infant performance in distinguishing between abstract structures such as ABB versus AAB (Dominey & Ramus, 2000) as described by Marcus et al. (1999). In the grammatical construction analog, these abstract structures correspond to the mapping of word order in the sentence onto semantic arguments in the meaning as illustrated in Figure 1. We thus demonstrated how the dual TRN/ARN system could be used for sentence comprehension. The sequence of closed class elements defining the construction was processed by the TRN (corresponding to a recurrent corticocortical network). Then, via modifiable corticostriatal synapses, the resulting pattern of cortical activity recovered the corresponding abstract structure for re-ordering the open class elements into the appropriate semantic argument order by the ARN (Dominey, Hoen, Blanc, & Lelekov-Boissard, 2003; Dominey, 2002).

The model essentially predicted a common neurophysiological basis for abstract mappings of the form BHM-HMB and form-to-meaning mappings such as “The ball was hit by Mary”: Hit(Mary, ball). Concretely, we predicted that brain lesions in the left perisylvian cortex that produce syntactic comprehension deficits would produce correlated impairments in abstract sequence processing. The first test of this prediction was thus to compare performance in aphasic patients on these two types of behavior. We observed a highly significant correlation between performance on syntactic comprehension and abstract structure processing in aphasic patients, as well as in schizophrenic patients (Dominey et al., 2003; Lelekov et al., 2000). We reasoned that if this correlation was due to a shared brain mechanism, then training in one of the tasks should transfer to improvement on the other. Based on this prediction, we subsequently demonstrated that reeducation with specific nonlinguistic abstract structures transferred to improved comprehension for the analogous sentence types in a group of aphasic patients (Hoen et al., 2003). In order to begin to characterize the underlying shared neural mechanisms, Hoen and Dominey (2000) measured brain activity with event-related potentials for abstract sequences in which the mapping was specified by a “function symbol” analogous to function words in sentences. Thus, in the sequences ABCxBAC and ABCzABC, the two function symbols x and z indicate two distinct structure mappings. We thus observed a left anterior negativity (LAN) in response to the function symbols, analogous to the LAN observed in response to grammatical function words during sentence processing

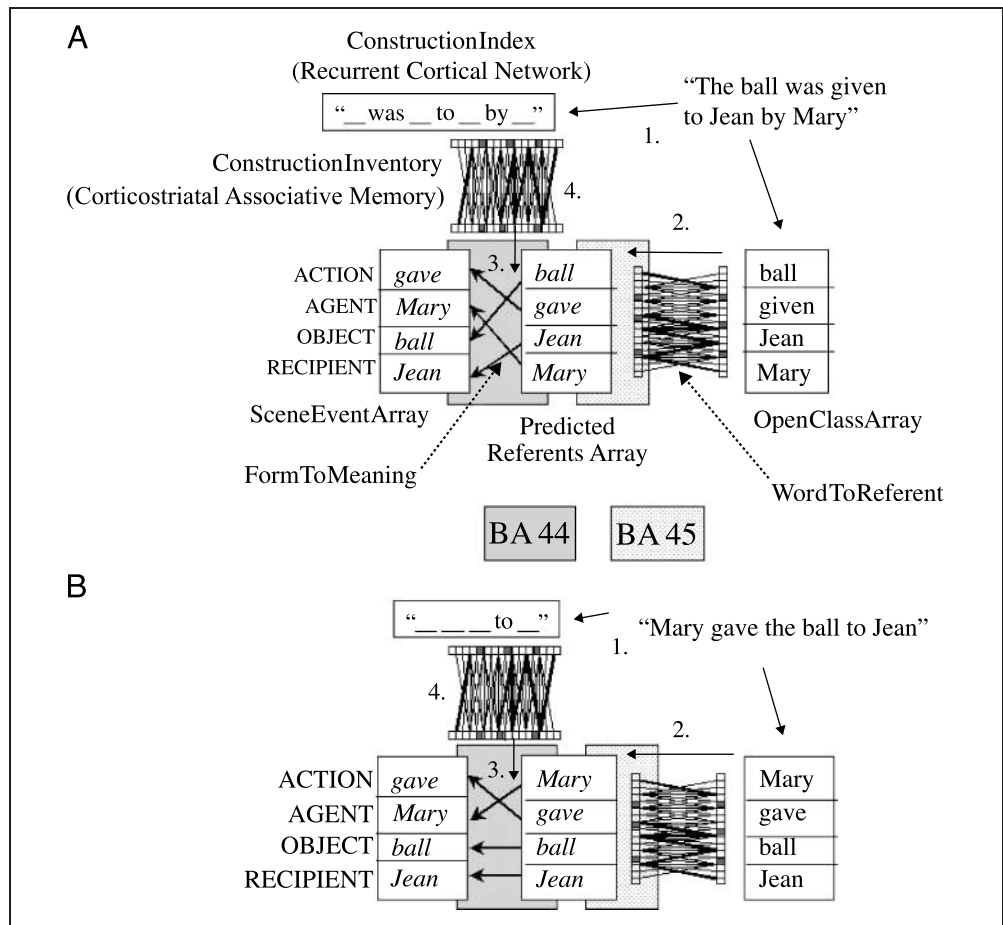
(Hoen & Dominey, 2000), again suggesting a shared neural mechanism. In order to neuroanatomically localize this shared mechanism, we performed a related set of brain imagery experiments comparing sentence and sequence processing using functional magnetic resonance imaging (fMRI). We observed that a common cortical network including Brodmann’s area (BA) 44 was involved in the processing of sentences and abstract structure in nonlinguistic sequences, whereas BA 45 was exclusively activated in sentence processing, corresponding to insertion of lexical semantic content into the transformation mechanism (Hoen, Pachot-Clouard, Segebarth, & Dominey, 2006).

This computational neuroscience approach allowed the projection of the construction grammar framework onto the corticostriatal system with two essential properties: first, construction identification by corticostriatal sequence recognition, and second, structure mapping based on the manipulation of sequences of frontal cortical working memory elements (Dominey & Hoen, 2006). Given the initial validation of this model in the neurophysiological and neuropsychological studies cited above, we can now proceed with a more detailed study of how this can contribute to an understanding of the possible implementation of grammatical constructions, as illustrated in Figure 2.

As the sentence is processed word by word, a process of lexical categorization identifies open and closed class words. This is not unrealistic, as newborns can perform this categorization (Shi, Werker, & Morgan, 1999), and several neural network studies have demonstrated lexical categorization of this type based on prosodic cues (Blanc, Dodane, & Dominey, 2003; Shi et al., 1999). In the current implementation, only nouns and verbs are recognized as open class words, with the modification of these by adjectives and adverbs left for now as a future issue. The meanings of the open class words are retrieved from the lexicon (not addressed here, but see Dominey & Boucher, 2005; Roy, 2002; Dominey, 2000; Siskind, 1996) and these referent meanings are stored in a working memory called the PredictedReferentsArray. The next crucial step is the mapping of these referent meanings onto the appropriate components of the meaning structure. In Figure 2, this corresponds to the mapping from the PredictedReferentsArray onto the meaning coded in the SceneEventArray. As seen in Figure 2A and B, this mapping varies depending on the construction type. Thus, the system must be able to store and retrieve different FormToMeaning mappings appropriate for different sentence types, corresponding to distinct grammatical constructions.

During the lexical categorization process, the structure of the sentence is recoded based on the local structure of open and closed class words, in order to yield a ConstructionIndex that will be unique to each construction type, corresponding to the cue ensemble of Bates et al. (1982) in the general case. The closed

Figure 2. Structure-mapping architecture. (A) Passive sentence processing: Step 1, lexical categorization—open and closed class words directed to OpenClassArray and ConstructionIndex, respectively. Step 2, open class words in OpenClassArray are translated to their referent meanings via the WordToReferent mapping. Insertion of this referent semantic content into the Predicted Referents Array (PRA) is realized in pars triangularis BA 45. Step 3, PRA elements are mapped onto their roles in the SceneEventArray by the FormToMeaning mapping specific to each sentence type. Step 4, This mapping is retrieved from ConstructionInventory (a corticostriatal associative memory) via the ConstructionIndex (a corticocortical recurrent network) that encodes the closed and open class word patterns that characterize each grammatical construction type. The structure mapping process is associated with activation of pars opercularis BA 44. In the current implementation, neural network associative memory for the ConstructionInventory is replaced by a procedural lookup table. (B) Active sentence. Note difference in ConstructionIndex and in FormToMeaning.



class words are explicitly represented in the ConstructionIndex, whereas open class words are represented as slots that can take open class words as arguments. The ConstructionIndex is thus a global representation of the sentence structure. Again, the requirement is that every different grammatical construction type should yield a unique ConstructionIndex. This ConstructionIndex can then be used as an index into an associative memory to store and retrieve the correct FormToMeaning mapping.

We have suggested that this mechanism relies on recurrent cortical networks and corticostriatal processing (Dominey et al., 2003). In particular, we propose that the formation of the ConstructionIndex as a neural pattern of activity will rely on sequence processing in recurrent cortical networks, and that the retrieval of the FormToMeaning component will rely on a corticostriatal associative memory. Finally, the mapping from form to meaning will take place in the frontal cortical region including BAs 44, 46, and 6. This corresponds to the SceneEventArray, consistent with observations that event meaning is represented in this BA 44 pars opercularis region, when events are being visually observed (Buccino et al., 2004), and when their descriptions are

listened to (Tettamanti et al., 2005). Hoen et al. (2006) provided strong evidence that for both the processing of grammatical and nonlinguistic structure processing rules, this frontal cortical region including BAs 44, 46, and 6 was activated (see Figure 3), indicating its role in linguistic and nonlinguistic structural mapping.

The proposed role of basal ganglia in rule storage and retrieval is somewhat related to the procedural component of Ullman's (2001, 2004, 2006) grammar processing model in which grammatical rules are encoded in specific (but potentially domain independent) channels of the corticostriatal system. Longworth, Keenan, Barker, Marslen-Wilson, and Tyler (2005) indicate a more restricted, non-language-specific role of the striatum in language in the selection of the appropriate mapping in the late integration processes of language comprehension. Neuropsychological evidence for the role of the striatum in such rule extraction has been provided in patients with Huntington's disease (a form of striatal dysfunction) that were impaired in rule application in three domains: morphology, syntax, and arithmetic (Teichmann et al., 2005). These data are thus consistent with the hypothesis that the striatum is




Figure 3. Comparison of brain activation in sentence processing and nonlinguistic sequence mapping tasks. Subjects read visually presented sentences and nonlinguistic sequences presented one word/element at a time, and performed grammaticality/correctness judgments after each, responding by button press. Red areas indicate regions activated by both tasks, including a prefrontal network that involves BAs 44, 46, 9, and 6. Green areas indicate cortical regions activated exclusively in the sentence processing task, including BAs 21, 22, 47, and 45. From Hoen et al. (2006) with permission.

involved in the computational application of rules including our proposed mappings.

In contrast, the integration of lexicosemantic content into this structure processing machinery, filling of the PredictedReferentsArray, corresponds to a more language-specific ventral stream mechanism that culminates in the pars triangularis (BA 45) of the ventral premotor area, consistent with the declarative component of Ullman's (2004) model. In this context, Hoen et al. (2006) observed that when the processing of sentences and nonlinguistic sequences was compared, BA 45 was activated exclusively by the sentence processing.

Goals of the Current Study

We have previously demonstrated that the model in Figure 2 can learn a variety of constructions that can then be used in different language interaction contexts (Dominey & Boucher, 2005a, b; Dominey, 2000). However, from a functional perspective, a model of grammatical construction processing should additionally address two distinct challenges of (1) cross-linguistic validity and (2) compositionality, which we will address in the current study.

Although the examples used above have been presented in English, the proposed model of grammatical construction processing should accommodate different types of languages. Thus, the first objective will be to demonstrate that the proposed system is capable of

learning grammatical constructions in English, French, and Japanese. Whereas English and French are relatively similar in their linguistic structure, Japanese is somewhat different in that it allows more freedom in word order, with information about thematic roles encoded in case markers.

Regarding the challenge of compositionality, as presented above, it appears that for each different type of sentence, the system must have a distinct construction. Considering Figure 1C, this indicates that every time a noun is expanded into a relative phrase, a new construction will be required. This is undesirable for two reasons: First it imposes a large number of similar constructions to be stored, and second, it means that if a sentence occurs with a relative phrase in a new location, the model will fail to understand that sentence without first having a <sentence, meaning> pair from which it can learn the mapping. Alternatively, the model could process constructions within constructions in a compositional manner. Thus, the sentence "The dog that chased the cat bit Mary" could be decomposed into its constituent phrases "The dog chased the cat" and "The dog bit Mary." The goal is to determine whether the model as presented above can accommodate this kind of compositionality.

During the processing of multiple sentences with this type of embedded relative clause, it will repeatedly occur that the pattern "*agent that action object*" will occur, and will map onto the meaning component *action(agent, object)*. The same kind of pattern finding that allows the association of a ConstructionIndex with the corresponding FormToMeaning mapping could also work on such patterns that reoccur within sentences. That is, rather than looking for patterns at the level of the entire sentence or ConstructionIndex, the model could apply exactly the same mechanisms to identify subpatterns within the sentence. This would allow the system to generalize over these occurrences such that when this pattern occurs in a new location in a sentence, it can be extracted. This will decompose the sentence into the principal and the relative phrases, with the outcome that both will more likely correspond to existing constructions (see Figure 1D). Clearly, this will not work in all cases of new relative clauses, but more importantly, it will provide a compositional capability for the system to accommodate a subclass of the possible new sentence types. This type of segmentation approach has been explored by Miikkulainen (1996), discussed below.

METHODS

In the first set of experiments (Experiment 1A–C) for each of the three languages English, Japanese, and French, respectively, we will expose the model to a set of <sentence, meaning> pairs in a training phase. We then validate that the model has learned these sentences by presenting the sentences alone and comparing the

generated meanings with the actual meanings. Finally, we test the generalization capability by creating a new corpus of sentences based on the learned constructions that were not used in training. We compare the meaning generated for each sentence with its actual meaning.

In the second set of experiments (Experiment 2A–D) we attempt to determine whether pattern matching within a sentence can be used to extract relative phrases. Specifically, these experiments test the hypotheses that (1) in English, French, and Japanese, certain relative phrases are reliably marked by recurring patterns that can be used to extract these relative phrases in a reliable manner so that (2) they can be processed in the standard manner with the construction model. The net result is the capability to handle structurally novel sentences by decomposing them into their phrasal components in a construction grammar setting.

Model Overview

The model is trained on <sentence, meaning> pairs and then is tested with new sentences to generate the appropriate meanings. Input sentences are processed word by word. Open class words are directed to the open class array (OCA), and are represented by an empty slot in the ConstructionIndex. Closed class words are copied into the ConstructionIndex, which thus preserves word order. Note that this lexical categorization can be reliably performed by infants and neural network simulators, based on the acoustic properties of open versus closed class categories (Blanc et al., 2003; Shi et al., 1999).

If an event meaning is provided as input (for learning the sentence-meaning mapping), the event meaning is encoded in a predicate-argument format in the scene event array (SEA). Words in OCA are translated to Predicted Referents via the WordToReferent mapping to populate the predicted referents array (PRA). In the current study, this is an identity mapping. In the critical function of the model, PRA elements are mapped onto their thematic roles in the SEA by the FormToMeaning mapping specific to each sentence type. This mapping is retrieved from the ConstructionInventory, based on the ConstructionIndex that characterizes each sentence type.

The link between the ConstructionIndex and the corresponding FormToMeaning mapping is established as follows. As each new <sentence, meaning> pair is processed during learning, the FormToMeaning mapping for that sentence is generated, by matching referents (in PRA) to scene elements (in SEA). The resulting FormToMeaning mapping thus essentially encodes the correspondence between word order in the sentence and thematic roles in the meaning. Given the FormToMeaning mapping for the current sentence, we can now store it in the ConstructionInventory, associated with the corresponding ConstructionIndex for that sentence. After training, we evaluate performance

by providing new sentences and testing whether the model can accurately generate the corresponding meaning. Given a new input sentence, its ConstructionIndex is extracted and used to recover the associated FormToMeaning mapping from the learned ConstructionInventory. The FormToMeaning mapping is then used to generate the meaning of the sentence by mapping elements in the PRA onto the SEA. This processing is summarized in the pseudocode in Table 1.

Now let us consider how the system could accommodate complex sentences as in Figure 1C (and the relative and conjoined sentences in Table 4). In a simple but noncompositional manner, the ConstructionInventory, FormToMeaning mapping, and SceneEventArray can be instantiated twice each to represent the two components of the dual scene. This noncompositional solution allows different relativized constructions to be learned, but does not generalize to new constructions, and raises issues concerning upper limits on the number of instantiations of the different structures. Thus, we now present the generalized compositional solution.

Processing Compositional Structure of Relative Phrases

As described above, the model will account for sentences with relative phrases by treating the entire sentence as a distinct construction, mapping the open class elements onto two distinct instances of the SceneEventArray as indicated in Figure 1C. As stated, this fails to account for an ability to understand sentences that employ a relative phrase in a *new* location without first learning the mapping for that specific construction via a <sentence, meaning> pair. Here we reveal how such compositional processing of relative phrases can be accounted for based on the use of the same underlying capabilities for lexical categorization and sequence recognition presented above. In particular, when a given word—or pattern of words—that marks relative phrases is encountered, the surrounding word group defining the relative phrase (specified for each language in Table 3) is extracted and processed as an independent sentence, followed by the processing of the remainder of the initial sentence. The pseudocode in Table 2 describes the process for extracting relative phrases. Thus, for example, with the following novel sentence, “The boy was seen by [the cat that hit the dog],” the system recognizes “that” and extracts the relative phrase “the cat hit the dog” and principal phrase “the boy was seen by the cat,” both of which are then processed as simple constructions that have been previously learned. We hypothesize that exploiting this pattern-based extraction of relative phrases, the model will correctly function for English, French, and Japanese. Table 3 specifies for each of these three languages the specific relative markers and relative phrase patterns for extraction. The important point is that the required

Table 1. Functional Characterization of the Construction Grammar Processing Model

```

while (sentence not empty)
{
  get next word
  if (word is open class) {
    put it in the next position in the OpenClassArray,
    append a blank slot to the ConstructionIndex}
  else if (word is closed class)
    append the word to the ConstructionIndex
  }
if (ConstructionIndex matches element i in ConstructionInventory)
  { //Comment: sentence employs an already learned construction
  extract FormToMeaning mapping from ConstructionInventory(i)
  apply FormToMeaning to OpenClassArray to
  generate SceneEventArray
  }
else if (sentence contains relative marker)
  { //Comment: sentence has a relative phrase in a novel position
  extract the relative phrase(s) as specified in Table 2.
  }
else
  { //Comment: sentence employs a new construction so learn it
  Store ConstructionIndex in ConstructionInventory
  Calculate FormToMeaning mapping by matching elements in OpenClassArray
  with elements in SceneEventArray
  Store FormToMeaning in ConstructionInventory
  }
}

```

processes of sequence pattern matching and subsequence processing are straightforward extensions of the corticostriatal sequence processing model and do not require any radically different processing from that already employed.

It is important to note that although this model is specified here in functional programming terms (and written in the “C” programming language), the component functions have all been demonstrated in neurocomputational models (Dominey et al., 2003; Dominey & Ramus, 2000; Dominey et al., 1998). In particular, the ability to recognize sequences and sequence elements,

required for processing the ConstructionIndex and the relative phrase makers, has been demonstrated in a recurrent neural network (Dominey, 1997; Dominey et al., 1995). Likewise, we demonstrated a hybrid neural network’s ability to fill working memory “slots” with specific sequence elements and then modulate the retrieval of the slot contents to transform these variable sequences for abstract sequence processing (Dominey & Ramus, 2000; Dominey et al., 1998) and syntactic comprehension (Dominey et al., 2003; Dominey, 2002). The ability to use the ConstructionIndex to apply a given form-to-meaning mapping has also been demonstrated

Table 2. Functional Characterization of the Relative Phrase Processing Model

```

extract_relative(sentence)
while (sentence not empty)
{
  from sentence end (or beginning in Japanese) get next word
  if (word == relative-marker) //Comment: rel-marker identified in Table 3
  {
    embedded_phrase = word group surrounding relative-marker
    //Comment: see word group templates in Table 3
    sentence = rest of sentence with word group and relative marker
    removed
  }
  apply construction-model to (embedded_phrase)
}
apply construction-model to (sentence)

```

Note that the search for relative markers is reversed for head-last (Japanese) versus head-first languages, corresponding to a language specific parameter, as indicated in Table 3.

Table 3. Language-specific Parameters for the Construction Processing Model

<i>Language</i>	<i>English</i>	<i>French</i>	<i>Japanese</i>
Closed class elements	the, a, to, from, by, for, was, at, that, before, after, and	le, la, a, à, été, par, au, aux	ga, wo, ni, niyotte, jibun, jishin
Relative marker	that, which, who	qui, que	Nonfinal verb
Head direction	Head first	Head first	Head last
Relative phrase structure	<u>X</u> <u>X</u> that _ _ _ the cat that chased the rat <i>chased(cat, rat)</i> *the cat that the rat chased <i>chased(rat, cat)</i>	<u>X</u> <u>X</u> qui/que _ _ _ _ *le chien que le chat a chassé <i>chassé(chat, chien)</i> le chien qui a chassé le chat <i>chassé(chien, chat)</i>	_ _ _ verb <u>X</u> <u>X</u> *dog wo oshita cat ga *dog wo oshita cat wo *dog wo oshita cat ni-yotte *dog ni-yotte oshita cat ga *cat ga oshita dog wo *cat ga oshita dog ga <i>oshita(cat, dog)</i>

The “relative phrase structure” indicates the word group that will be extracted. The components marked “X” will also be left in the principal phrase. Relative phrases marked with * correspond to forms that are acceptable only as relative phrases.

in a neural network model (Dominey et al., 2003; Dominey, 2002). Thus, in the current study we employ these neurally plausible functions without reimplementing them in neural networks.

EXPERIMENT 1 RESULTS: CANONICAL GRAMMATICAL CONSTRUCTIONS

Experiment 1A: English

Given a training corpus of <sentence, meaning> pairs specified in Table 4, the model should learn the grammatical constructions that define the mappings for these pairs. More importantly, the model should demonstrate a generalization capability so that for new sentences generated from the learned constructions, it should be able to extract their meanings. Thus, after learning the construction associated with the <sentence, meaning> pair <“John hit the ball that broke the window; broke(ball, window), hit(John, ball)”>, when given the test sentence “Mary chased the dog that bit the boy,” the model should produce the meaning “bit(dog, boy), chased(Mary, dog).”

In this context, we exposed the model to the training corpus of 38 sentences (each corresponding to a different grammatical construction) outlined in Table 4. We then generated a new generalization corpus based on the same constructions, in which all of the open class elements from the training corpus were systematically replaced by new open class elements. For each of these test sentences, the model successfully identified the corresponding construction, applied the corresponding sentence-to-meaning mapping, and generated the correct meaning.

That the model can accommodate these 38 different grammatical constructions with no modifications indicates its capability to generalize. This translates to a (partial) validation of the hypothesis that across languages, thematic role assignment is encoded by a limited set of parameters including word order and grammatical marking, and that distinct grammatical constructions will have distinct and identifying ensembles of these parameters (Bates et al., 1982; MacWhinney, 1982). However, these results have been obtained with English, which is a relatively fixed word order language, and a more rigorous test of this hypothesis would involve testing with a free word order language such as Japanese.

Experiment 1B: Japanese

The current experiment tests the model with sentences in Japanese. Unlike English, Japanese allows extensive liberty in the ordering of words, with information about thematic roles and grammatical functions indirectly marked by postpositional function word particles including -ga, -ni, -wo, -yotte, and -to. In the “active” Japanese sentences, the postpositional function words -ga, -ni, and -wo explicitly mark agent, recipient, and object, whereas in the passive, these are marked, respectively, by -ni-yotte, -ni, and -ga. For both the active and passive forms, there are four different legal word order permutations that preserve and rely on this marking. Japanese thus provides an interesting test of the model’s ability to accommodate such freedom in word order. An example of this word order flexibility of Japanese with respect to English is illustrated in Table 5 with the English passive ditransitive forms that can be

Table 4. English Constructions

<i>Example Sentences and Meanings</i>	<i>Grammatical Constructions</i>
1. The block pushed the cylinder. Push(block, cylinder)	1. Agent verb object. (Active) Verb(<i>agent</i> , object)
2. The cylinder was pushed by the block. Push(block, cylinder)	2. Object was verbed by agent. (Passive) Verb(<i>agent</i> , <i>object</i>).
3. The block gave the cylinder to the moon. Give(block, cylinder, moon)	3. Agent verbed object to recipient. (Dative) Verb(<i>agent</i> , object, recipient)
4. The cylinder was given to the moon by the block. Give(block, cylinder, moon)	4. Object was verbed to recipient by agent. (Dat passive) Action1(<i>agent1</i> , object2, recipient3).
<i>Dual-event Relative Constructions</i>	
6. The block that pushed the cylinder touched the moon. Push(block, cylinder), Touch(block, moon)	6. Agent1 that verb1ed object2 verb2ed object3. Action1(<i>agent1</i> , object2), Action2 (<i>agent1</i> , object3)
18. The cylinder that was pushed by the block gave the cat to the dog. Push(block, cylinder), give(cylinder, cat, dog).	18. Obj4 was act2ed from ag3 to recip1 that act1ed obj2 Action1(<i>agent1</i> , object2), Action2 (<i>agent3</i> , object4, recipient1)
<i>Dual-event Conjoined Constructions</i>	
27. The block pushed the cylinder and the moon. Push(block, cylinder), Push(block, moon)	27. Agent1 action1 object1 and object. Action1(<i>agent1</i> , object1), Action1(<i>agent1</i> , object2)
30. The moon and the block were given to the cylinder by the cat. Give(cat, moon, cylinder), Give(cat, block, cylinder).	30. Object2 and object3 were action1ed to recipient4 by agent1. Action1(<i>agent1</i> , object2, recipient4), Action1(<i>agent1</i> , object3, recipient4)

Each numbered sentence is an example of a specific abstract grammatical construction type whose meaning is provided in an event(argument) format following the sentence corresponding to that meaning. Each construction can generalize to new sentences in which the open class elements are replaced.

expressed in four different common manners in Japanese (Constructions 9–12).

Using the same method as described in the previous experiment, we thus expose the model to <sentence, meaning> pairs generated from 26 Japanese constructions, a sample of which is described in Table 5. We predicted that by processing the -ga, -ni, -yotte, and -wo particles as closed class elements, the model would be able to discriminate and identify the distinct grammatical constructions and learn the corresponding mappings. Indeed, the model successfully discriminates between all of the 26 construction types based on the Construction-Index unique to each construction type, and associates the correct FormToMeaning mapping with each of them. As for the English constructions, once learned, a given construction could generalize to new untrained sentences formed by replacing the open class elements in the training sentences with a different distribution of open class elements (see Dominey & Inui, 2004).

This demonstration with Japanese is an important validation that at least for this subset of constructions, the construction-based model is applicable both to fixed word order languages such as English, as well as free word order languages such as Japanese. This also provides further validation for the proposal of Bates et al. (1982) and MacWhinney (1982) that thematic roles are indicated by a constellation of cues including grammatical markers and word order.

Experiment 1C: French

As for the English and Japanese studies, the model was trained on a set of French sentences (illustrated in Table 6) that were paired with their meanings. The model was then tested using the same sentences with no input meanings to verify that it could retrieve the corresponding meanings. Finally, it was exposed to a

Table 5. Japanese Constructions^a

<i>Single-event Constructions</i>	<i>Grammatical Construction</i>
<i>The block hit the circle (1)</i>	
1. block-ga circle-wo tataita.	1. Agent-ga object-wo verb.
2. circle-wo block-ga tataita. tataita(block, circle)	2. Object-wo agent-ga verb. Verb(agent, object).
<i>The circle was hit by the block (2)</i>	
3. Circle-ga block-ni tatakareta.	3. Object-ga agent-ni verb.
4. Block-ni circle-ga tatakareta. tatakareta (block, circle)	4. Agent-ni object-ga verb. Verb(agent, object).
<i>The circle was given to the triangle by the block (4)</i>	
9. Circle-ga block-ni-yotte triangle-ni watasareta.	9. Object-ga agent-ni-yotte recipient-ni verb.
10. Block-ni-yotte circle-ga triangle-ni watasareta.	10. agent-ni-yotte object-ga recipient-ni verb.
11. Block-ni-yotte triangle-ni circle-ga watasareta.	11. agent-ni-yotte recipient-ni object-ga verb.
12. Triangle-ni circle-ga block-ni-yotte watasareta. watasareta (block, circle, triangle)	12. recipient-ni object-ga agent-ni-yotte verb. Verb(agent, object, recipient)
<i>Dual-event Relative Constructions</i>	
<i>The block that hit the circle pushed the triangle (6)</i>	
13. Circle-wo tataita block-ga triangle-wo oshita. Hit(block, circle), Pushed(block, triangle)	13. Obj1-wo verb1 Agent1-ga Obj2-wo verb2. Verb1(Agent1, Obj1), verb2(Agent1, Obj2)
<i>The block was hit by the triangle that pushed the circle (7)</i>	
14. Block-ga circle-wo oshita triangle-ni-yotte tatakareta.	14. obj2-ga obj1-wo verb1 agent1-ni-yotte verb2.
15. Circle-wo oshita triangle-ni-yotte block-ga tatakareta. Pushed(triangle, circle), Hit(triangle, block)	15. obj1-wo verb1 agent1-ni-yotte obj2-ga verb2. Verb1(agent1, circle), Verb2(agent1, block)
<i>The block that hit the circle was pushed by the triangle (8)</i>	
16. Circle-wo tataita block-ga triangle-ni-yotte osareta.	16. obj1-wo verb1 agent1-ga agent2-ni-yotte verb2.
17. Triangle-ni-yotte circle-wo tataita block-ga osareta. Hit(block, circle), Pushed(triangle, block)	17. agent2-ni-yotte obj1-wo verb1 agent1-ga verb2. Verb1(agent1, obj1), verb2(agent2, agent1)
<i>The block hit the triangle that pushed the circle (9)</i>	
18. Block-ga circle-wo oshita triangle-wo tataita.	18. Agent2-ga obj1-wo verb1 agent1-wo verb2.
19. Circle-wo oshita triangle-wo block-ga tataita. Pushed(triangle, circle), Hit(block, triangle)	19. obj1-wo verb1 agent1-wo agent2-ga verb2. Verb1(agent1, obj1), Verb2(agent2, agent1)
<i>The block that was hit by the circle pushed the triangle (10)</i>	
20. Circle-ni-yotte tatakareta block-ga triangle-wo oshita. Hit(circle, block), Push(block, triangle)	20. agent1-ni-yotte verb1 obj1-ga obj2-wo verb2. Verb1(agent1, obj1), verb2(obj1, obj2)
<i>Dual-event Conjoined Constructions</i>	
<i>The block hit the circle and the triangle (27)</i>	
21(27). Block-ga circle-to triangle-wo tataita.	21(27). agent-ga obj1-to obj2-wo verb.
22(27). Circle-to triangle-wo block-ga tataita. Hit(block, circle), hit(block, triangle)	22(27). Obj1-to obj2-wo agent-ga verb. Verb(agent, obj1), verb(agent, obj2)

Numbered examples of specific abstract grammatical construction type in Japanese whose meaning is provided in an event(argument) format following the sentence(s) corresponding to that meaning. The corresponding English construction is indicated in parentheses.

^aHit = tatakau, hit = tataita, be hit = tatakareru, was hit = tatakareta; give = ataeru, gave = watahita, be given = ataerareru, was given = watasareta; push = osu, pushed = oshita, be pushed = osareru, was pushed = osareta; believe = shinjiru, believed = shinjita; itself = jibun or jishin, it = sore.

Table 6. French Constructions

<i>Example Sentences and Meanings</i>	<i>Grammatical Constructions</i>
<i>The boy pushed the dog.</i>	
1. Le garçon a poussé le chien. poussé(garçon, chien)	1. Agent a verb object. (Active) Verb(agent, object)
<i>The dog was pushed by the boy.</i>	
2. Le chien a été poussé par le garçon. poussé(garçon, chien)	2. Object a été verbed par agent. (Passive) Verb(agent, object).
<i>The boy gave the block to the dog.</i>	
3. Le garçon a donné le bloc au chien. donné(garçon, bloc, chien)	3. Agent a verb le object au recipient. (Dative) Verb(agent, object, recipient)
<i>The block was given to the dog by the boy.</i>	
4. Le bloc a été donné au chien par le garçon. donné(garçon, bloc, chien)	4. Object a été verbed au recipient par agent. (Dat passive) Verb(agent, object, recipient)
<i>Dual-event Relative Constructions</i>	
<i>The boy who pushed the dog touched the block.</i>	
5. Le garçon qui a poussé le chien a touché le bloc. poussé(garçon, chien), touché(garçon, bloc)	5. Agent1 qui a verb1ed object2 a verb2ed object3. Action1(agent1, object2), Action2(agent1, object3)
<i>The boy who was pushed by the chief gave the block to the dog.</i>	
6. Le garçon qui a été poussé par le chef a donné le bloc au chien. poussé(chef, garçon), donné(garçon, bloc, chien)	6. Obj1 qui a été verb1 par agent1 a verb2 obj2 au recip1 Verb1(agent1, object1), verb2(object1, object2, recip1)
<i>Dual-event Conjoined Constructions</i>	
<i>The chief pushed the boy and the block.</i>	
7. Le chef a poussé le garçon et le bloc. poussé(chef, garçon), poussé(chef, bloc)	7. Agent1 a verb object1 et object. Verb1(agent1, object1), verb1(agent1, object2)
<i>The block and the table were given to the chef by the master.</i>	
8. Le bloc et la table ont été donnés au chef par le maître. donnés(maître, bloc, chef), donnés (maître, table, chef).	8. Object2 et object3 ont été verb1 to recipient4 par agent1. Verb1(agent1, object2, recipient4), verb1(agent1, object3, recipient4)

new set of sentences made up from the same constructions, but with a new set of open class elements that were not used in training. For these validation sentences, the model also reliably performed the mapping from sentence to meaning for all sentences.

EXPERIMENT 2 RESULTS: COMPOSITIONAL CONSTRUCTIONS FOR RELATIVE PHRASES

In the previous experiments, for each of the three languages, a limited set of sentences including those

with relative phrases was learned, based on a noncompositional mechanism. For these sentences, the verbs in the main and relative phrases in the OpenClassArray became associated with the event components of two distinct copies of the SceneEventArray, corresponding to constructions like that in Figure 1C. In this manner, each different relative construction was learned explicitly, with no capability to generalize to sentences with relative phrases in novel positions. The current experiments address this limitation, testing the hypothesis that pattern-based extraction of relative phrases will correctly function for English, French, and Japanese.

Experiment 2A: English

For the English experiments, the model was first trained on five sentence types corresponding to the active and passive transitive and dative constructions (1–4 in Table 4), and the relative “The boy the dog hit,” along with their corresponding meaning representations. We then tested performance on the following sentences. In each, the principal phrase is underlined, and the relative phrase(s) is/are indicated in square brackets:

1. The boy was seen by [the cat that hit the dog].
2. The boy saw [the cat that hit [the dog] that bit the man].
3. [The boy that the cat bit] told the man.

When Sentence 1 was presented to the trained model, the system did not recognize the construction type and thus attempted to determine if there was an embedded phrase structure revealed by the presence of the embedding marker “that.” Applying phrase structure extraction thus yielded the embedded phrase, “The cat hit the dog,” and the remaining principal phrase, “The boy was seen by the cat.” These two phrases were both recognized as corresponding to learned constructions by the system, and the corresponding meanings were returned: Hit(cat, dog), and seen(cat, boy). Similarly for Sentence 2, application of the phrase structure extraction yielded three phrases: “The dog bit the man,” “The cat hit the dog,” and “The boy saw the cat,” for each of which the corresponding meanings were extracted based on the previously learned constructions. Likewise, Sentence 3 was recognized as corresponding to no known construction, and the phrase structure extraction yielded the relative phrase “the boy the cat bit” and the principal phrase “the boy told the man,” both of which correspond to learned constructions. These results thus demonstrate that by extracting relative phrases that can then be processed as learned constructions, the system displays the compositional capability to correctly process new grammatical constructions that it has not previously encountered in training. Indeed, although we provide the detailed explanation for these three sentences, the model can generalize to an open set of sentences in which any noun can be expanded into a relative phrase such as “The boy that hit the cat that the girl saw gave the ball to the man that threw the ball.”

Experiment 2B: Japanese

For the Japanese constructions, we take the same approach, first training the system on a small set of simple phrases that can then be recombined to construct a variety of more complex sentences with relative phrases. The training sentences were Examples 1 and 2 from Table 5, and the two passives, “Dog ni-yotte cat ga verb” and “Cat ga dog ni-yotte verb,” both with the meaning

“verb(dog, cat).” We also trained the system with the relative phrase structures for Japanese indicated in Table 3.

Japanese is a verb-final language, and relative phrases are marked by the presence of a verb in a nonfinal position as with the verb *tataita* in Sentence 13 of Table 5: “[Circle-wo *tataita* block-ga triangle-wo *oshita*,” with the corresponding meaning: *tataita*(block, circle), *oshita*(block, triangle) or “The block that hit the circle pushed the triangle.” Thus, rather than basing detection of relative phrases on an explicit marker, in Japanese the relative phrase extractor simply attempts to locate verbs in nonfinal positions and then extracts the surrounding words, as indicated in Table 3, that constitute the embedded relative phrase. We tested the model on seven new sentences below with relative phrases in which the nonfinal verb is indicated as “relverb.” Each was first found by the model to have a Construction-Index that was unknown and was decomposed by the model into a combination of the learned single-verb constructions. To illustrate the results, we present each of the seven sentences followed by the two phrases extracted by the model and the corresponding meanings (in italics) that were generated.

1. Boy ga cat wo relverb dog ni-yotte tatakareta.
Cat wo relverb dog ni-yotte. *Relverb(dog, cat)*
Boy ga dog ni-yotte tatakareta. *hit(dog, boy)*
2. Cat wo relverb dog ni-yotte boy ga tatakareta
Cat wo relverb dog ni-yotte. *Relverb(dog, cat)*
Dog ni-yotte boy ga tatakareta. *hit(dog, boy)*
3. Cat wo relverb boy ga dog ni-yotte osareta
Cat wo relverb boy ga. *Relverb(boy, cat)*
Boy ga dog ni-yotte osareta. *pushed(dog, boy)*
4. Dog ni-yotte cat wo relverb boy ga osareta.
Cat wo relverb boy ga. *Relverb(boy, cat)*
Dog ni-yotte boy ga osareta. *pushed(dog, boy)*
5. Boy ga cat wo relverb dog wo tataita
Cat wo relverb dog wo. *Relverb(dog, cat)*
Boy ga dog wo tataita. *hit(boy, dog)*
6. Cat wo relverb dog wo boy ga tataita
Cat wo relverb dog wo. *Relverb(dog, cat)*
Dog wo boy ga tataita. *hit(boy, dog)*
7. Cat ni-yotte relverb boy ga dog wo oshita
Cat ni-yotte relverb boy ga. *Relverb(boy, cat)*
Boy ga dog wo oshita. *pushed(boy, dog)*

Experiment 2C: French

Like English, French allows for explicit marking of relative phrases with the two closed class words “qui” and “que,” and so the embedded structure extraction was based on recognition of one of these words as indicated in Table 3. We trained the model on the single-verb construction sentences of Types 1–4 from Table 6, and one relativized structure:

<“le chat le chien a chassé”; chassé (chien, chat)>.
<“the cat the dog chased”: *chased(dog, cat)*>

We then tested the system on the following relative sentences:

1. Le chat qui a chassé le rat a frappé le poisson.
The cat that chased the rat bit the fish.
2. Le chat que le chien a chassé a mangé le bloc.
The cat that the dog chased ate the block.
3. Le fromage que le garçon a lancé a été donné au rat par le chat.
The cheese that the boy threw was given to the rat by the cat.

Again, as predicted, these relative constructions were not recognized in the ConstructionInventory, and the embedded structure extractor was thus employed. Each of these sentences was successfully decomposed into a combination of the four basic sentence types and the two embedded phrase types illustrated in Table 3, and the corresponding meaning was produced for each.

Experiment 2D: Extension to the Miikkulainen (1996) Corpus of Relativized Sentences

Experiments 2A–C illustrate the compositional capability of the model to extract embedded relative phrase structure, but do not indicate how well it would scale to larger corpora and more complex sentences. In this context, Miikkulainen (1996) developed a corpus of 49 sentences with up to four clauses per sentence arranged in a variety of grammatical structures against which he tested his model. We thus used this set of 49 sentences in English and translated the sentences to Japanese. Interestingly, because of the left-branching, head-last structure of Japanese, the unpacking of the relative phrases must proceed from leftmost to rightmost. This is in contrast to English and French, whose right-branching, head-first structure requires the unpacking to proceed from right to left. Indeed, as French is essentially isomorphic to English in this context of relative phrase structure, we limit the current exercise to English and Japanese sentences.

For both English and Japanese, the model instances were first pretrained on a small set of sentences as specified in Experiments 2A and 2B. The models were then tested on the 49 English and Japanese sentences, respectively. For both languages, the models successfully extracted the embedded phrases that could then be processed as previously learned constructions. Examples of two sentences in English and Japanese and their corresponding meanings that were extracted by the model are presented here:

7. The girl saw the boy who chased the cat who the dog who bit the girl bit.
Girl-ga girl-wo bit dog-ga bit cat-wo chased boy-wo saw.

Bit(dog, girl), bit(dog, cat), chased(boy, cat), saw(girl, boy)

46. The girl who the dog who the boy pushed bit saw the boy who chased the cat.
Boy-ga pushed dog-ga bit girl-ga cat-wo chased boy-wo saw.
Pushed(boy, dog), bit(dog, girl), chased(boy, cat), saw(girl, boy)

GENERAL DISCUSSION

The results of these experiments demonstrate that the principals inherent in the grammatical construction model are adequate for processing sentences in English, Japanese, and French, including the ability to handle novel grammatical constructions with embedded relative phrases.

Psychological Reality of the Model

The first principle inherent in the model is that instead of representing <sentence, meaning> mappings in terms of a generative grammar, they are represented directly in a structured inventory of grammatical constructions that are nothing more than these mappings (Tomasello, 2003; Goldberg, 1995). Growing evidence both from studies of human language development (Tomasello, 2000, 2003), and adult processing (Ferreira, 2003; Sanford & Sturt, 2002) indicate that a substantial component of language behavior can be accounted for in this manner. That is, a large part of language production and comprehension is based on the reuse (including recombination) of existing templates, in a context in which the templates (i.e., grammatical constructions) can be learned by straightforward mechanisms as illustrated in Figure 2.

Although this use of learned grammatical constructions can account for a certain level of language performance, it is clear that human language also includes the creative ability to generate and understand wholly novel sentence forms. One of the advantages of the generative grammar framework is precisely that this capability is accounted for, with the associated difficulty of specifying how a generative grammar capability could have evolved. On the other hand, the path from primate sensorimotor sequence learning capabilities to the grammatical construction capabilities presented here is much more direct (see Dominey, 2005a, 2005b). With respect to nativist issues, the current model is consistent with the usage-based construction grammar account in which language acquisition is largely based on learning rather than an essentially complete system of syntactic representation from the beginning (reviewed in Tomasello, 2000). The essential innate capabilities we postulate are mechanisms for (1) abstract (variable based) structure mapping and for

(2) pattern-based embedded structure processing. Interestingly, these mapping and embedding capabilities could have their origins in sensorimotor processing, required for transformations (mapping) between multiple sensory and motor domains and for embedding motor programs within larger sensorimotor plans, respectively (see also Dominey, 2005b). The additional requirement for language may be the ability to insert symbolic meanings into this transformation capability as performed by BA 45 (*pars triangularis*) in our model.

What has been lacking, however, is a specification of how grammatical constructions can be combined in order to generate new constructions on the fly, using these same sequence processing mechanisms. The current research takes a first concrete step in this direction. In three distinct languages, we have demonstrated how relative phrase structure (for limited categories of relative phrases) can be recognized and extracted, allowing novel constructions to be understood in terms of their simpler component constructions. This approach should extend to accommodate more generalized relative phrase structure, as well as prepositional and other categories of embedded phrases.

This ability to recombine constructions provides the desired property of allowing infinite sentences to be represented by a finite set of compositional constructions. At the same time, the system preserves the defining property of construction grammar that distinct constructions will be invoked when the required meaning cannot be encoded by a combination of existing constructions (see Goldberg, 1995). Thus, when variations in word order are used in order to place the pragmatic focus on a particular item, as in the Japanese examples 9–12 in Table 5, distinct constructions will be employed to capture these differences. Likewise, the English active and passive forms correspond to distinct constructions as they have different pragmatic focus, indicated by the underlined arguments in Examples 1 and 2 of Table 4. However, the experiments with the Miikkulainen (1996) corpus of 49 relative phrase sentences illustrated how a small set of constructions can be used in a compositional manner to generate a large set of relative sentences. The construction grammar bet here is that the strategy of (a) invoking new constructions when meaning cannot be expressed with existing constructions and (b) exploiting compositionality of existing constructions will provide a manageable construction inventory with complete expression and comprehension capability.

Division of Labor in the Construction Inventory

Construction grammar assumes no strict division between the lexicon and syntax in that both lexical and syntactic constructions pair form with meaning (see

Goldberg, 1995). However, construction grammar does not deny the existence of distinctly syntactic constructions. Although words and abstract argument constructions thus share a common theoretical existence as constructions, from a neural implementation there are necessary differences. In particular, phrasal arguments (such as *agent* in a transitive construction) are like abstract variables that can be instantiated by an open set of nouns. In the context of a debate over the neural processing of such abstract variables (Marcus et al., 1999, and subsequent commentaries), we have demonstrated that a neural network that is capable of learning fixed symbol sequences (i.e., without variables) fails in learning abstract sequences in which variable slots can be instantiated by different symbols because it lacks a variable representation capability (Dominey & Ramus, 2000). As suggested by Marcus et al. (1999), we demonstrated that the manipulation of these variable sequences requires additional machinery, including a form of working memory where the variable contents are temporarily instantiated (Dominey & Ramus, 2000).

The same issue holds for constructions. In an early developmental phase, an utterance like “Gimme the ball” could be understood as an undifferentiated “idiomlike holophrase” construction (see Tomasello, 2003). No lexical categorization of open versus closed class words would occur, and the entire utterance would contribute to the definition of the ConstructionIndex. Subsequent developmental categorization of an open class category of concrete nouns would allow verb-based constructions such as “Gimme X.” Although the mechanisms allowing this category formation remain to be investigated, high variability in the position following “Gimme” in the input sentences could allow statistical mechanisms to determine that elements in this position are members of a form of equivalence class. A pertinent account of such distributional analysis is presented in Cartwright and Brent (1997).

Once this categorization occurs, “Gimme” would still contribute to the ConstructionIndex, whereas the object position would now be an open class argument. In other words, if each utterance corresponds to a distinct construction, then a “single mechanism” can accommodate the form-to-meaning mappings. However, if argument variables are to be introduced that can be instantiated by a set of open class elements, then an additional mechanism for handling these variables and their contents is required. Indeed, as we illustrated with the relative phrase processing, these open class arguments can even be instantiated by constructions as well as open class words, thus contributing to compositionality. In this sense, the current model represents what we consider a necessary refinement that is wholly consistent with construction grammar, and indeed provides insight into a consistent explanation for compositionality.

Temporal Dynamics and Resolution of Detail

The theoretical foundation of the model is that the sentence-to-meaning mapping is performed based on the configuration of closed class elements (the ConstructionIndex) in the sentence (consistent with Bates et al., 1982; MacWhinney, 1982). In the current implementation, to reduce processing complexity, we wait until the entire sentence has been seen before extracting the corresponding mapping. However, as the ConstructionIndex is built up incrementally, there is nothing that prevents an implementation in which the mapping process begins to take place from the outset as the ConstructionIndex is successively built up. In such an implementation, the search for a particular construction would begin at the onset of sentence processing, continuously using the evolving ConstructionIndex to retrieve the current best guess at the construction, thus leading to garden path interpretations in appropriate sentences.

In this sense, the model is consistent with incremental parsing. In the current study, however, we do not explore real-time effects observed with such an implementation, restricting the analysis to the sentence level of granularity. The question then arises as to whether there is adequate information available to permit incremental processing. For example, in Japanese, the relative clauses are marked by a nonfinal verb. This implies that the parser must wait until it sees a verb and then beyond in order to determine the presence of a relative phrase, and so the results of the parse may change leading to garden path effects (see Yamashita, 1997). However, Miyamoto (2002) observed that in some cases, marking on the noun phrases indicates phrase boundaries prior to the arrival of the verb, so that the detection of the relative phrase structure can actually precede the nonfinal verb (Miyamoto, 2002). Thus, for example, a clause in Japanese cannot contain more than one accusative *o*-marked NP; hence, when a second one is encountered, this signals a phrase boundary. In general, if delay is decreased, then the required information must come from an alternative source (such as the double *o* constraint). Thus, in building up the ConstructionIndex incrementally, when such information is not available then the incremental processing will induce temporary failures including well-documented garden path effects.

Underlying Neurophysiology

Part of the importance of the current approach is that the underlying primitive functions that provide this construction processing capability can be linked in a plausible manner to known neurophysiological systems. In this context, the sequential structure of the ConstructionIndex can be implemented in recurrent corticocortical networks (Dominey et al., 2003). The

resulting pattern of cortical activity projects to the caudate nucleus of the striatum, and via these modifiable corticostriatal synapses, retrieves the appropriate FormToMeaning mapping that is then implemented in the frontal transformation processing network that includes BA 44. We have simulated how a recurrent cortical network can encode the sequential structure of closed and open class words in a sentence, corresponding to the ConstructionIndex, in order to retrieve the appropriate FormToMeaning mapping based on corticostriatal associative memory (Dominey et al., 2003; Dominey, 2002, 2005a, 2005b). Again, this is consistent with recent studies indicating an involvement in the corticostriatal system in the selection and application of rules in multiple domains (Longworth et al., 2005; Teichmann et al., 2005; Ullman, 2004, 2006).

Once the ConstructionIndex has been used to retrieve the appropriate FormToMeaning mapping, its implementation relies on a non-language-specific transformation processing mechanism that corresponds to a local cortical network including BAs 44, 46, 9, and 6 (Hoen et al., 2006). Primate neuroanatomy and human brain imagery indicate that at least part of this network, in particular BA 46, corresponds to the frontal component of the dorsal visual stream (Ungerleider, Courtney, & Haxby, 1998), consistent with its proposed role here in structural transformation processing.

This result is consistent with previous fMRI studies that showed activation of these areas in sentence comprehension tasks (e.g., Kuperberg et al., 2003; Baumgaertner, Weiller, & Büchel, 2002), particularly for sequential or structural processing aspects in sentences (Newman, Just, Keller, Roth, & Carpenter, 2003; Dapretto & Bookheimer, 1999; Kang, Constable, Gore, & Avrutin, 1999; Inui et al., 1998). This region was also often found activated in multimodal nonlinguistic tasks that specifically implicate temporal sequencing aspects (Drummond, Brown, & Salamat, 2003; Marshuetz, Smith, Jonides, DeGutis, & Chenevert, 2000). Above all, this is consistent with observations that event meaning is represented in this area 44 pars opercularis region (corresponding to our SceneEventArray), when events are being visually observed (Buccino et al., 2004) and when their descriptions are listened to (Tettamanti et al., 2005).

In contrast, for sentence comprehension the integration of lexicosemantic content into PredictedReferentsArray for subsequent FormToMeaning processing corresponds to a ventral stream mechanism, including BAs 22, 21, 47, and 45 (Hoen et al., 2006) that culminates in the pars triangularis (BA 45) in the inferior frontal gyrus (IFG) region (Ungerleider et al., 1998), consistent with the declarative component of Ullman's (2004) model. Interestingly, although this area (BA 45) was specifically activated in the language task in our experiment (Hoen et al., 2006), it is more generally characterized as participating in object or semantic (vs. spatial) working memory functions (reviewed in

Ungerleider et al., 1998), consistent with its proposed role here for semantic integration. This observation is in agreement with various functional imaging experiments that have shown activation of the anterior part of the ventral prefrontal cortex related to semantic or thematic processes in sentence comprehension (Sakai, 2005; Newman et al., 2003; Hashimoto & Sakai, 2002; Cooke et al., 2001; Dapretto & Bookheimer, 1999).

We have previously explored agrammatic aphasic sentence processing in an earlier (but equivalent) version of the model (Dominey, 2002). The lesion of the BA 44 system was simulated by introducing noise to the ConstructionIndex and by introducing a bias toward the statistically more common canonical (e.g., active vs. passive) sentence forms. This yielded a profile of sentence comprehension performance highly correlated ($R^2 = .89$) with that of the nine agrammatic patients we studied (Dominey, 2002). It is worth noting that implementing the ConstructionInventory as an associative mechanism rather than an indexed table would likely yield this bias effect based on the statistics of the training corpus.

Related Neurocomputational Parsing Models

The current proposal is consistent with and complementary to several existing neurocomputational models of language processing. Friederici (2002) analyses auditory sentence processing and concludes that BA 45/47 is more involved in semantic processing for thematic role assignment, whereas BA 44 is involved, not in the processing of syntax per se, but rather in aspects of syntactic memory and structure building in linguistic and nonlinguistic domains. Grodzinsky (2000) argues that Broca's area is involved in a more specialized form of syntactic transformation processing. These views correspond rather well with our characterization of BA 44 for processing structural transformations of linguistic and nonlinguistic sequences, and BA 45 for inserting semantic content into this mapping process for thematic role assignment. We further suggest that this transformation processing is guided by templates that are extracted from a corticostriatal memory system, based on the configuration of closed class elements coded in a recurrent corticocortical network (corresponding to the ConstructionIndex). This is consistent with, and extends, Ullman's (2004, 2006) characterization of a corticostriatal procedural memory for grammatical structure processing at the word level. Although these comparisons are made within the domain of language comprehension, it is noteworthy that syntactic processing in language production activates cortical regions adjacent to and overlapping BA 44 (Indefrey et al., 2001). The common point consistent with all of these models is that BA 45 appears to be associated with semantic processing, whereas BA 44 is associated with grammatical structure processing. This position has

recently been reaffirmed by Sakai (2005), who proposed that the left IFG region extending from the triangular part (F3t or BA 45) to the orbital part (F3O or BA 47) is the putative region for the selection and integration of semantic information, whereas the opercular and triangular parts (F3op/F3t or BAs 44 and 45) of the left IFG and the left lateral premotor cortex (BAs 6, 8, and 9; mainly in BA 8) are selectively related to grammatical processing.

Extraction of Relative Phrase Structure

Part of the novelty of our approach is that the extraction of relative phrase structure relies on these same capabilities. Identification within a sequence of an element (or group of elements) that indicates a relative phrase (such as "that," a nonfinal verb, etc.) is a basic recognition function that is clearly neurophysiologically feasible. Based on this recognition, the embedded sequence is extracted and is then processed by the construction model. Thus, relative phrase processing wholly reuses the existing construction processing capability in order to apply this capability to the extracted phrasal constituents. Our cross-linguistic study reveals that the mechanism we propose works equally well for head-initial and head-final languages (i.e., English and French, vs. Japanese, respectively) always proceeding from deepest to shallowest embedding. Interestingly, Christophe, Nespore, Gausti, and van Ooyen (2003) have shown how this direction parameter may be established based on prosodic cues that are perceptually salient to the infant by 2–3 months of age.

In this context of prosodic cues for syntactic structure (e.g., Blanc et al., 2003; Shi et al., 1999), the gross lexical categorization of all open class elements (nouns, verbs, etc.) into a single class is a source of information loss, but allows a first-level approximation for learning simple constructions. The model was confronted with this issue in the processing of more complex, Japanese relative sentences. In the current study, unlike the English and French sentences that used explicit grammatical function words for marking relative phrases, relative phrases in the Japanese sentences were marked by the presence of a verb in a nonfinal position in the sentence. We thus introduced a finer level of detail in the lexical categorization so that verbs are recognized as a separate category. Again, this is an extension that reuses the existing capabilities of the model and requires no structural modifications; but places stricter requirements on the lexical categorization capability. A similar case can also arise in English in the absence of an explicit marker, in which a relative phrase can be indicated by the presence of two successive noun phrases. This indicates that lexical categorization at the level of open versus closed class elements is not sufficient, and that a further degree of granularity is required for the general processing of relative phrases. In other words, if we wish

to fractionate constructions, we need more visibility into their internal structure, as revealed by a finer grain of lexical categorization.

Miikkulainen (1996) developed a related system called SPEC that was especially designed to address the problem of generalizing relative clause structures with new sentences. The central component is the simple recurrent network (SRN) that reads distributed word representations as its input and generates case-role representations as its output. The subsymbolic parser for embedded clauses (SPEC)'s generalization capability is based on simplifying the SRN's task through several architectural innovations, particularly: (a) introducing a segmenter network that breaks the input sequence into smaller chunks; and (b) introducing a stack network that memorizes constituents over intervening embedded clauses. The SRN thus plays a role similar to the basic grammatical construction processing mechanism in the current model, and the segmenter corresponds to the phrase extraction mechanism. An important difference between the two is the use of a stack network. In SPEC, the current phrase is "pushed" onto a stack, whereas the embedded phrase is being processed, and is then "popped" off for subsequent processing. Nested phrase structures thus involve successive pushes as the parser recurses into this structure, and then successive pops as the nested phrases are successively completed.

The model that we present here reformulates this issue. The sentence is represented as a linear sequence, and nested phrases are extracted from this linear sequence in a "deepest-first" order by always choosing the deepest relative phrase marker remaining in the sentence, based on the head direction. This guarantees that there is no further nesting in the extracted phrase. Each extracted phrase is processed by the construction model, and then the remainder of the original sequence is processed. Rather than a stack with potentially deep nesting, the current approach relies on a working memory to hold the main phrase while the embedded phrase (which is guaranteed to have no further embedded structure) is being processed, and both stack (Miikkulainen, 1996; Voegtlin & Dominey, 2005) and our working memory have feasible neural implementations. We thus resolve this parsing problem without the use of a distinct stack system, and indeed can process the 49 sentence Miikkulainen corpus in English and a Japanese translation. Interestingly, the 49 English sentences were all decomposed into two simple constructions, whereas the Japanese required five. Thus, despite the free word ordering in Japanese, the variability in the 49 Miikkulainen sentences can still be captured with a small number of constructions.

The current approach to sentence processing is unusual in that it can decline the availability of well-documented computational techniques for parsing such as unification-based approaches (e.g., Vosse & Kempen, 2000), and instead attempts to determine what can be

done with very simple structure-mapping capabilities. We can begin to see, however, that in the extraction of embedded relative phrases, the properties of individual lexical items (e.g., "that" in English, "qui" and "que" in French) are exploited. At the extreme, this is related to lexically driven approaches including that of Vosse and Kempen (2000), lexical functional grammar (LFG; Bresnan, 2001), and head-driven phrase structure grammar (HPSG; Pollard & Sag, 1994), in which lexical entries include type descriptions that form a hierarchy that, combined with grammatical rules, allows for unification-based parsing. Such systems can display a monotonicity property such that as an input string is incrementally processed, the amount of information that contributes to the parse increases (Bresnan, 2001). This allows for the assignment of structures to incomplete input representations. Similarly, in our system, as the Construction-Index is incrementally built up, it can be used to retrieve the corresponding construction with increasing fidelity.

A closely related approach is that of embodied construction grammar (Bergen & Chang, in press) in which construction-based utterance understanding involves *analysis* to determine which constructions the utterance instantiates, and mental *simulation* using the body's perceptual and motor systems according to the parameters specified by those constructions. In all of these related systems, however, there is a significant prespecification of the required grammatical rules, lexical entries, and construction definitions. In our approach, lexical categories are prespecified, and the system must be trained on form-meaning pairs, from which it can then learn the corresponding constructions.

In conclusion, the current results suggest a possible framework in which grammatical construction processing can be considered in terms of cortical and subcortical processing. The model is demonstrated to accommodate three typologically distinct languages using one common architecture, and is demonstrated to provide a principled account for the processing of embedded compositional phrase structure using the same architecture. We anticipate that only small modifications to the generic "phrasal extraction" mechanism will be required to handle different potentially recursive expansions (e.g., prepositional phrases). Future research should investigate the scalability of this approach.

Open Issues

A crucial issue that any proposed model of sentence processing must address is how that model will scale to variability in the input. One source of variability is the structure of the noun phrase. We have begun to address how NPs that are instantiated by relative phrases as in "The cat that the dog chased scared the boy" can be addressed in a general manner. By using a form of pattern extraction for unpacking relative phrases, the model can handle relative phrases at arbitrary NP positions

in new sentences. We thus illustrated that the processing of relative phrases in Japanese and English (with French being isomorphic to English in this analysis) scales quite well to a large set of sentences borrowed from Miikkulainen (1996). A second source of variability is in the modification of NPs with adjectives, VPs with adverbs, and both with prepositions. In the current implementation we do not address this, so that these modifiers are simply ignored. Our future research will attempt to demonstrate that the regularities governing these phrasal structures can be accommodated by the mechanism proposed here for processing relative phrases.

Another open issue has to do with interactions between syntax and semantics. By fully segregating open and closed classes, we get a powerful generalization capability with a loss of interaction between lexical semantics and syntax. The solution is partially provided in the discussion above concerning item-based constructions, in which open class elements partially disguise themselves as closed class elements so that they can contribute to the definition of the construction, and thus provide an interaction between semantics and syntax. Future research should develop this interaction both at the sentence and discourse levels (Spivey & Tanenhaus, 1998).

In conclusion, the big picture has been characterized in an interesting manner by Bever, Sanz, and Townsend (1998), who said, "The relation between pattern frequencies, semantics and syntax remains the central problem of language processing. Almost any set of simplifying assumptions about how to study that integration and how it works is likely to be incomplete or wrong. The damn thing is probably much more complex than our models will allow." We must clearly acknowledge indeed that language is a complex beast that probably involves a number of mechanisms including lexically driven parsing, function-word template parsing, and more laborious formal methods in parallel. In an effort to move forward, the current research has attempted to characterize a construction-based approach in terms of its neurophysiological, behavioral, and cross-linguistic validity.

Acknowledgments

This study was supported in part by the HFSP Organization, the French ACI NIC and TTT Projects, LAFMI, and the ESF.

Reprint requests should be sent to Peter Ford Dominey, Institut des Sciences Cognitives, CNRS UMR 5015, 67, Boulevard Pinel, 69675 Bron Cedex, France, or via e-mail: dominey@isc.cnrs.fr and www.isc.cnrs.fr/.

REFERENCES

- Bates, E., McNew, S., MacWhinney, B., Devescovi, A., & Smith, S. (1982). Functional constraints on sentence processing: A cross-linguistic study. *Cognition*, *11*, 245–299.
- Baumgaertner, A., Weiller, C., & Büchel, C. (2002). Event-related fMRI reveals cortical sites involved in contextual sentence integration. *Neuroimage*, *16*, 736–745.
- Bergen, B., & Chang, N. (in press). Embodied construction grammar in simulation-based language understanding. In J.-O. Östman & M. Fried (Eds.), *Construction grammar(s): Cognitive and cross-language dimensions*. Johns Benjamins.
- Bever, T. G., Sanz, M., & Townsend, D. J. (1998). The emperor's psycholinguistics. *Journal of Psycholinguistic Research*, *27*, 261–284.
- Blanc, J. M., Dodane, C., & Dominey, P. F. (2003). Temporal processing for syntax acquisition: A simulation study. In R. Alterman & D. Kirsh (Eds.), *Proceedings of the 25th annual meeting of the cognitive science society* (pp. 145–150). Boston: Cognitive Science Society.
- Bresnan, J. (2001). *Lexical-functional syntax*. Oxford: Blackwell.
- Buccino, G., Vogt, S., Ritzl, A., Fink, G. R., Zilles, K., Freund, H. J., et al. (2004). Neural circuits underlying imitation learning of hand actions: An event-related fMRI study. *Neuron*, *42*, 323–334.
- Cartwright, T. A., & Brent, M. R. (1997). Syntactic categorization in early language acquisition: Formalizing the role of distributional analysis. *Cognition*, *63*, 121–170.
- Christophe, A., Nespor, M., Gausti, M. T., & van Ooyen, B. (2003). Prosodic structure and syntactic acquisition: The case of the head-direction parameter. *Developmental Science*, *6*, 211–220.
- Clark, E. V. (2003). *First language acquisition*. Cambridge, UK: Cambridge University Press.
- Cooke, A., Zurif, E. B., DeVita, C., Alsop, D., Koenig, P., Detre, J., et al. (2001). Neural basis for sentence comprehension: Grammatical and short-term memory components. *Human Brain Mapping*, *15*, 80–94.
- Dapretto, M., & Bookheimer, S. Y. (1999). Form and content: Dissociating syntax and semantics in sentence comprehension. *Neuron*, *24*, 427–432.
- Dominey, P. F. (1995). Complex sensory-motor sequence learning based on recurrent state-representation and reinforcement learning. *Biological Cybernetics*, *73*, 265–274.
- Dominey, P. F. (1997). An anatomically structured sensory-motor sequence learning system displays some general linguistic capacities. *Brain and Language*, *59*, 50–75.
- Dominey, P. F. (1998a). A shared system for learning serial and temporal structure of sensori-motor sequences? Evidence from simulation and human experiments. *Cognitive Brain Research*, *6*, 163–174.
- Dominey, P. F. (1998b). Influences of temporal organization on transfer in sequence learning: Comments on Stadler (1995) and Curran and Keele (1993). *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *24*, 234–248.
- Dominey, P. F. (2000). Conceptual grounding in simulation studies of language acquisition. *Evolution of Communication*, *4*, 57–85.
- Dominey, P. F. (2002). A model of learning syntactic comprehension for natural and artificial grammars. In E. Witruk, A. D. Friederici, & Lachmann (Eds.), *Basic mechanisms of language, reading and reading disability* (pp. 61–77). Dordrecht: Kluwer Academic.
- Dominey, P. F. (2005a). Emergence of grammatical constructions: Evidence from simulation and grounded agent experiments. *Connection Science*, *17*, 289–306.
- Dominey, P. F. (2005b). From sensorimotor sequence to grammatical construction: Evidence from

- simulation and neurophysiology. *Adaptive Behavior*, *13*, 347–361.
- Dominey, P. F., Arbib, M. A., & Joseph, J. P. (1995). A model of cortico-striatal plasticity for learning oculomotor associations and sequences. *Journal of Cognitive Neuroscience*, *7*, 311–336.
- Dominey, P. F., & Boucher, J. D. (2005a). Developmental stages of perception and language acquisition in a perceptually grounded robot. *Cognitive Systems Research*, *6*, 243–259.
- Dominey, P. F., & Boucher, J. D. (2005b). Learning to talk about events from narrated video in the construction grammar framework. *Artificial Intelligence*, *167*, 31–61.
- Dominey, P. F., & Hoen, M. (2006). Structure mapping and semantic integration in a construction-based neurolinguistic model of sentence processing. *Cortex*, *42*, 476–479.
- Dominey, P. F., Hoen, M., Blanc, J.-M., & Lelekov-Boissard, T. (2003). Neurological basis of language and sequential cognition: Evidence from simulation, aphasia and ERP studies. *Brain and Language*, *86*, 207–225.
- Dominey, P. F., & Inui, T. (2004). A developmental model of syntax acquisition in the construction grammar framework with cross-linguistic validation in English and Japanese. Available at <http://acl.ldc.upenn.edu/coling2004/W2/pdf/proceedings.pdf>.
- Dominey, P. F., Lelekov, T., Ventre-Dominey, J., & Jeannerod, M. (1998). Dissociable processes for learning the surface and abstract structure sensorimotor sequences. *Journal of Cognitive Neuroscience*, *10*, 734–751.
- Dominey, P. F., & Ramus, F. (2000). Neural network processing of natural language: I. Sensitivity to serial, temporal and abstract structure of language in the infant. *Language and Cognitive Processes*, *15*, 87–127.
- Drummond, S. P., Brown, G. G., & Salamat, J. S. (2003). Brain regions involved in simple and complex grammatical transformations. *NeuroReport*, *14*, 1117–1122.
- Ferreira, F. (2003). The misinterpretation of noncanonical sentences. *Cognitive Psychology*, *47*, 164–203.
- Friederici, A. D. (2002). Towards a neural basis of auditory sentence processing. *Trends in Cognitive Sciences*, *6*, 78–84.
- Goldberg, A. (1995). *Constructions: A construction grammar approach to argument structure*. Chicago: University of Chicago Press.
- Grodzinsky, Y. (2000). The neurology of syntax. *Behavioral and Brain Sciences*, *23*, 1–71.
- Hashimoto, R., & Sakai, K. L. (2002). Specialization in the left prefrontal cortex for sentence comprehension. *Neuron*, *35*, 589–597.
- Hoen, M., & Dominey, P. F. (2000). ERP analysis of cognitive sequencing: A left anterior negativity related to structural transformation processing. *NeuroReport*, *11*, 3187–3191.
- Hoen, M., Golembiowski, M., Guyot, E., Deprez, V., Caplan, D., & Dominey, P. F. (2003). Training with cognitive sequences improves syntactic comprehension in agrammatic aphasics. *NeuroReport*, *14*, 495–499.
- Hoen, M., Pachot-Clouard, M., Segebarth, C., & Dominey, P. F. (2006). When Broca experiences the Janus syndrome. An ER-fMRI study comparing sentence comprehension and cognitive sequence processing. *Cortex*, *42*, 605–623.
- Indefrey, P., Brown, C. M., Hellwig, F., Amunts, K., Herzog, H., Seitz, R. J., et al. (2001). A neural correlate of syntactic encoding during speech production. *Proceedings of the National Academy of Sciences, U.S.A.*, *98*, 5933–5936.
- Inui, T., Otsu, Y., Tanaka, S., Okada, T., Nishizawa, S., & Konishi, J. (1998). A functional MRI analysis of comprehension processes of Japanese sentences. *NeuroReport*, *9*, 3325–3328.
- Kang, A. M., Constable, R. T., Gore, J. C., & Avrutin, S. (1999). An event-related fMRI study of implicit phrase-level syntactic and semantic processing. *Neuroimage*, *10*, 555–561.
- Kuperberg, G. R., Holcomb, P. J., Sitnikova, T., Greve, D., Dale, A. M., & Caplan, D. (2003). Distinct patterns of neural modulation during the processing of conceptual and syntactic anomalies. *Journal of Cognitive Neuroscience*, *15*, 272–293.
- Lelekov, T., Franck, N., Dominey, P. F., & Georgieff, N. (2000). Cognitive sequence processing and syntactic comprehension in schizophrenia. *NeuroReport*, *11*, 2145–2149.
- Longworth, C. E., Keenan, S. E., Barker, R. A., Marslen-Wilson, W. D., & Tyler, L. K. (2005). The basal ganglia and rule-governed language use: Evidence from vascular and degenerative conditions. *Brain*, *128*, 584–596.
- MacWhinney, B. (1982). Basic syntactic processes. In S. Kuczaj (Ed.), *Language development: Vol. 1, Syntax and semantics* (pp. 73–136). Hillsdale, NJ: Erlbaum.
- Marcus, G. F., Vijayan, S., Bandi Rao, S., & Vishton, P. M. (1999). Rule learning by seven-month-old infants. *Science*, *283*, 77–80.
- Marshuetz, C., Smith, E. E., Jonides, J., DeGutis, J., & Chenevert, T. L. (2000). Order information in working memory: fMRI evidence for parietal and prefrontal mechanisms. *Journal of Cognitive Neuroscience*, *12*, 130–144.
- Miikkulainen, R. (1996). Subsymbolic case-role analysis of sentences with embedded clauses. *Cognitive Science*, *20*, 47–73.
- Miyamoto, E. T. (2002). Case markers as clause boundary indicators in Japanese. *Journal of Psycholinguistic Research*, *31*, 307–347.
- Newman, S. D., Just, M. A., Keller, T. A., Roth, J., & Carpenter, P. A. (2003). Differential effects of syntactic and semantic processing on the subregions of Broca's area. *Brain Research, Cognitive Brain Research*, *16*, 297–307.
- Pollard, C., & Sag, I. A. (1994). *Head-driven phrase structure grammar*. Chicago: University of Chicago Press.
- Roy, D. (2002). Learning visually grounded words and syntax for a scene description task. *Computer Speech and Language*, *16*, 353–385.
- Sakai, K. L. (2005). Language acquisition and brain development. *Science*, *310*, 815–819.
- Sanford, A. J., & Sturt, P. (2002). Depth of processing in language comprehension: Not noticing the evidence. *Trends in Cognitive Sciences*, *6*, 382–386.
- Shi, R., Werker, J. F., & Morgan, J. L. (1999). Newborn infants' sensitivity to perceptual cues to lexical and grammatical words. *Cognition*, *72*, B11–B21.
- Siskind, J. M. (1996). A computational study of cross-situational techniques for learning word-to-meaning mappings. *Cognition*, *61*, 39–91.
- Spivey, M., & Tanenhaus, M. (1998). Syntactic ambiguity resolution in discourse: Modeling effects of context and lexical frequency. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *24*, 1521–1543.
- Teichmann, M., Dupoux, E., Kouider, S., Brugieres, P., Boisse, M. F., Baudic, S., et al. (2005). The role of the striatum in rule application: The model of Huntington's disease at early stage. *Brain*, *128*, 1155–1167.
- Tettamanti, M., Buccino, G., Saccuman, M. C., Gallese, V., Danna, M., Scifo, P., et al. (2005). Listening to action-related sentences activates fronto-parietal

- motor circuits. *Journal of Cognitive Neuroscience*, 17, 273–281.
- Tomasello, M. (2000). Do young children have adult syntactic competence? *Cognition*, 74, 209–253.
- Tomasello, M. (2003). *Constructing a language: A usage-based theory of language acquisition*. Cambridge, MA: Harvard University Press.
- Townsend, D. J., & Bever, T. G. (2001). *Sentence comprehension: The integration of habits and rules*. Cambridge: MIT Press.
- Ullman, M. T. (2001). A neurocognitive perspective on language: The declarative/procedural model. *Nature Reviews Neuroscience*, 2, 717–726.
- Ullman, M. T. (2004). Contributions of memory circuits to language: The declarative/procedural model. *Cognition*, 92, 231–270.
- Ullman, M. T. (2006). Is Broca's area part of a basal ganglia thalamocortical circuit? *Cortex*, 42, 480–485.
- Ungerleider, L. G., Courtney, S. M., & Haxby, J. V. (1998). A neural system for human visual working memory. *Proceedings of the National Academy of Sciences, U.S.A.*, 95, 883–890.
- Voegtlin, T., & Dominey, P. F. (2005). Linear recursive distributed representations. *Neural Networks*, 18, 878–895.
- Vosse, T., & Kempen, G. (2000). Syntactic structure assembly in human parsing: A computational model based on competitive inhibition and a lexicalist grammar. *Cognition*, 75, 105–143.
- Yamashita, H. (1997). The effects of word-order and case marking information on the processing of Japanese. *Journal of Psycholinguistic Research*, 26.