

# **From Holophrases to Abstract Grammatical Constructions:**

## **Insights from Simulation Studies**

Peter F. Dominey

The construction grammar framework provides a rich formalism for the characterization of development in language acquisition, in part because it allows for a graceful transition from idiom-like holophrases, to fully abstract argument constructions. From the perspective of learning complexity, this means that rather than mastering the grammar used by an adult to generate or understand an utterance, the child can exploit more direct correspondences between fixed utterances and their meanings in the form of holophrases, and then progressively apply procedures for generalization in order to develop an abstract argument construction repertoire. The current chapter presents a neuro-computational model that addresses aspects of the construction grammar formalism, and uses this model to explain a possible trajectory from idiom-like holophrases to progressively more abstract argument constructions.

The goal of the current research is thus to test a theory of form to meaning mapping that takes substantial input from the proposals of usage based construction grammar (Goldberg 1995, 1998; Croft 2001; Clark 2003, Tomasello 1999, 2003) analogical mapping from sentence to meaning (Fisher 1996), and cue coalitions that specify or identify the mappings (Bates et al. 1982). In particular, the model will provide a functional framework for analysis of the progression of schematicity (Croft 2001) or abstraction (Tomasello 2003) in the evolution of grammatical constructions from fixed structure “holophrases” - i.e. unparsed holistic utterances that correspond directly to a meaning - to fully productive abstract schematic constructions. The

essential idea that will be developed is that there are two complimentary mechanisms for mapping form to meaning. The first, which will be exploited during early development in the holophrase period, involves a direct link between the utterance structure representation and the meaning. The second involves a process that can “cut and paste” arguments into the representation of the utterance structure. The combination of the two mechanisms allows the development of progressively abstract argument constructions in which the form of the construction is processed by the first mechanism, and the insertion of specific arguments into the construction is processed by the second. The advantage of this approach is that the system can get started early with holistic form to meaning mappings before learning about word categories and the possible roles they can play in more abstract constructions.

Within the nativist vs. usage-based debate, this work is positioned near the later. In the nativist context Chomsky (1959) defined a research program for linguistics in which language acquisition was characterized in terms of grammar induction, and significantly circumvented the behaviorist emphasis on meaning. It was in this context that he developed the “poverty of the stimulus” argument, that for learning a grammar, the input to the child was highly underspecified and did not adequately constrain the space of possible grammars. To address this problem, he proposed a Universal Grammar (Chomsky 1965), and defined language acquisition in terms of determining the UG parameters for the target language based on limited input. In this context, proponents of UG tend to argue for the “continuity hypothesis” which holds that there is a functional continuity between the grammar of the child and that of the adult. The child has access to the adult UG and once a given parameter is set, it becomes available in a generalized manner.

Usage-based approaches, from which the current work is inspired, have challenged these assumptions. In this context, research in the developmental psychology of language acquisition

has provided strong arguments for the richness of the stimulus in a usage based framework for language acquisition (Clark 2003, Tomasello 2003). In fact, there are extensive findings to indicate the importance of social pragmatic interaction in the acquisition of a language (reviewed in Tomasello 2003). Far from simply being a question of formal grammar identification, language acquisition involves the construction of a shared space for attention and meaning that allows the child to learn highly specific form-meaning relations that later become generalized. Grammatical constructions define the functional mappings between the surface forms of utterances and the corresponding structural forms of the semantic representations of the meaning of these utterances. With respect to the continuity hypothesis, this usage based framework stipulates that the initial form-to-meaning mappings will be quite specific and non-generalizing, corresponding to the holophrase period attributed both to development and evolution (Wray 2000, Tomasello 2003, Dominey 2004).

This reveals a version of the “discontinuity hypothesis” in which knowledge of grammar is first rudimentary, mapping entire utterances to fixed meanings in holophrase constructions (see Tomasello 2000). Next, free arguments are introduced into these “holophrase” constructions, allowing variation in the specification of agent, object and recipient roles. Later still, the grammatical constructions become truly compositional. Examples of such successively abstract constructions are illustrated in Table 1.

In this context, consider observations from the language production perspective that have been reported by Tomasello and his colleagues and other researchers (reviewed in Tomasello 2003): In the months following their first birthday, in addition to single word utterances, Western middle class children also learn to use unparsed adult expressions such as “I wanna do it” or “Lemme see” as holophrases. By 18 months, children begin to use “pivot schemas” such as “More \_\_\_\_\_” or “\_\_\_\_\_ gone” where the empty slots are filled by open-class or content words.

These soon give way to more generalized item-based constructions including verb islands (Tomasello 1999). These are verb-based constructions (such as “Cut \_\_\_” or “Draw \_\_\_”) in which the slots can be filled in with arguments that are specific for the given verb. Lieven, Pine and Baldwin (1997) observed that between 1 and 3 years of age, children used most of their verbs in only one type of predicative construction. The majority of children’s multiword utterances were produced through the use of a small set of such constructions. By around 4 years of age, these verb-specific constructions are superseded by more productive abstract constructions including the intransitive, transitive, ditransitive, dative, etc.

As form - meaning pairs, constructions include the types listed above, as well as words, idioms, and argument structure constructions that, together, provide the basic means for expression in language (Goldberg 1998). In this communicative context, Goldberg (1995, 1998) proposes a tight correspondence between the structure of perceptual events that are basic to human experience, and constructions for the corresponding basic sentence types. Children’s acquisition of constructions would then require the matching of patterns extracted from the basic scenarios of their human experience, and patterns extracted from sentential forms in their language. The next section defines a model that can perform this matching, and I then provide a walk-through example of sentence processing, with the proper treatment of the transition from holophrase to abstract argument construction.

### **Form to Meaning Mapping (FM<sup>2</sup>) Model**

The model architecture is presented in Figure 1. While this aspect will not be further pursued here, the model was initially developed to explain the functional neurophysiology of abstract (rule based) non-linguistic sequence learning (Dominey et al. 1998), and has generated predictions concerning certain functional relations between language and cognitive sequence

processing that have been supported in behavioral (Lelekov et al. 2000; Hoen et al. 2003), ERP (Dominey et al. 2003; Hoen & Dominey 2004), and fMRI (Hoen et al. 2005) studies.

The model is based on the following principles: (1) Language acquisition can be characterized as learning the mappings from grammatical form to meaning (i.e. grammatical constructions) that should allow productive generalization with the set of constructions learned; (2) Within the sentence, the construction is encoded or identified by the relative configuration or pattern of open and closed class elements, that can thus be used as an index by which the corresponding construction for that sentence type can be learned and retrieved.

The initial point of departure for these simulation studies was to demonstrate how the learning of lexical semantics (word meaning) and phrasal semantics (the use of grammatical structure for assigning lexical items to their appropriate thematic roles) could proceed in parallel and interact. Globally, this corresponds to first learning a few concrete, observable nouns and verbs from cross-situational statistics (Siskind 1996). That is, the system simply associates each word with each perceived entity, and exploits the statistical property that words will be paired more often with their proper referents than with other referents. Then, using this lexical knowledge, the system will begin to extract statistical regularities concerning the relations - in distinct grammatical constructions - between the relative position of open class items in the sentence, and their corresponding referents' roles in the meaning representation. Figure 1 illustrates this graphically, and a deeper explanation will follow. The point is that in this configuration, the system can directly learn abstract argument constructions, without passing through a holophrase or item-based transition period. I will thus first describe the operation of the model in this mode, and will then address the important intermediate steps which are in fact the object of this chapter.

From a behavioral perspective, during learning, the model is presented with <sentence, meaning> pairs, and it should learn the word meanings, and the set of grammatical constructions that define the sentence to meaning mappings in the input training set. During testing, the model should demonstrate that it can use this knowledge to understand new sentences that were not presented in the training set, but were generated from the same lexicon and the same set of grammatical constructions. In particular, the model should demonstrate systematicity, such that words that have only been experienced in particular grammatical roles (e.g., subject in an SVO sentence) will be correctly processed when they appear in new legal grammatical positions (e.g. the same word now used as an object in an SVO sentence).

In the following section, I first describe the model in terms of its full capability to accommodate abstract argument constructions. It will then be used to provide a possible explanation for the trajectory to this state from an initial holophrase state, via an intermediate item-based state.

\*\*\* Insert Figure 1 About Here \*\*\*

**Input Representations:** The two inputs to the system are sentences, and meanings. In general I refer to “sentences”, but the verbal input can actually be more general, and can take the form of sentence fragments, and “ungrammatical” utterances.

**Sentence Input:** Words in the input sentences are presented to the model in their sentential order, and the words are identified by the model as being open or closed class elements. This lexical categorization is among the early language-related perceptual distinctions learned by infants based on perceptible cues in the auditory signal (Shi et al. 1999, Höhle & Weissenborn 2003) along with differences in statistical distributions of words in these

categories (Morgan & Demuth 1996), and is built into the model in the abstract construction processing mode. In this context, we have recently demonstrated that in French and English, the temporal profile of the fundamental frequency ( $F_0$ ) of the speech signal provides a reliable cue for categorizing open and closed class words (Blanc et al. 2003). Related simulation studies in which prosody is symbolically encoded have also demonstrated successful results (Morgan et al. 1996).

The result of this early discrimination capacity applied to the child's target language is subsequently expressed in adulthood. Indeed, in adults, extensive data from event related potentials, brain imagery and psycholinguist studies indicate that these lexical categories are processed by distinct and dissociated neurophysiological pathways (e.g. Kluender & Kutas 1993; Friederici, 1985; Pulvermüller, 1995; Brown et al. 1999).

In the model, words are represented as 25 element vectors<sup>1</sup>. The content (or open class) words will be encoded in the Open Class Array (OCA) that contains 6 fields, each a 25-element vector that encodes one open class word. Word ordering is preserved in the OCA, as illustrated in the example in Figure 1. Closed class words (e.g. was, to, by) are encoded in the Construction Index, a 25 element vector, by an algorithm that preserves the identity and order of arrival of the input closed class elements, and their relative positions with respect to open class elements in the sentence. This is functionally equivalent to the template representation illustrated in Figure 1. Additionally, each sentence is initiated by a closed class start symbol and terminated by a closed class end symbol. The essential role of the ConstructionIndex is to provide a unique identifier for each construction type, based on the cue competition model of Bates et al. (1982).

---

<sup>1</sup> Content words coded with single bit-on in the range 1-16, and function words in 17-25. The single bit coding limits the lexicon to 25 words. The lexicon can be expanded by increasing the vector size, or by using multiple bits per word.

**Meaning Input:** If language acquisition is the learning of a mapping between sentences and meanings, then the infant must have some pre-linguistic capacity for representing this meaning. From this perspective, already at 6 months of age, children are capable of processing causal events with agents, objects and actions, and using these representations to understand simple action scenarios that involve goal-directed reaching for objects (e.g. Woodward 1998). Similarly, infants this age display rather sophisticated knowledge of the physical properties of objects that allows them to "parse" and understand dynamic scenes with multiple objects (Carey and Xu 2000). This implies the existence of conceptual representations of events that can be instantiated by non-linguistic (e.g. visual) perceptual input prior to the development of language. These conceptual representations will form the framework upon which the mapping between linguistic and conceptual structure can be built. This approach requires that at least a primitive conceptualization capability that can deal with events in a predicate-argument format exists in a pre-linguistic state. Fisher (1996) identified the requirement that event representations should take a predicate-argument structure that is related to the grammatical structure of the verb onto which they will be mapped. Similarly, in elaborating the structural relations between linguistic and conceptual forms, Jackendoff considered that predicate/argument structure is a central feature of semantic structures (Jackendoff 2002: 123). This type of abstract predicate/argument event representation is central to the structure to meaning mapping for grammatical constructions as characterized by Goldberg (1995, 1998, 2003).

Thus, the "meaning" onto which the sentence is to be mapped takes this predicate(argument) form, encoded in the Scene Event Array (SEA) of Figure 1 that consists of two sub-arrays so as to accommodate complex events (this configuration is illustrated in Figure 2). Each sub-array contains fields corresponding to *action*, *agent*, *object*, *recipient/source*. Each field is a 25-element vector representing its respective meaning element. The SEA thus allows

representation of simple events (e.g. give, take, push, touch), as well as their combination in hierarchical events, described below. In the example illustrated in Figure 1 the predicate is *gave*, and the arguments corresponding to agent, object and recipient are *Marie, Ball, John*.

**Model Overview and Processing Example:** I now present an overview of the model, and define the representations and functions of each component of the model using the example sentence “The ball was given to Jean by Marie,” and the corresponding meaning “gave(Marie, Ball, John)” in Figure 1. Again the point here is to illustrate how abstract argument constructions can be processed. Once this is done, we will return to the issue of how to get to this point from an initial holophrase state.

The first step in the sentence-meaning mapping process is to extract the meaning of the open class words and store them in the Predicted Referents Array (PRA). The word meanings are extracted from the real-valued WordToReferent matrix that stores learned mappings from input word vectors to output meaning vectors. The second step is to determine the appropriate mapping of the separate items in the PredictedReferentsArray onto the predicate and argument positions of the SceneEventArray. This is the “form to meaning” mapping component of the grammatical construction. PRA items are thus mapped onto their roles in the Scene Event Array (SEA) by the FormToMeaning mapping, specific to each construction type. FormToMeaning is thus a 6x6 real-valued matrix. This mapping is retrieved from ConstructionInventory, based on the ConstructionIndex that encodes the closed class words that characterize each sentence type. The ConstructionIndex is a 25 element vector, and the FormToMeaning mapping is a 6x6 real-valued matrix, corresponding to 36 real values. Thus the ConstructionInventory is a 25x36 real-valued matrix that defines the learned mappings from ConstructionIndex vectors onto 6x6 FormToMeaning matrices.

Again, while the goal of this chapter is to apply this model to the problem of understanding the transition from holophrases to increasingly abstract constructions, the initial goal of this model was to gain insight into how such a system could accommodate a variety of abstract constructions of varying complexity. Thus, I will first demonstrate how the model can handle the general case illustrated with a ditransitive passive construction, and will then back up and show how to get there from holophrases. Thus, consider the sentence "The ball was given to Jean by Marie." As illustrated in Figure 1, the contents of the open class array for this sentence will be "ball, given, Jean, Marie." Retrieval of the corresponding referents from WordToReferent for each of the open class words yields the contents " *ball, given, Jean, Marie*" in the predicted referents array (PRA). With respect to the meaning component, the contents of the scene event array (SEA) will be "*given, Marie, ball, Jean*" for the respective components "action(agent, object, recipient)."<sup>2</sup> Now what has to be learned is the mapping between the PRA and the SEA that constitutes the grammatical construction. Comparing the PRA and the SEA, it can be seen that element 1 (*ball*) of the PRA maps onto element 3 of the SEA. PRA Element 2 (*given*) maps onto SEA element 1, PRA element 3 maps onto SEA element 4, and PRA 4 maps onto SEA 2. This generates the resulting mapping or transformation PRA(1,2,3,4) -> SEA(3,1,4,2) for this passive sentence type. In the model this is represented as the vector transformation in the FormToMeaning matrix that maps PRA onto SEA. When <sentence, meaning> pairs are presented during learning such that the system is familiar with all of the open class words, this mapping can be extracted by observation, and then stored in the ConstructionInventory, indexed by the ConstructionIndex, for future use in generating meanings from input sentences.

---

<sup>2</sup> Italicised words correspond to semantic representations, and plain text corresponds to lexical acoustic input representations.

If there were only one sentence type in the language, then the model would learn only one such FormToMeaning mapping. However, there are numerous sentence types or grammatical constructions in a language, each with its own FormToMeaning (PRA to SEA) mapping. Thus the question arises, how can the system keep track of, and manipulate, distinct constructions. This is where the ConstructionIndex becomes invaluable. A central hypothesis in the model is that each distinct grammatical construction, or sentence type, will have a unique ConstructionIndex based on lexical category, word order and grammatical marking information extracted from the sentence and encoded in ConstructionIndex. During training, the system can thus learn, and store in the associative ConstructionInventory memory, the association between each Construction Index, and its corresponding FormToMeaning map. During post-learning comprehension testing, the ConstructionIndex for the current sentence can then be used to retrieve the corresponding FormToMeaning mapping from the ConstructionInventory.

Once a construction has been learned, the model will be able to understand any novel sentence based on that construction. By definition the novel sentence will have the ConstructionIndex corresponding to its proper construction, allowing the appropriate FormToMeaning map to be retrieved from the ConstructionInventory and applied to the elements of the PredictedReferentsArray. Of course this assumes that the open class elements in the novel sentence are familiar, i.e. that they are represented in WordToReferent. This means that the model generalizes in a systematic manner. Thus, for example, a word learned as the subject of an active SVO sentence can be correctly processed as the object of a relative clause in a new input sentence (if the relativized construction has been previously learned).

As illustrated in Figure 2, in addition to simple <sentence, meaning> pairs such as <"The block pushed the ball", *push(block, ball)*>, we will also consider complex pairs such as <"The block that pushed the ball touched the triangle", *push(block, ball), touch(block, triangle)*> that

employ a relativised sentence and a dual-event scene.

\*\*\* Insert Figure 2 About Here \*\*\*

To accommodate the dual scenes for such complex events, the data structures that perform the form to meaning mapping are instantiated twice each, to represent the two components of the dual scene as illustrated in Figure 2. In the case of simple scenes, the second component of the dual scene representation is null.

I have indicated how knowledge of the mapping of grammatical structure onto event structure requires semantic knowledge of word meaning, corresponding to "semantic bootstrapping." That is, the quality of the FormToMeaning mapping relies on the quality of the PredictedReferentsArray which in turn relies on the WordToReferent mapping. This FormToMeaning knowledge can reciprocally be used to improve the quality of the WordToReferent mapping. The WordToReferent learning algorithm initially associates each open class word with each referent in the scene, and relies on cross situational statistics to yield an advantage for the correct mappings. If the FormToMeaning mapping is known, however, this knowledge can be used so that a given word is only associated with the scene referent (e.g. the agent of the action) that is specified by the FormToMeaning mapping. This "syntactic bootstrapping" will then refine the quality of the WordToReferent mappings, which in turn refines the quality of the FormToMeaning mappings. In this scenario, semantic and syntactic bootstrapping coexist in a cooperative and synergistic manner (Dominey 2000).

The model has been demonstrated to learn a variety (>30) of English grammatical constructions, including the ability to generalize the knowledge of these constructions to sentences not seen during learning, and has also been validated in cross-linguistic manner with

Japanese (Dominey & Inui 2004; Dominey & Boucher 2005). More technical details of the model can be found in Dominey & Boucher (2005) and Dominey (2006).

**From Holophrase via Verb Island to Abstract Construction:** In simulations using the model as described above, the model manipulates sentence to meaning mappings that correspond to abstract argument constructions (Goldberg 1995, Tomasello 2003). However, as outlined above, developmental psycholinguists have conducted a number of experiments that document a developmental transition with respect to the degree to which children can use grammatical constructions with the same abstraction that adult users employ (reviewed in Tomasello 2000, 2003 and Clark 2003). These studies indicate that the abstract construction phase is preceded by the use of fixed, less abstract constructions. In this context, the child's use of a well formed grammatical construction might result from her abstract knowledge of grammatical structure, but it might also arise from repeating a contextually appropriate utterance that the child has heard adults use. This leads one to ask whether the model can provide insight into the human progression from holophrases, via verb islands to the stage of abstract schemas.

In order to address this developmental issue, we must re-examine the mechanism by which grammatical constructions are represented in the model. Recall that each construction in the ConstructionInventory is referenced by a single unique ConstructionIndex that reflects the internal organization of that construction. In the abstract construction framework, open class elements are represented in the ConstructionIndex only in terms of their relative positions with respect to the closed class words, and importantly the actual identity of these open class words is not encoded in the ConstructionIndex. Of equal importance here is the point that we have allowed the system from the outset to have knowledge of open vs closed class elements, and to

treat open class elements such that they can be inserted into well defined "slots" in the learned constructions. Thus, two distinct transitive sentences made up of different open class elements share the same ConstructionIndex, and the same form to meaning mapping. Indeed, it is this capability, described and presented above, that allows the model to generalize in a systematic manner to new sentences of the transitive form (and all other learned grammatical construction forms).

Now let us consider the other extreme of the fixed-abstract continuum, in which the same two transitive sentences are treated as two distinct holophrases, i.e. two distinct constructions in which the entire meaning is directly encoded in the construction, with no additional specification provided by open class words. In this case the system would make no distinction between open and closed class words, and thus both types would directly participate in specification of the ConstructionIndex. For this reason, each sentence would be considered as a holistic and distinct construction. Thus, this corresponds to the situation in which the system fails to treat open class elements as a category of elements that can instantiate variables or slots in abstract constructions. In this context open and closed class words are uniformly treated as part of a holophrase, each of which corresponds to a distinct construction. Functionally, this would consist in open class words being processed both in the Open Class Array, and in the ConstructionIndex. In this manner, the OpenClassArray and ConstructionIndex become bound together to form a single holistic unit, with open class words bound into their grammatical constructions. Under these conditions all distinct input sentences would be processed as distinct constructions.

What then would be the mechanism of transition from holophrase to a more abstract construction type? Consider the hypothetical transition in Table 1 from the holophrase "Gimme the ball" to the verb island construction "Gimme X" where "X" can be instantiated by different concrete noun arguments. There are at least two possible, related mechanisms for how this

transition could take place. If the child is exposed to a suitable number of variants like “Gimme the ball” “Gimme the block,” etc., then a statistical learning mechanism could detect where there is variability in the verbal input, leading to a special “argument” status for this site of high variability. In parallel, with respect to the representation of the corresponding meaning, the child will observe a statistically significant sample of physical events in which different objects can be seized by the adult speaker, thus helping to identify the lexical category of physical objects that can be seized in the “Gimme X” context. This process of categorization leads to the development of different categories of objects that can be acted upon in different ways (Tomasello 2003), as is functionally related to the early learning of concrete nouns (Gillette et al. 1999). Thus, by the cross situational statistics both of the speech input, and of the corresponding physical/perceptual meaning, elements that can fill the object slot in the “Gimme X” construction (corresponding to concrete nouns, but not verbs) will come to be processed as a variable category or "slot" with respect to the ConstructionIndex. More generally, in this intermediate phase, these “filler” nouns would take on a new status in which they would no long contribute to the definition of specific constructions. In terms of the model, they would no longer contribute to the ConstructionIndex. This would lead to an item-based “verb island” phase in which constructions are based around verbs (that continue to contribute their identity to the ConstructionIndex) with free variables for the arguments.

The subsequent emergence of verbs as a functional category (from the operation of cross-situational statistics on the speech and the perceptual representations) and the resulting generalization on verbs would allow for the full abstract construction capability that we have demonstrated above. In this context, the progressive emergence of generalized functional lexical categories is functionally linked with the emergence of progressively more abstract constructions that generalize over those categories. This developmental construction framework was validated

with a series simulations using the sentences based on those in Table 1.

\*\*\*\*\* Insert Figure 3 About Here \*\*\*\*\*

In the first simulation of the “holophrase” developmental phase (Figure 3A), there is no “lexical categorization” and no distinction between open and closed class words. Input processing was structured so that all of the words (open and closed class) directly contributed to the ConstructionIndex. Under these conditions the model learned each sentence, such as “Gimme the ball” as a distinct construction, mapping each construction onto its meaning.

In the second simulation of an item based “verb island” phase (Figure 3B), the input was processed as in the holophrase condition, with the exception that certain nouns – corresponding to those that referred to objects that could be “thrown” were now processed as “open class” words. This results in two important changes. First, these words were no longer directly represented in the ConstructionIndex. Instead, their status as open class words and their relative sentence positions were encoded in the ConstructionIndex as described in the previous section. Second, the meaning of these words is transmitted via the pathway from the OpenClassArray to PredictedReferentsArray via the WordToReferent mapping, and then to the SceneEventArray via the FormToMeaning mapping. This use of these two pathways is illustrated in Figure 3B.

Figure 3C illustrates a further abstraction with a kind of verb-based construction in which “give” can take recipient and object arguments. Under these conditions we would expect constructions to be organized around specific verbs as in the verb island scenario (Tomasello 2003). Indeed, we demonstrated that when verbs were still not processed as open class elements, the model learned 4 distinct verb island constructions corresponding to *touch*, *push*, *take* and

*give*, that could generalize and take different arguments in the agent, object and recipient roles.

Finally, when verbs are processed as open class elements, we arrive at the abstract argument construction stage. Figure 3D illustrates the simulation conditions for a canonical ditransitive construction with *agent, action, recipient, object* ordering in which the open vs. closed class distinction is now completed, and no open class words are encoded in the construction index. Note that in 3D, the absence of any explicit function words still allows the creation of a ConstructionIndex based on the unmarked begin- and end- sentence markers, and the counting of open class words that intercede between them. Figure 3E illustrates a non-canonically ordered ditransitive, in which the object is at the head of the sentence. Comparing Figure 3D and 3E allows us to appreciate how different ConstructionIndices will each invoke their specific FormToMeaning mapping, so that the arguments in the PredictedReferentsArray are properly assigned to their functional roles in the SceneEventsArray. The progression illustrated in A-E of Figure 3 demonstrates one view – based on the model - of the successive transition from processing sentences as holistic units, then as item based constructions, and finally as abstract argument constructions.

## **Discussion**

Before three years of age, a child will not use a verb in a transitive construction if she has not heard it used that way. However, she will regularly use other verbs in transitive forms that she has heard in such forms (reviewed in Tomasello 1999, 2000, 2003). The item-based verb island phases that I simulated (illustrated in Figure 3B and 3C) can help to provide some insight into this phenomenon in the construction grammar framework. Consider the sentences “The boy took the block ” and “The boy took the block to the girl.” After the model has been trained with

sentences such as “The boy took the ball” in the verb island phase (that is, where verbs still contribute to the definition of constructions in the ConstructionIndex), the model will accept a new use of *took* in a transitive, but not in a ditransitive construction. Similarly, if the model has been exposed to sentences with “gave” used in a ditransitive mode, it will accept new ditransitive uses of *gave*. Thus, during this “verb island” phase, the way that verbs can be processed in new sentences depends on how those verbs have been used in the input, with some verbs being accepted as ditransitive, while others are not. Subsequently, once in the abstract phase in which verbs are processed as open class elements, the ditransitive construction becomes “liberated”. That is, it can accept verbs such as *took* as a ditransitive verb even if this ditransitive use of *took* had not been experienced in training.

The analysis of this situation I presented was perhaps a bit mechanistic and leaves some important questions open. In particular, Tomasello (2003) identified two important theoretical questions in this context: (1) What is the nature of the cognitive processes that enable children to build up abstract constructions? And (2) How and why do children make just the generalizations they do, appropriately constraining their generalizations? We can now consider the first question on the nature of the cognitive processes that enable children to make this transition from holophrases to abstract argument structure constructions. These simulations indicate that at least two processes are necessary. The first corresponds to determining the elements in the construction that can be instantiated by other elements, i.e. the abstract arguments like X in “Gimme X”. The second corresponds to an abstract structure-mapping capability in which, once these variables have been instantiated by specific arguments, they are then manipulated according to an abstract mapping, e.g. that X in “Gimme the X” can be instantiated by “ball” and then mapped onto the object role of *give*.

The first cognitive function required involves deciding what goes into the variable slots, i.e.

over what set can we generalize? As mentioned above, when the category of concrete nouns becomes abstracted, this provides the basis for the partially productive item-based constructions, illustrated by the verb island constructions. The subsequent abstraction of the category of verbs allows the increasingly productive abstract constructions to emerge. Thus, the development of these functional categories may provide insight into the determination of what is a generalized variable category, when and why. Note that in the current simulations we linked this to an ordered development of the categories concrete noun/thing and verb/event, though in the more complex human "model" the distinctions may be different. This development may be directly linked to the development in vocabulary growth for these categories. These issues have also been addressed in terms of the functionally based distributional analysis that initially yields local functional categories that later give way to more abstract paradigmatic categories (Tomasello 2003). Thus, rather than relying exclusively on external lexical category formation, we can also look for variability in the input, in terms of which elements in a sentence vary. These varying elements will be candidates for arguments in progressively more abstract constructions. The point is that during the transition phases when different words become recognized as elements that can fill variable-slots in constructions there will be transition item-based constructions in which certain nouns will only appear in constructions with specific verbs, leading to item-based phenomena such as verb islands.

In this context of learning and categorization, two related studies on 7-9 month old children's ability to extract structure from sound sequences are of interest here. Saffran et al. (1996) showed that brief exposure (2 minutes) to repeating sound sequences was sufficient for 9 month olds to learn the repeating sequences, and to distinguish them from sequences made up of the same sounds presented in a different order. Keep this in mind as a mechanism for constructing holophrases. In a related study, Marcus et al. (1999) exposed 7 month old children to similar

sound sequences. However, this time, the sequences were constructed according to an “abstract” rule, of the form ABB (or AAB). During the two minutes, different sound sequences adhering to the rule were presented, and after this exposure the children could distinguish between new sequences that obeyed or violated the rule. Together these two experiments provide interesting data on the children’s ability to perceive and represent variability in the input. Where the variability in the input is low with respect to the words themselves and their serial order, the infants will bind the words together in a “holophrase” as revealed in the Saffran et al. (1996) experiments. However, when the variability in the words is high, but low in the relations between words, infants will bind a variable slot into the representation, as revealed in the Marcus et al. (1999) experiments. Thus the operation of these types of mechanisms on the verbal input should also contribute to the formation of holophrases and the transition to the insertion of abstract arguments, based precisely on the statistical pattern of variability in the input.

Interestingly, the second requirement - an abstract mapping capability - is functionally related to the abstract sequence manipulation capability demonstrated by 7-month old infants in the "ABB" vs. "AAB" experiments of Marcus et al. (1999). In a series of simulation studies, we determined that the Saffran-like sequence learning could be achieved by a recurrent network that was initially developed to simulate the neural activity in the prefrontal cortex of non-human primates in a sequence repetition task (Dominey et al. 1995). This model failed, however to perform tasks that required the Marcus-like representation, with which sequences such as ABCBAC and DEFEDF would be represented as sharing the abstract structure 123213. This required the addition of a working memory for storing recent sequences of elements and comparing them with new elements in order to detect and encode repetitive abstract structure (Dominey et al. 1998). We thus demonstrated that the original recurrent network simulated the Saffran results but failed on the Marcus results, which required an abstract structure processing

model (Dominey & Ramus 2000).

Perhaps more interestingly, we subsequently demonstrated the capability of the same abstract structure processing architecture to perform thematic role assignment where open class sentence elements are mapped onto their canonically structured scene referents (Dominey et al. 2003), in a manner functionally equivalent to that in the current model in Figure 1. To illustrate, consider the following sentences

- (1) John gave the ball to Mary.
- (2) The ball was given to Mary by John.

For both of these sentences the agent, object, recipient are respectively John, Ball, and Mary. In a behavioral task (Caplan et al. 1985) in which the participant should indicate the agent, object and recipient, in that order for sentence (1), this response sequence can be generated simply by repeating the nouns in the same order in which they occurred in the input. We can thus characterize the structure of the transformation from input to output as JBM-JBM, or 123-123. Alternatively, in the analogous type of analysis for sentence (2) the input noun order does not correspond to the required agent, object, recipient response, and thus must undergo a structural transformation corresponding to BMJ-JBM. Viewed in this manner, the thematic role assignment task of Caplan et al (1985) can be reformulated behaviorally, and simulated, as an abstract cognitive sequencing task, in which subjects must learn the appropriate abstract structure for each of the nine different grammatical constructions used in the task, and the different constructions are discriminated based on their closed class element configuration (Dominey et al. 2003, see also Chang 2002).

This suggests that the required abstract structure mapping mechanism is not language

specific, and is part of the general inventory of cognitive mechanisms available to the higher primates (Hauser, Weiss & Marcus 2002). This type of rule based structure mapping has also been extensively elaborated and investigated in diverse structure-mapping contexts by Gentner and her colleagues (e.g. Gentner & Medina 1998), and plays an important role in Fisher's (1996) analogical mapping architecture for language learning.

The second question concerns how and why children make the generalizations they do, constraining their abstract constructions to (and not beyond) their proper application. The complete answer to this question is clearly beyond the scope of the current research. However, I see the inherent conservatism of a construction-based system such that at the level of abstract argument constructions, the freedom available within the appropriate argument categories provides productive generalization, without an unbounded generation of new, potentially faulty, construction types. That is, after exposure to a training corpus, the model will generalize over the constructions represented in the corpus. It will be able to process sentences corresponding to those in these constructions, but nothing further.

With respect to generalization, this situation is at the same time the strength and the weakness of the current level of generativity and abstraction. Generalization can occur productively, but only within the bounds of previously learned constructions. The strength is that this conservative generalization is nonetheless a generalization, and it is constrained not to go beyond the bounds of acceptability in the target language. The weakness, of course, is that new construction types can only be added by learning from a well formed <Sentence, Meaning> pair. In the current formalism the system cannot accommodate new constructions on the fly, but must see at least one example before generalizing. Again, this appears to be consistent with a developmental period around 2 to 3 years of age, when children's utterances are primarily derived from constructions attested in speech they have heard (Clark 2003).

In this context, however the current system fails to understand novel constructions through dynamic construction building, since a new construction must be heard once in order to establish an appropriate mapping. This is partially related to the idea that new constructions, e.g. a sentence with a relative clause, are learned as entirely new constructions, with no connection to any existing constructions from which they are (or rather, could be) built - so it lacks a certain kind of pattern detection. The solution will likely rely on an appropriate use of statistical pattern detection so that, for example, the system will realize that encountering the word "which" indicates that what follows is often a relative clause, so that relative clauses in novel positions could be handled (Miikkulainen 1996). Thus, pattern detection would occur not just at the sentence level (as it currently does in the formation of the ConstructionIndex) but also at the level of noun phrases. In the case of the relative clause, after learning that a relative clause can occur in an NP, the model would then be capable of correctly processing a relative clause in any legal NP position in a new construction. This modularization of abstract construction management is likely the next step in the developmental sequence from holophrases, via item constructions to abstract constructions, and on to modular abstract constructions.

In conclusion, I hope that this theoretical simulation-based analysis can provide some useful insight and/or points of discussion or reflection for researchers in language acquisition. It is clear that many important aspects of language development (e.g. the contribution of bound verb morphology to the ConstructionIndex) are glossed over or entirely missing. Hopefully, however, what has been included is relevant, and will provide the basis for the development of more accurate models of language processing.

**Acknowledgements:** This work was supported in part by the ESF OMLL, the HFSP, the French ACI NIC and TTT Projects.



## 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.
- Blanc, J.M., Dodane, C., & Dominey, P.F.( 2003). Temporal Processing for Syntax Acquisition: A simulation study. *25<sup>th</sup> Annual Meeting of the Cognitive Science Society*. Boston, MA.
- Brown, C. M., Hagoort, P., & ter Keurs, M. (1999) Electrophysiological signatures of visual lexical processing : Open- and closed-class words. *Journal of Cognitive Neuroscience*. 11 , 261-281
- Carey, S., Xu, F. (2001) Infant's knowledge of objects: beyond object files and object tracking. *Cognition*, 80, 179-213.
- Chang, F. (2002) Symbolically speaking : a connectionist model of sentence production. *Cognitive Science* 93, 1 – 43.
- Chomsky, N. (1959) A Review of B. F. Skinner's *Verbal Behavior*." *Language*, 35, 26-58.
- Chomsky, N. (1965) *Aspects of a theory of syntax*. Cambridge: MIT Press.
- Clark, E.V. (2003) *First language acquisition*. Cambridge: Cambridge University Press.
- Croft, W. (2001) *Radical construction grammar: syntactic theory in typological perspective*. Oxford: Oxford University Press.
- Dominey, P.F. (2000) Conceptual grounding in simulation studies of language acquisition, *Evolution of Communication*, 4, 57-85.
- Dominey, P.F. (2003) Learning grammatical constructions in a miniature language from narrated video events, Proceedings of the 25<sup>th</sup> Annual Meeting of the Cognitive Science Society, Boston, MA. PAGES. PLACE: PUBLISHER

- Dominey, P.F. (2004) From holophrases to abstract grammatical constructions in development and evolution, *Evol Lang 2004*, Leipzig. [spell out what this is ??] AND SO ON
- Dominey, P.F. (2005) TBD
- Dominey, P.F., Arbib MA, Joseph JP (1995) A Model of Cortico-Striatal Plasticity for Learning Oculomotor Associations and Sequences. *J Cog Neuroscience*, 7:3, 311-336
- Dominey, P.F., Boucher, J.D. (2005) Developmental stages of perception and language acquisition in a perceptually grounded robot, In press, *Cognitive Systems Research*
- Dominey, P.F., Hoen M (2005) Structure Mapping and Semantic Integration in a Construction-Based Neurolinguistic Model of Sentence Processing, In Press, *Cortex*
- Dominey, P.F., Hoen M, Lelekov T, Blanc JM (2003) Neurological basis of language in sequential cognition: Evidence from simulation, aphasia and ERP studies, *Brain and Language*, 86 (2003) 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, Proceedings of the CoLing Workshop on Psycho-Computational Models of Language Acquisition, Geneva, pp. 33-40.
- 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 :6 734-751.
- 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(1) 87-127
- Fisher, C. (1996) Structural limits on verb mapping: The role of analogy in children's interpretation of sentences. *Cognitive Psychology*, 31, 41-81

- Friederici, A.D. (1985) Levels of processing and vocabulary types: evidence from on-line comprehension in normals and agrammatics. *Cognition* 19, 133-166.
- Gentner, D., Medina, J. (1998) Similarity and the development of rules. *Cognition*, 65, 263-297.
- Gillette, J., Gleitman, H., Gleitman, L. and Lederer, A. (1999) Human simulation of vocabulary learning. *Cognition*, 73, 135-151.
- Goldberg, A. (1995) *Constructions*. University Chicago Press, Chicago and London.
- Goldberg, A. (1998) Patterns of experience in patterns of language, In (Michael Tomasello Ed) *The new psychology of language*, Vol 1, 203-19
- Goldberg, A. (2003) *Constructions: a new theoretical approach to language*, Trends in Cognitive Science, Volume 7, Issue 5 , May 2003, Pages 219-224
- Hirsh-Pasek, K., Golinkof, R.M. (1996) *The origins of grammar: evidence from early language comprehension*. MIT Press, Boston.
- Hoen, M., Pachot-Clouard, M., Segebarth, C., Dominey, P.F. (2005) When Broca experiences the Janus syndrome. An er-fMRI study comparing sentence comprehension and cognitive sequence processing. In Press, *Cortex*
- Höhle, B., Weissenborn, J. (2003) German-learning infants' ability to detect unstressed closed-class elements in continuous speech, *Developmental Science*, 6 (2) 122-127.
- Jackendoff, R. (2002) *Foundations of Language: Brain, Meaning, Grammar, Evolution*. Oxford University Press. 2002
- Kluender, R., Kutas, M. (1993). Bridging the gap : evidence from ERPs on the processing of unbounded dependencies. *J Cog Neurosci* 5:196-214.
- Kotovsky, L., Baillargeon, R. (1998) The development of calibration-based reasoning about collision events in young infants, *Cognition*, 67, 311-351

- Lieven, E., Pine, J., Baldwin, G. (1997) Lexically-based learning and early grammatical development. *Journal of Child Language*, 24, 187-220.
- Marcus, G.F., Vijayan, S., Bandi Rao, S., Vishton, P.M. (1999) Rule learning by seven-month-old infants. *Science* Jan 1;283(5398):77-80
- Morgan, J.L., Demuth, K. (1996) *Signal to syntax: bootstrapping from speech to grammar in early acquisition*. Lawrence Erlbaum, Mahwah NJ, USA.
- Pinker, S. (1984) *Language learnability and language development*, Cambridge, MA, Harvard University Press
- Pinker, S. (1987) The bootstrapping problem in language acquisition. In B. MacWhinney, ed., *Mechanisms of language acquisition*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Pulvermüller, F. (1995) Agrammatism: behavioral description and neurobiological explanation. *J Cog Neuroscience*, 7, 2, 165-181.
- Roy, D., Pentland, A. (2002). Learning Words from Sights and Sounds: A Computational Model. *Cognitive Science*, 26(1), 113-146.
- Shi, R., Werker, J.F., Morgan, J.L. (1999) Newborn infants' sensitivity to perceptual cues to lexical and grammatical words, *Cognition*, Volume 72, Issue 2, B11-B21.
- Siskind, J.M. (1996) A computational study of cross-situational techniques for learning word-to-meaning mappings, *Cognition* (61) 39-91.
- Skinner, B.F. (1957) *Verbal Behavior*, Reprinted by the B. F. Skinner Foundation in 1992 and 2002
- Steels, L. (2001) *Language Games for Autonomous Robots*. IEEE Intelligent Systems, 16(5) 16-22, New York: IEEE Press.
- Tomasello, M. (1999) The item-based nature of children's early syntactic development, *Trends in Cognitive Science*, 4(4) :156-163

Tomasello, M. (2000) Do young children have adult syntactic competence? *Cognition* 74 (2000) 209-253

Tomasello, M. (2003) *Constructing a language: A usage-based theory of language acquisition*. Harvard University Press, Cambridge.

Woodward, A.L. (1998) Infants selectively encode the goal object of an actor's reach. *Cognition* 69 1-34.

Wray, A. (2000) A protolanguage with no declaratives and no names. *Evol Lang* 2000, Paris.

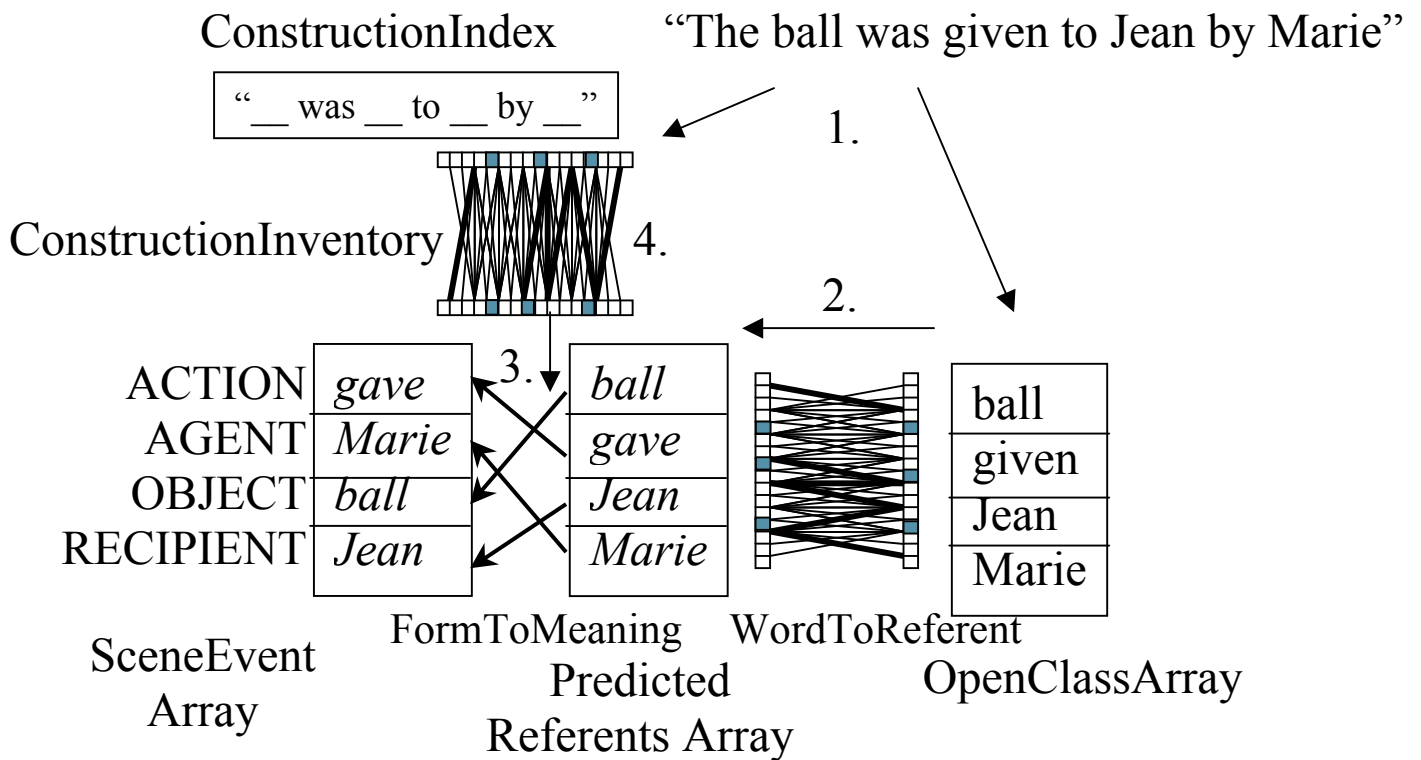


Figure 1. Model Overview: Processing of active and passive sentence types in A, B, respectively. 1. On input, open and closed class words are segregated. Open class words populate the Open Class Array (OCA), while closed class words populate the ConstructionIndex. Visual Scene Analysis populates the Scene Event Array (SEA) with the extracted meaning as scene elements. 2. Words in OCA are translated to Predicted Referents via the WordToReferent mapping to populate the Predicted Referents Array (PRA). 3. PRA elements are mapped onto their roles in the Scene Event Array (SEA) by the FormToMeaning mapping, specific to each sentence type. 4. This mapping is retrieved from Construction Inventory, via the ConstructionIndex that encodes the closed class words that characterize each sentence type. Words in sentences, and elements in the scene are coded as single ON bits in respective 25-element vectors.




Figure 2. Processing of relative sentences. The relativized input sentence describes two distinct events that are linked by the common argument, ball. Thus, by adding a second Event representation in the SceneEventArray component, the meaning structure onto which the sentence can map is made available. The ConstructionIndex thus becomes associated with a pair of mapping structures, one that maps elements from the PRA onto the first event representation, and the second that maps elements from the PRA onto the second event representation..

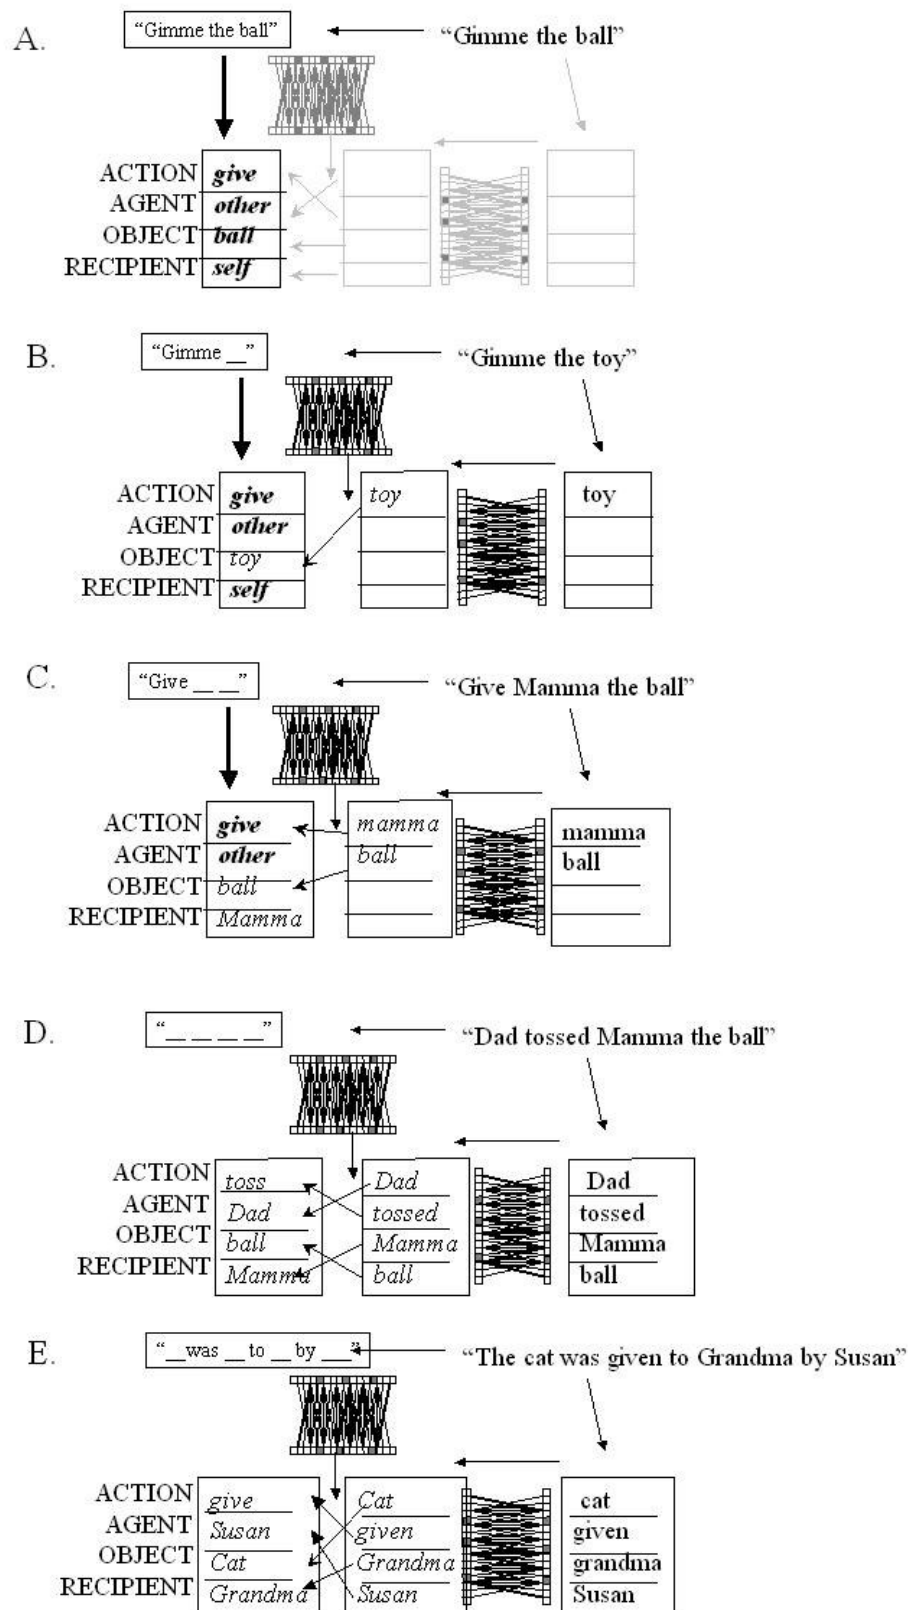


Figure 3. Construction templates for Holophrase and Verb Island phases. A. In the holophrase, the utterance itself defines the construction, with a direct correspondence between the utterance and the meaning. B. In the verb island, free variables are introduced for the nouns corresponding in this case to the agent, object and recipient thematic roles. Nouns no longer contribute to the identity of the construction (i.e. the ConstructionIndex), while verbs continue to do so.

Example	Meaning	Construction Type
	Event(Agent, Object, Recipient)	
1. Gimme the ball.	Give(Caretaker, ball,child)	Holophrase
2. Gimme <i>object</i> 2.1 Gimme the block	Give(Caretaker, <i>object</i> ,child)	Item Based Constr:
3. Give <i>recipient object</i> 3.1 Give John the block	Give(__, <i>object,recipient</i> )	Item Based Constr
4. You give <i>recipient object</i> 4.1 You give John the block	Give(Caretaker, <i>object,recipient</i> )	Item Based Constr
5. <i>agent</i> give <i>recipient object</i> 5.1 <i>agent</i> give <i>recipient object</i>	give( <i>agent,object,recipient</i> )	Item Based Constr
6. <i>Trans recipient object</i> 6.1 Throw John the ball.	<i>Trans(__, object, recipient)</i>	Abstr Arg Constr
7. <i>agent Trans recipient object</i> 7.1 Mary throws John the ball.	<i>Trans(agent,object,recipient)</i>	Abstr Arg Constr
8. <i>agent Trans object to recipient</i> 8.1 Mary throws the ball to John.	<i>Trans(agent,object,recipient)</i>	Abstr Arg Constr
9. <i>object was Trans to recipient</i> <i>by agent</i> 9.1 The ball was thrown to John by Mary.	<i>Trans(agent,object,recipient)</i>	Abstr Arg Constr

Table 1.

Table 1 Caption: Illustration of example constructions that vary in their respective levels of abstraction. The Example Column displays examples of the different construction types, and in the case of constructions with abstract arguments, specific examples are also presented. Terms expressed in italics correspond to abstract variable elements that can be replaced or instantiated by specific appropriate open class elements. The Meaning Column displays the corresponding meaning representations in a *Predicate(Agent, Object, Recipient)* format. The Construction Type Column gives a global categorization of three different levels of construction abstraction. Illustration of a hypothetical progression from holophrase to abstract non-canonical argument construction.