BUCLD 35 Proceedings
To be published in 2011 by Cascadilla Press
Rights forms signed by all authors

Frozen Chunks and Generalized Representations: The Case of the English Dative Alternation

Erin Conwell, Timothy J. O'Donnell, and Jesse Snedeker

1.Introduction

A fundamental and celebrated property of language is its productivity: using language we can understand and express an unbounded number of novel thoughts. However, it has long been acknowledged that this productivity presents a fundamental problem for the learner (Baker, 1979). How is the learner to know which of the many *possible* generalizations consistent with the data are actually productive in their own language? How can learners strike the balance between regularity and idiosyncrasy?

Two distinct theoretical approaches to this problem make different assumptions about the starting state of the learner and predict different developmental trajectories. One approach assumes that a rich and abstract linguistic system is available very early in development, perhaps even at birth. The learner's task is to tailor this system to the particular language being learned. We will call this the "early generativity" approach. There are many different versions of early generativity theories (for examples, see Chomsky, 1981; Pinker, 1989; Fisher, 2002; Snyder, 2007). However, these approaches are unified by the assumption that abstractions of approximately the right grain--such as "noun," "verb," "subject" and "object"-- are available to the learner from an early age and that the role of linguistic input is simply to indicate how these abstractions are manifested in the particular language being learned. The fundamental learning problem for theories that posit early generativity is how the child *constrains* the potential productivity in the system. Such theories provide a starting state that accounts for the productive linguistic processes in *all* of the world's languages, but must also provide a way for the learner to determine which *particular* processes are productive in her own language.

Another class of approaches assumes that the learner begins with no linguistic abstractions but instead gradually forms these generalizations from the linguistic input using domain-general cognitive and social abilities. These theories fall under the umbrella of "usage-based" approaches (e.g., Tomasello, 2003; Lieven, Behrens, Speares & Tomasello, 2003; Goldberg, 2003; Braine, 1963). Such approaches typically emphasize the gradual emergence of

*Department of Psychology, Harvard University. Corresponding author is EC, econwell@gmail.com.

piecemeal generalizations, resulting in stages of development in which children's grammatical representations are based on specific lexical items or fixed slot-and-frame type constructions rather than grammatical abstractions such as "noun" and "verb." These approaches try to avoid the potential problem of overgeneralization by hypothesizing that children only generalize when such generalizations are strongly licensed by the data (Tomasello, 2000).

Both the early generativity and the usage-based approaches are aimed at dealing with the fundamental linguistic tension between productive computation, on one hand, and reuse of both language and item-specific idiosyncratic structure, on the other. The early generativity approach emphasizes the importance of providing the right kinds of generalizations; while the usage-based approach emphasizes the necessity of licensing such generalizations from the data.

The present paper uses the dative alternation as a test case for exploring these two theoretical approaches. We focus on the dative because there is considerable evidence on children's comprehension and production of datives, particularly at around 3 years of age, and because it has been studied extensively by proponents of both theories. First, we review this prior work. Then we conduct a corpus analysis to characterize children's experience with the dative alternation, focusing on how the distributions of arguments in the input could account for the prior data on early acquisition and use.

Finally, we provide a computational analysis of the dative alternation using a Bayesian model known as *Fragment Grammar* (O'Donnell, et al. 2009; O'Donnell, 2011). Fragment Grammar is a general model of productivity and reuse which treats the balance of these two factors as an inference problem in its own right. The model can be seen as asking the question: what would an agent who was making optimal use of information conclude about productivity and reuse from the patterns in the input data? We find that when this model is trained on a corpus of datives for child-directed speech, it correctly predicts the patterns of generalization reported in the prior studies. Critically, the model limits its generalizations in the same ways as young children, even though its input is encoded using abstract syntactic and semantic categories. In the conclusion, we discuss what these results imply for both early generativity and usage-based theories.

2. The Dative Alternation

The English dative alternation has received extensive attention in the language development community, both from theoretical and empirical perspectives. Dative verbs describe events of transfer and typically take two post-verbal arguments: the theme (that which is transferred) and the recipient.

Many dative verbs appear in both the prepositional dative (1a) and in the double object dative form (1b). However, not all verbs that appear in one structure are able to appear in both (2a, b).

(1) a. Okijima gave the ball to Papelbon.
b. Okijima gave Papelbon the ball.
(2) a. Tito explained the strategy to Beckett.
b. *Tito explained Beckett the strategy.

Children's processing and use of the dative alternation at the age of 3 years has been studied using a range of methodologies. Corpus studies of datives in speech to and by children find that the double object form is much more frequent in natural production (Snyder & Stromswold, 1997; Campbell & Tomasello, 2001; Viau, 2007). In comprehension studies, 3-year-old children show verb-general priming of both dative forms (Thothathiri & Snedeker, 2008). In studies using novel verbs, children of this age often fail to understand double object datives containing two full noun phrases, but successfully comprehend prepositional datives with novel verbs (Rowland & Noble, *in press*). In novel verb elicitation studies, 3-year-olds productively use the prepositional dative form, but are conservative in their use of novel verbs in the double object dative (Conwell & Demuth, 2007).

Neither of the two theoretical approaches to argument structure acquisition discussed above can gracefully account for the full pattern of empirical findings. The robust generalization of the prepositional dative to novel verbs and the verb-general priming for both dative structures are consistent with the early generativity approach. Both findings indicate that young children must have somewhat abstract representations of these structures. However, the early generativity approach fails to explain why young children's comprehension and production of datives with novel verbs is robust for the prepositional dative, but quite weak for the double object form—that is, why the two constructions generalize differently. A usage-based approach might posit that abstract representation is constructed more rapidly for the prepositional dative and the double-object datives still rely on verb-specific representations (e.g., Rowland & Noble, *in press*). However, verb-general priming is considered strong evidence for abstract representation of a given structure (Pickering & Branigan, 1998) and 3-year-old children do show verb-general priming of both structures. Thus neither approach neatly captures all of these results.

We suggest that the nature of the evidence that children receive from their linguistic input may account for these results by simultaneously supporting representations that are item based and representations that are abstract. Our evidence for this is two-pronged. First, we will present a corpus analysis showing that children's experience with the dative alternation may support the storage of lexically-specific representations for the double object dative and the storage of more abstract templates for the prepositional dative. Then we will use Fragment Grammar to demonstrate that an ideal learner, given naturalistic input, will, in fact, arrive at this mixture of abstract and item-based forms. Such a mixture predicts the empirical results on the dative alternation at age 3.

3. Corpus Analysis

Corpus analyses of the acquisition of the dative alternation uniformly find that the prepositional dative is less frequent than the double object dative in speech to young children by a ratio of approximately 3:1. Furthermore, these studies typically find that the double object form appears earlier in children's spontaneous speech (Campbell & Tomasello, 2001; Snyder & Stromswold, 1997; Viau, 2007). None of these studies, however, has looked closely at the patterns of argument use in each structure. Previous work finds that argument variation predicts argument structure learning by children (Naigles & Hoff-Ginsberg, 1998). Perhaps a closer examination of the kinds of arguments used in each dative alternation will provide insight into the patterns of comprehension and production of dative verbs at age 3.

3.1 Corpora

Our corpus analysis examined child-directed speech from the 6 mothers in the Providence Corpus (Demuth, Culbertson & Alter, 2006). This corpus begins with each child's first words and continues for two years (approximately ages 1;0-3;0). All children are native English-learners and their parents are native speakers of English. For all children, mother is primary caregiver and interlocutor.

3.2 Procedure

To identify all potential dative utterances, we compiled a list of possible dative verbs based on Levin (1993). All maternal utterances containing those verbs were extracted from the corpora and hand-coded for type of dative (PP or DO) and for type of first post-verbal argument (pronoun, proper noun or full noun phrase).

3.3. Results

At a coarse level, the distribution of prepositional and double object datives in these corpora replicates previous results. Of dative utterances, 72% took the double object form, while 28% were prepositional.

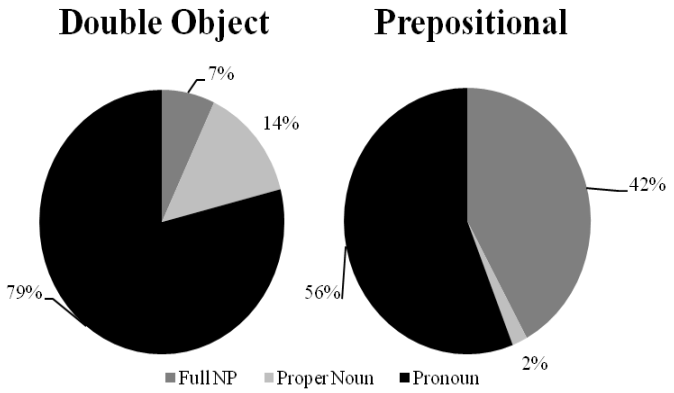

Figure 1. Frequency of types of first post-verbal arguments seen in each dative structure.

An examination of the types of first post-verbal argument used in each form, however, shows that each structure has a different level of argument variation. Figure 1 shows the distribution of first post-verbal arguments in prepositional and double object datives. We consider the first post-verbal argument because it has the greatest potential for outright lexicalization with the verb. Additional data on the distribution of second post-verbal arguments is being collected. For both forms, the majority of first post-verbal arguments are pronouns. However, this trend is weaker for the prepositional dative than for the double object dative. Also, the prepositional dative is more likely than the double object form to have a full noun phrase immediately following the verb. Because pronouns are a closed class while the number of full noun phrases is potentially infinite, these results suggest that prepositional datives have a great deal more variation in their first post-verbal arguments than do double object datives. Can this differential variability in argument type help to account for the empirical data on the nature of dative representations at the age of 3?

Our corpus analysis finds that the prepositional dative form appears with a wide range of first post-verbal arguments, including a large proportion of full noun phrases. This diversity may help children to conclude relatively early in development that an abstract and general version of the prepositional dative construction is usable in novel contexts, even though it is less frequent than the double object form. On the other hand, the reduced variability of first postverbal arguments in the double object dative may cause children to assume that this structure has more limited productivity. Indeed, this is consistent with evidence that 3-year-old children show asymmetric productivity with the prepositional and double object forms. These children may have abstract representations of both forms, but the double object is so dominated by a few arguments that they do not assume real productivity in that form. Rather, their early double object

utterances may be based primarily on highly frequent, routinized forms (e.g., "give me X" or "show you Y") because they believe that this construction is primarily usable in these restricted contexts.

Abstract grammatical representations and item-based or collocated forms need not be mutually exclusive (Jackendoff, 2003; Fisher, 2002; Goldberg, 1995). Some grammatical forms in adult speech are highly productive while others are more routinized or collocated. Children's grammars might be based on a similar mixture of abstract forms that are present early but are over-ridden by stored or item-based structures. The challenge, then is to find the balance productivity and storage to best explain the pattern in their language, which in turn will inform their conclusions about the grammar they are learning.

4. Computational Model

In this section we give an intuitive overview of the Fragment Grammar model. It is beyond the scope of the present paper to present mathematical details, but a full exposition can be found in (O'Donnell, 2011; O'Donnell et al., 2009). The Fragment Grammar model treats the problem of productivity and reuse as an inference in a Bayesian framework. Fundamentally, the model asks the question: what patterns in the data imply that structures will be reused and what patterns in the data imply that future novelty and variability is likely?

A productive computation is one which can give rise to novel forms. In a Bayesian setting, if a system hypothesizes that some (sub)computation is productive, it must reserve probability mass for hitherto unseen structures. On the other hand, if a Bayesian system hypothesizes that some sequence of computations will be reused together, it must reserve probability mass to that particular sequence as a whole. Since there is only a finite budget of probability, this necessarily leads to a trade-off: a probabilistic system hypothesizes reusability at the cost of generalization and productivity at the cost of reusability. The Fragment Grammar model can be seen as optimizing this tradeoff for a given dataset. Importantly, the currency which is optimized is not the cost of computation or memory; rather the system optimizes its ability to correctly predict which computations will generate novel structures in the future and which sequences of computations will be reused again as units.

The learner starts with an initial computational system that is able to handle all of the possible generalizations in the domain of interest. For the present work the starting computational system is formalized as a probabilistic context-free grammar (PCFG). [2] A PCFG for the verb-phrase structure of the dative alternation is shown in Figure 2. The values θ represent the probability of each rule.

[2] In O'Donnell (2011), we show how these techniques can be generalized to arbitrary probabilistic generative systems.

θ_{VP_1} VP $\longrightarrow$ V NP–*recipient* NP–*theme*
θ_{VP_2} VP $\longrightarrow$ V NP–*theme* PP
θ_{V_1} V $\longrightarrow$ *give*
θ_{V_2} V $\longrightarrow$ *show*
... ... $\longrightarrow$ ...
θ_{PP} PP $\longrightarrow$ *to* NP–*recipient*
... ... $\longrightarrow$ ...
θ_{NP_1} NP $\longrightarrow$ *me*
θ_{NP_2} NP $\longrightarrow$ *you*
θ_{NP_3} NP $\longrightarrow$ *Mommy*
... ... $\longrightarrow$ ...

Figure 2. A probabilistic context-free grammar (PCFG) for the dative alternation.

Note that this PCFG suppresses much syntactic and semantic structure which is irrelevant to the dative alternation, while making other kinds of information explicit, such as the thematic role of each argument. These choices were deliberate. We wished isolate the problems of determining the relative productivity of the double object and prepositional dative constructions, and whether each showed evidence for item-specific reuse. To do this we made available to the model all and only the information relevant to these two constructions. We, of course, do not propose this particular computational system as a realistic starting state for the child learner. Under any theoretical account much of the information made explicit in this system of rules, such as the thematic roles of arguments, would have to be learned or induced by the child. Likewise, there are many unmodeled aspects of verb-argument structure which are important such as selectional restrictions on subjects. However, by adopting this aggressive simplification of the domain we will be able to show that the patterns in the input data predict the right kinds of reuse of double object and prepositional dative constructions *despite* the bias in favor of productive generalization of both in the starting state. A PCFG gives rise to the familiar tree-shaped computations shown in Figure 3.

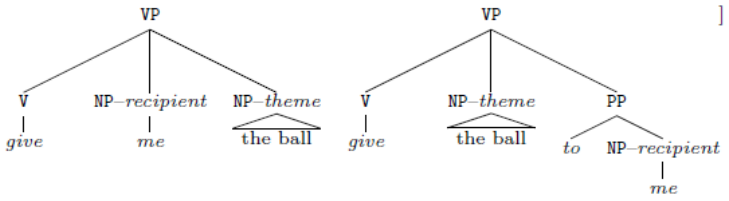

Figure 3. Examples of some of the trees used as input to the Fragment Grammar model.

The Fragment Grammar model determines which sets of subtrees should be stored to best explain the input data. Figure 4 shows three of the many possible sets of subcomputations that the Fragment Grammar model might store for a

small input corpus. Each of these sets of subcomputations should be understood as a *hypothesis* or a *prediction* about future reuse and novelty in the system. The first row of Figure 4 shows a solution which preserves the maximal productivity in the starting computation system. Under this hypothesis all structures are built in a completely combinatorial way from minimal-sized components. The second row of Figure 4 shows a hypothesis that combines the double object construction with the indirect object *me*, thus predicting that this combination is likely to be seen again in the future. Note, that this solution leaves both the verb and direct object positions to be computed on-the-fly. Thus true novelty and variability is predicted in these positions.

Figure 4. Examples of some possible predictions made by Fragment Grammar.

The third row of Figure 4 shows a different prediction. Under this hypothesis a verb specific structure containing both the verb *give* and the indirect object *me* has been stored; this stored structure might correspond to the frequent collocation, *gimme*. Note the tradeoff with the preceding line. The stored frame in the third line cannot be used with the verb *show*, as could the stored double object construction in the preceding line. Thus by increasing belief that future uses of the double object construction would involve **both** *give* and *me*, belief that future uses of the double object construction involving *show* have been reduced. This illustrates the complex set of tradeoffs optimized by the model. By storing structures that hypothesize future repetition of some specific form, the model must reduce its belief in future novelty and variation in the corresponding structure.

Mathematically, Fragment Grammar is a generalization of Adaptor Grammar presented by Johnson and colleagues (2007). Fragment Grammar is also closely related to the Data Oriented Parsing (DOP) formalism (Bod et al., 2003). Fragment Grammar differs from classical versions of DOP in that it tries to find the *best* subset of trees, whereas classical DOP approaches attempted to store *all* of the subtrees that were consistent with the data. More recently, the

optimality perspective has found its way into DOP and other approaches to learning grammars of tree fragments (tree-substitution grammars) for syntactic parsing (Cohn et al., 2009; Post & Gildea, 2009; Zuidema, 2007). The models presented in Cohn et al. (2009); Post & Gildea (2009) are especially similar to the current framework. See O'Donnell (2011) for an in–depth discussion and comparison of these various alternate models.

4.2 Procedure

All dative utterances from the maternal speech in the Providence Corpus (Demuth, et al., 2006) were converted to trees labeled with thematic roles. Verbs were converted to lemmas. (See Figure 3 for examples.) These utterances were run through the Fragment Grammar model described above. The output of the model is the grammar consisting of the various fragments of structure which best approximates the distribution of forms in the input data. Because this grammar is a generative system itself, a large number of different quantities can be computed based on it—e.g., the probability that given a novel verb the double object of prepositional dative construction will be preferred.

Because our interest is in the patterns of productivity in argument structure, we considered only sentence-level fragments. Other fragments and lexical items were not examined. After training, we gave the model utterances containing only novel words. We then determined the conditional probability that the model assigned to either the double object or prepositional dative construction given these novel words.

4.3 Results

On the basis of this input data, the Fragment Grammar stored 163 tree fragments at the sentence level. Storage of a fragment indicates that there was evidence in the input that this structure should be treated as a unit (e.g., frequent use). The table below lists a few of the top-ranked (i.e., most probable) fragments, as well as the rankings of two fragments that are of particular interest to the present question: the fully abstract, non-item-based prepositional and double object forms.

Fragment	Rank
Give you THEME	1
Give me THEME	3
VERB THEME to RECIPIENT	59
VERB RECIPIENT THEME	157

Notice that the most highly ranked fragment is "give you THEME," with "give me THEME" ranked third. This shows that the input data strongly support routinized or collocated "give me" and "give you" forms. Indeed, these forms are ranked well above any more abstract structures. The fully abstract

prepositional dative falls just out of the top third, while the fully abstract double object form is very low ranked. Intuitively, the model has learned that the double object construction is more likely to appear in various routinized or collocated fragments than it is in its more abstract form. The prepositional dative, by contrast, shows the opposite pattern. This means, among other things, that, when evaluated on novel verbs or novel arguments, the model will prefer the prepositional dative. In fact, the model prefers the prepositional dative at a rate of 7.5 to 1 for such forms. This suggests that the patterns of usage in the input data given do not support robust storage of a double object construction, while they do support storage of a prepositional construction. It should be emphasized that this is the case even though the underlying computational system provides a fully abstract representation of both types of datives from the beginning.

5. Discussion

The acquisition of the English dative alternation has been used to support both theoretical accounts that posit early generativity and those that posit an extended period of item-driven grammar. This is in part because the empirical

data on the dative alternation at age 3 show that children have some abstract representation of the alternation (e.g., Thothathiri & Snedeker, 2008; Conwell & Demuth, 2007) but also that their representations for the double object construction do not seem to be as robustly productive as might be expected (Conwell & Demuth, 2007; Rowland & Noble, in press). Specifically, although children hear many more double object datives than prepositional datives and produce them earlier in spontaneous speech, they have difficulty producing and comprehending the double object dative with novel verbs. At the same time, they show priming effects for both structures, suggesting that the abstractions are available in their linguistic system.

Our corpus analysis and our computational model suggest that this difference in productivity is caused by differences in the argument variability between the two constructions. The double object form is much more likely than the prepositional to have a pronoun immediately following the verb, creating a region of low variability and therefore low expectation of future novelty in that region. Regions of low variability suggest to the learner that a particular sequence is not productive but rather that it should be stored and used as a unit. If a construction was not used in novel ways in the past, why should it be used that way in the future? On the other hand, if a particular noun occurred unexpectedly often in a construction, then the learner can bet on it occurring again.

The present model instantiates a theory in which early abstraction is hemmed in by a system which stores commonly occurring computations. Such a system can account for the previous observations on the comprehension and production of datives in young children. First, the early emergence of the double object form in children's early speech may be largely based on highly frequent,

possibly collocated forms such as "give me." In contrast, the verb-general priming effects for both dative forms can be explained by the availability of abstract representations of these structures in the child's grammar.

The model never *eliminates* the possibility of a potential abstraction; it just comes to believe that it is unlikely to be used productively and should be "frozen" into other structures. These studies use prime and target sentences that do not share arguments or verbs, thus the only common structure that they share is the abstract one. Nevertheless, this model is able to explain the differential generalizability of the two constructions to novel verbs: specifically that young children generalize more readily to the prepositional dative, and less readily to the double object dative (in both comprehension and production). In our model this occurs because the abstract double object form is less frequently used, and thus weaker than, the abstract prepositional form, reflecting the learners inference, from her linguistic experience, that the double object dative is rarely productive. Thus such a model provides a way of reconciling a set of observations that cannot be accounted for by usage based theories or explained in its entirety by the early generativity hypothesis.

More generally, this work provides an explicit means of describing the role of the input in accounts of language acquisition that posit very early generativity. Usage-based accounts have emphasized the gradualness of abstraction and how the process is driven by the data itself. The Fragment Grammar model adopts a very different perspective on abstraction and linguistic experience, treating productivity and reuse as targets of inference in their own right, rather than as byproducts of the architecture of the system. The simulation results show that when productivity and reuse are properly understood as learning problems in their own right, the data itself constrains generalization of the double object construction. This is despite the fact that abstract, adult-like structures were available a priori (see O'Donnell, 2011, for examples from other domains). Of course, a more realistic learning model would have to acquire much of the information which we have provided in the start state. However, this would not change the fact that the pattern in the input data supports generalization of the prepositional dative but not the double object dative.

The data and simulations presented here show that even a very abstract starting state will produce item-specific structures when productivity is treated as a learning problem in its own right and such structures are warranted by the input data. Therefore, evidence that children have or use item-based forms cannot be taken as evidence that they lack grammatical abstractions. On the other hand productive use of such abstractions must be licensed by the data.

Of course, any complete, computational theory of language acquisition must ultimately specify the space of generalizations possible for the learner. Early generativity theories have emphasized the importance of providing a space that is of the approximately appropriate grain-size from early in learning. Moreover, they have often provided explicit and detailed accounts of the representations involved (e.g., Pinker, 1984).

We believe that the account provided here takes the best of both worlds. In one sense, by treating the problem of productivity as an inference in its own right whose solution is directly projected from the data, the approach is fundamentally usage based. On the other hand, the approach is also radically generative in both the traditional and modern linguistic senses. It provides a precise specification of the space of possible generalizations right from the beginning. The problem is to decide which *potential* generalizations are *productive* ones.

References

Akhtar, Nameera & Tomasello, Michael. (1997). Young children's productivity with word order and verb morphology. *Developmental Psychology, 33*, 952-965.

Anderson, John R. (1990). *The adaptive character of thought.* Hillsdale, NJ: Lawrence Erlbaum Associates.

Baker, C.L. (1979). Syntactic theory and the projection problem. *Linguistic Inquiry, 10*, 533-581.

Bod, R., Scha, R., & Sima'an, K. (Eds.). (2003). *Data-oriented parsing*. Stanford, CA: CSLI.

Braine, Martin D. S. (1963). The ontogeny of English phrase structure: The first phase. *Language, 39*, 1-13.

Campbell, Aimee. & Tomasello, Michael. (2001). The acquisition of English dative constructions. *Applied Psycholinguistics, 22*, 253-267.

Chomsky, Noam. (1981). *Lectures on Government and Binding.* Dordrecht: Forris.

Cohn, Trevor, Goldwater, Sharon, & Blunson, Phil (2009). Inducing compact but accurate tree-substitution grammars. In *Proceedings of the North American conference on computational linguistics*.

Conwell, Erin & Demuth, Katherine. (2007). Early syntactic productivity: Evidence from dative shift. *Cognition, 103*, 163–179.

Demuth, Katherine, Culbertson, Jennifer & Alter, Jennifer. (2006). Word-minimality, epenthesis, and coda licensing in the acquisition of English. *Language & Speech, 49*, 137-174.

Fisher, Cynthia. (2002). The role of abstract syntactic knowledge in language acquisition: a reply to Tomasello (2000). *Cognition, 82*, 259-278.

Goldberg, Adele E. (1995). *Constructions: A construction grammar approach to argument structure*. Chicago, IL: University of Chicago Press.

Goldberg, Adele E. (2003). *Constructions: A new theoretical approach to language. Trends in Cognitive Sciences, 7*, 219-224.

Goodman, Noah D., Mansinghka, Vikash K., Roy, Daniel, Bonawitz, Keith, & Tenenbaum, Joshua B. (2008). Church: a language for generative models. In *Uncertainty in artificial intelligence*. Helsinki, Finland: AUAI Press.

Jackendoff, Ray. (2002). *Foundations of Language: Brain, meaning, grammar, evolution.* Oxford: Oxford University Press.

Johnson, Mark, Griffiths, Thomas L., & Goldwater, Sharon (2007). Adaptor Grammars: A frame-work for specifying compositional nonparametric Bayesian models. In *Advances in neural information processing systems 19.* Cambridge, MA: MIT Press.

Levin, Beth. (1993). *English Verb Classes and Alternations: A preliminary investigation.* Chicago, IL: University of Chicago Press.

Lieven, Elena V. M., Behrens, Heike, Speares, Jennifer & Tomasello, Michael. (2003). Early syntactic productivity: A usage-based approach. *Journal of Child Language, 30*, 333-370.

Lieven, Elena., Pine, Julian. & Baldwin, Gillian. (1997). Lexically based learning in early grammatical development. *Journal of Child Language, 24*, 187-219.

Naigles, Letitia. R. (2002). Form is easy, meaning is hard: Resolving a paradox in early child language. *Cognition, 86*, 157-199.

Naigles, Letitia. R. & Hoff-Ginsberg, Erika. (1998). Why are some verbs learned before other verbs? Effects of input frequency and structure on children's early verb use. *Journal of Child Language, 25*, 95-120.

O'Donnell, Timothy J. (2011). *Productivity and reuse in language*. Unpublished doctoral dissertation, Harvard University.

O'Donnell, Timothy J., Goodman, Noah D., & Tenenbaum, Joshua B. (2009). Fragment grammars: Exploring computation and reuse in language (Tech. Rep. No. MITCSAIL- TR-2009-013). Cambridge, Ma: MIT Computer Science and Artificial Intelligence Laboratory Technical Report Series.

Pickering, Martin J. & Branigan, Holly P. (1998). The representation of verbs: Evidence from syntactic priming in language production. *Journal of Memory and Language, 39*, 633-651.

Pinker, Steven. (1984). *Language learnability and language development*. Cambridge, MA: Harvard University Press.

Pinker, Steven. (1989) *Learnability and Cognition: The acquisition of argument structure*. Cambridge, MA: MIT Press.

Post, Matthew, & Gildea, Dan (2009). Bayesian learning of a tree substitution grammar. In *Proceedings of the joint conference of the 47th annual meeting of the ACL and the 4th international joint conference on natural language processing of the AFNLP*.

Rowland, Caroline F. & Noble, Claire H. (*in press*). The role of syntactic structure in children's sentence comprehension: Evidence from the dative. *Language Learning and Development.*

Snyder, William. (2007). *Child Language: The Parametric Approach*. Oxford University Press.

Snyder, William & Stromswold, Karin. (1997). The structure and acquisition of the English dative constructions. *Linguistic Inquiry, 28*, 281-317.

Thothathiri, Malathi & Snedeker, Jesse. (2008). Syntactic priming during language comprehension in three- and four-year old children. *Journal of Memory and Language, 58*, 188-213.

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

Viau, Joshua. (2007). *Possession and spatial motion in the acquisition of ditransitives*. Unpublished doctoral dissertation, Northwestern University.

Zuidema, Willem (2007). Parsimonious data-oriented parsing. In *Proceedings of the 2007 joint conference on empirical methods in natural language processing and computational natural language learning (EMNLP-CONLL 2007).*