

A Computational Model of Early Argument Structure Acquisition

Afra Alishahi, Suzanne Stevenson

Department of Computer Science, University of Toronto

Received 27 July 2006; received in revised form 5 November 2007; accepted 5 November 2007

Abstract

How children go about learning the general regularities that govern language, as well as keeping track of the exceptions to them, remains one of the challenging open questions in the cognitive science of language. Computational modeling is an important methodology in research aimed at addressing this issue. We must determine appropriate learning mechanisms that can grasp generalizations from examples of specific usages, and that exhibit patterns of behavior over the course of learning similar to those in children. Early learning of verb argument structure is an area of language acquisition that provides an interesting testbed for such approaches due to the complexity of verb usages. A range of linguistic factors interact in determining the felicitous use of a verb in various *constructions*—associations between syntactic forms and properties of meaning that form the basis for a number of linguistic and psycholinguistic theories of language. This article presents a computational model for the representation, acquisition, and use of verbs and constructions. The Bayesian framework is founded on a novel view of constructions as a probabilistic association between syntactic and semantic features. The computational experiments reported here demonstrate the feasibility of learning general constructions, and their exceptions, from individual usages of verbs. The behavior of the model over the timecourse of acquisition mimics, in relevant aspects, the stages of learning exhibited by children. Therefore, this proposal sheds light on the possible mechanisms at work in forming linguistic generalizations and maintaining knowledge of exceptions.

Keywords: First language acquisition; Argument structure constructions; Computational modeling; Bayesian modeling

1. Introduction

In order to achieve mastery of their native language, children must acquire knowledge of the regularities that govern the usage of various linguistic forms, as well as the exceptions to

them. Such patterns are ubiquitous across all areas of language. In morphology, for example, there are word-formation rules such as the regular past tense in English (*move/moved*), which has various exceptions (*go/went, break/broke*). In syntax, there are word order rules such as the head-first noun-adjective order in French (*fille intelligente*), also with numerous lexical exceptions (*belle fille*). It continues to be a subject of much debate how children manage to discover and use the general patterns, while keeping track of the exceptions—especially in the absence of explicit instruction and in the face of incomplete and noisy input.

One aspect of language that provides a rich testbed for exploring such issues is the usage of verbs—specifically, what syntactic constructions a verb can occur in. Verbs pose particular challenges to children in this regard, due to the complexity of their possible usages and the interacting semantic and syntactic factors that determine both the general patterns and the exceptions. In particular, verb argument structure is a complex aspect of language for a child to master, as it requires learning the relations of arguments to a verb and how those arguments are mapped into valid expressions of the language. In this article, we describe a computational model that mimics early acquisition of verb argument structure in a child, and sheds light on the possible mechanisms at work in forming generalizations and maintaining knowledge of exceptions.

1.1. Child learning of verb argument structure

Appropriate usage of verbs is guided by the knowledge of verb argument structure. Verb argument structure details the semantic relations of a verb to its arguments—that is, the participants in the event the verb describes—and how those arguments are mapped into valid syntactic expressions of the language. This complex aspect of language exhibits both general patterns across semantically similar verbs, as well as more idiosyncratic mappings of verbal arguments to syntactic forms. For example, as shown in Examples 1 through 3, many verbs of movement allow both an intransitive form (as in the (a) examples) and a related transitive form in which the agent causing the movement is expressed as the subject (as in Examples 1b and 2b). However, for the verb *fall*, the transitive form is not acceptable, as shown in Example 3b:

- 1a. The chain dropped into the box.
- 1b. The girl dropped the chain.
- 2a. The ball rolled into the room.
- 2b. The child rolled the ball.
- 3a. The book fell into the fire.
- 3b. *The boy fell the book.

Thus, children must learn both the general argument structure patterns exemplified by the Examples 1 and 2, and the finer-grained semantic restrictions that govern the exceptions, as in Example 3. Given the number of verbs in the language, and the wide range of syntactic structures that can be used to express verbal arguments, the acquisition of such knowledge is quite complex, and numerous questions remain concerning how children attain an adult level of competence.

It appears that children are sensitive to such regularities from an early age, producing novel utterances that obey the mapping of arguments to syntactic positions in their language (Bowerman, 1982; MacWhinney, 1995; Pinker, 1989). Moreover, children may even over-generalize the observed patterns, thereby producing incorrect forms such as *Don't you fall me down* (Bowerman, 1982). Eventually, such erroneous verb usages discontinue. In this way, children exhibit a “U-shaped learning curve” seen in other areas of language learning (Marcus et al., 1992). A U-shaped learning curve indicates an early phase of correct imitation of observed forms, such as *Ball roll*, *Adam roll ball*, and *Ball fall* (the left tip of the U). This is followed by a period of generalization of observed patterns in which children may incorrectly apply the regularities they have perceived to inappropriate forms, such as *Adam fall ball*, lowering their proportion of correct forms (the bottom of the U). Finally, children recover from such over-generalizations, leading once more to a higher rate of correct forms (the right tip of the U). As is generally the case in language learning (Marcus, 1993), explicit negative evidence in the form of corrective input does not play a role in children’s cessation of incorrect forms.

1.2. Theories of argument structure acquisition

The observed pattern of argument structure acquisition leads to two questions that must be answered in a theory of early verb learning:

- How does general knowledge emerge from knowledge of specific verb usages, and what form does this knowledge of argument structure regularities take?
- How do general and specific knowledge interact over the course of acquisition, leading to the observed developmental stages (i.e., the U-shaped learning curve)?

In addition to theoretical proposals, explicit computational models that instantiate a coherent set of answers to these questions must be explored, in order to demonstrate the cognitive feasibility of an account of the data.

With regard to the first question, learning in this domain has been suggested to rely on rich innate knowledge of argument structure regularities—in particular, mappings of semantic arguments to syntactic positions (e.g., Grimshaw, 1990; Pinker, 1989). However, a number of usage-based proposals have argued that children learn such regularities from the input alone, without guidance in the form of innate principles (e.g., Akhtar, 1999; Bowerman, 1982; Tomasello, 2000). Moreover, recent psycholinguistic evidence supports the view that children initially learn verb-argument patterns on an item-by-item (verb-by-verb) basis, before forming a conceptualization of more general syntactic structures (e.g., MacWhinney, 1982, 1987; Tomasello, 2000, 2003). Given this evidence, the proposed innate principles that draw on general structural knowledge are less plausible.

With regard to the second question, of particular interest has been how general and specific knowledge interact to yield a U-shaped learning curve. The mechanisms that children use to recover from making over-generalization errors has especially been the subject of much debate (Bowerman, 1988, 1996; MacWhinney, 2004; Pinker, 1984, 1989). This problem is of special interest because negative evidence (corrective feedback) plays little to no role in

the process of recovery; only additional positive evidence of correct usages is necessary for “unlearning” of over-generalized rules (Marcus, 1993). However, novel forms for a verb cannot be ruled out simply on the grounds that the form has not been observed before, because adult-level competence includes an ability to comprehend the meaning of unusual utterances, or to generate novel utterances for unusual situations. Accounting both for this aspect of language, known as *productive generalization*, and for recovery from over-generalization, necessitates a detailed usage-based theory of the mechanisms at play in argument structure acquisition. A family of usage-based theories that has grown in prominence in work on the acquisition of verb argument structure is *construction grammar*. Construction grammars posit that, in addition to the idiosyncratic meanings associated with individual words or morphemes, meaning may also be directly associated with syntactic forms (e.g., Fillmore, Kay, & O’Connor, 1988; Lakoff, 1987; Langacker, 1999). In particular, an argument structure construction is a mapping between underlying verb-argument relations and the syntax used to express them, as in Examples 4 and 5 from Goldberg (1995):

4. Subj V Obj Obj₂ ⇔ X CAUSE Y RECEIVE Z (e.g., Pat faxed Bill the letter).
5. Subj V Oblique ⇔ X MOVE Y (e.g., The fly buzzed into the room).

Usage-based theories, in general, and those founded on argument structure constructions, in particular, have recently received much attention in the language acquisition community. These proposals raise the question of whether a psychologically plausible usage-based learning strategy can be successfully exploited in a computational model. In particular, explicit models must be explored, both of the underlying learning mechanisms and of the use of the acquired knowledge. Computational experiments must show the feasibility of learning general patterns of the language from individual input without built-in knowledge of argument structure regularities and without the use of negative evidence.

1.3. A Bayesian model for early argument structure acquisition

We present an implemented computational model of the early stages of argument structure acquisition in children. We take a statistical usage-based approach in which the model learns word order regularities in the expression of arguments solely given examples of utterances paired with appropriate semantic descriptions. Our system is shown to mimic children’s behavior in forming word order generalizations and in recovering from over-generalizations without the use of negative evidence. Generalization in our model is achieved through an unsupervised Bayesian algorithm that groups similar argument structure usages. The same Bayesian process underlies language use—comprehension and production—in our model, giving a unified mechanism for language acquisition and processing.

In response to the first theoretical question, concerning the form of general knowledge and how it emerges, we propose a variation on the construction grammar approach. Our model induces probabilistic associations between syntactic and semantic features of a verb; these associations generalize over the syntactic patterns and the fine-grained semantics of both the verb and its arguments. The resulting associations have a similar function to the constructions of Goldberg (1995), and indeed we refer to the associations formed by our

model as constructions. However, our representation enables a new view on the nature of constructions: Each construction is not simply a form-meaning pair, but rather a grouping of verb usages that share similar syntactic and semantic features, inducing a probability distribution over the values for these features. The probabilistic nature of argument structure constructions in the model enables it to capture both statistical effects in language learning and adaptability in language use.

In addressing the second question, a key property of our computational model is the interaction between the specific knowledge of observed verb usages and the constructions that generalize over such observations. This interaction between verb-based and construction-based knowledge plays a crucial role in language use in the model, which is formulated as a Bayesian prediction problem. For example, in production, the model predicts the syntax to express an intended semantics, whereas in comprehension it predicts (some of) the semantics of a perceived utterance. The emergent constructions in the model enable it to generalize observed patterns of association to new or low-frequency situations, whereas knowledge of specific verb usages is used when it is more established. This interaction leads to behavior reminiscent of the U-shaped learning curve observed in children.

A preliminary version of this model was originally presented in Alishahi and Stevenson (2005a, 2005b). In this article, we describe an improved version of the model, present additional experiments, and discuss the results and their import for modeling of child language acquisition. We explain the model and its modules in more detail in the immediately following sections. Successive sections present and discuss computational experiments that demonstrate the match between the behavior of the model and that of children in learning argument structure regularities and in using that knowledge in language comprehension and production.

2. The representation of probabilistic constructions

The input to our model consists of natural language utterances, each paired with a logical representation of the meaning that the utterance conveys. In section 2.1, we describe our assumptions regarding the input. The details of the input pairs are explained in section 2.2. The model extracts an argument structure frame from each input pair, as described in section 2.3. Whereas the input pairs are a simulation of the data that children use in language learning, the argument structure frames are the internal representation of the model for that data, which might differ for different tasks. Therefore, all through the experiments in this article, we use the scene-utterance pairs as the original form of the input to our model. In section 2.4, we describe the probabilistic representation of a construction in our model, which is a group of such argument structure frames that share similar features.

2.1. Basic assumptions

Following Jackendoff (1990), we assume that, upon observing a simple event, the child can construct a conceptual structure to represent that event. Moreover, if the child hears a sentence while watching the scene, he or she can establish a link between the linguistic description of

the event and its conceptual representation. We use such pairings of utterance and meaning representation as the basic input to our model.

One well-known problem in language acquisition is that of *referential indeterminacy* (also termed *referential uncertainty*) in which the child may perceive many aspects of the scene that are unrelated to the perceived utterance (e.g., Gleitman, 1990; Quine, 1960). For example, when the speaker utters *Mom put toys in boxes*, the child may also form a mental representation of the description of various toys in the scene, of the events of Mom picking up and moving toys, and even of unrelated entities and events such as a kitten playing with a string. A number of computational models have dealt with the problem of picking the right meaning for an utterance, as well as learning the meaning of individual words (e.g., as in Fleischman & Roy, 2005; Siskind, 1996; Yu, Ballard, & Aslin, 2005). In our work, we assume that the (non-trivial) task of picking out the appropriate semantics for the utterance from the full representation of the perceived aspects of the scene has been performed, and that for each word the child knows the corresponding meaning (i.e., semantic symbol). Learning the correct meaning for both verbs and nouns has been suggested to be based on cross-situational observation (Fisher, Hall, Rakowitz, & Gleitman, 1994; Pinker, 1989). It has been argued that, for certain verbs or situations, using cross-situational learning is not enough.¹ However, for the purpose of learning the simple verbs discussed in this article (see section 5.2), we make the simplifying assumption that the child can learn the core (idiosyncratic) meaning of simple verbs before learning their full argument structure, although the detailed meaning (which we call the *semantic primitives* of the verb) are learned later.

We also assume that at the point of learning modeled by our system, the child is able to recognize linguistically relevant semantic properties such as the roles played by the arguments in the event. (For example, in an event described by *Mat gave Pat the hammer*, the child can detect Mat as the agent, Pat as the recipient, and the hammer as the theme of the “giving” event.) However, we recognize that resolving referential indeterminacy, learning word-to-meaning mappings, and determining potential argument roles are processes that are likely interleaved with argument structure acquisition. Future work will need to integrate these steps into the model, rather than assuming that the input to the model is the output of a prior processing step that extracts these aspects of the scene. In fact, we address the problem of learning semantic roles in Alishahi and Stevenson (2007b), where we present a modified version of the model that acquires a general conception of the semantic roles of predicates based only on exposure to individual verb usages.

Another property of our model is that, in the long run, it does not remember the verb usages that it has seen in the input. The model extracts certain features from each input pair in the form of an argument structure frame (see section 2.3 for details), and then the input pair is discarded. Among the extracted features are the semantic categories of the arguments that participate in a scene, but the actual words that represent those arguments are not kept. In a way, our model is “memoryless” (i.e., immediately after extracting a frame from an input pair, it forgets the words that appear in each argument slot). Although this enables the model to make appropriate generalizations, it limits the ability of the model to learn more fine-grained semantic properties of the arguments for each verb or construction, as well as collocations (or word co-occurrences).

Scene-Utterance Input Pair		
Put _[cause, move] (MOM _(agent) , TOYS _(theme) , In _[] (BOXES _(destination)) _(destination)) <i>Mom put toys in boxes.</i>		
Extracted Argument Structure Frame		
head verb	<i>put</i>	
verb semantic primitives	<i><cause, move></i>	
number of arguments	3	
arguments	roles	<i><agent, theme, destination></i>
	categories ^a	<i><human, concrete, dest-pred></i>
syntactic pattern	arg1 verb arg2 arg3	

Fig. 1. An input pair and extracted frame for a verb argument structure. Note: ^aExtracted from a representation of the child's ontology.

2.2. Scene-utterance input pairs

An individual verb usage consists of a scene-utterance pair, which we also refer to as an input Pair, because a sequence of such pairs comprises the input to the learner. The scene component of an input pair is a logical form representation of the aspect of the observed scene that the utterance (the perceived linguistic datum) describes. For example, the top panel of Fig. 1 shows a sample input pair. In the scene representation, each predicate (shown in bold) is indexed by a (possibly empty) set of semantic primitives known to describe the action, state, or relation expressed by the predicate. Each entity (shown in all upper case) is indexed by the participant role that the entity plays with respect to the predicate. The corresponding utterance indicates the sequence of words used to express the semantics of the scene. For now, our utterance representation omits such aspects as tense marking on the verb and grammatical function words such as determiners. (Note, however, that prepositions are included as predicates because they play an important role in argument structure.)

2.3. Argument structure frames

From each scene-utterance pair, our system extracts an argument structure frame. The bottom panel of Fig. 1 shows the frame corresponding to the main predicate, **Put**, of the input pair in the top panel of the figure. Each frame records the head word (verb or preposition), the number of arguments, the semantic primitives that describe the predicate, the participant roles of the arguments with respect to the predicate, and the ontological category of each of the arguments. The frame also records the syntactic pattern of a predicate usage; currently, this syntactic pattern encodes the relative positions of the predicate word and each of its arguments. Because prepositions are predicates, an argument structure frame is also extracted for each use of a preposition in the utterance. An extracted argument structure frame that has been previously unobserved for its head word (verb or preposition) is stored in the lexical entry of the head word, and its frequency count is initialized to one. If the frame has been

previously observed for the head word, then its already stored frequency count is incremented. Highly similar frames may be merged into a single compatible frame in the lexical entry. For example, if two frames are identical except that the semantic categories of the arguments of one frame are more general than those of the other, the more specific frame is replaced with the more general one. An extracted frame that has been unobserved previously is input to our Bayesian learner that forms groupings of frames into constructions.

2.4. Constructions as groups of frames

A construction in our model is a group of argument structure frames that are “similar enough,” according to the probabilities over their features, to be grouped together. The notion of “similar enough” is described in detail in the next section of the article; here we focus on the general properties of constructions and their representation in the lexicon. Due to the formulation of our learning model, frames with similar syntactic and semantic properties are usually grouped together in one construction. The values for some features, such as the syntactic pattern and the thematic role assignments, come from a limited set due to the regular nature of natural languages. Therefore, there is often a dominant syntactic pattern and thematic role assignment among frames in each construction. However, due to the probabilistic nature of our model, a few “outliers” might be found in each construction as well. Features with more varied values (such as semantic categories of the arguments and semantic primitives of the verbs) are less uniform. However, it is often the case that frames in a construction share close values for such features—for example, the semantic categories of the first argument in a group of frames that represent a transitive construction might be (*human*), (*animal*), (*living_thing*), and so on. Therefore, the primary property of constructions in our model is that they determine a probabilistic association between syntactic and semantic features. For example, usages such as *Jay got a tower* and *Kay made a tower* may yield frames that form a (transitive) construction. Although the frames share the verb semantic primitive *act*, they differ in others (*possess* for the former, and *become* for the latter). If this observation holds across a number of usages that exhibit this form, then we would find a higher probability for the primitive *act* given this construction than for the other semantic primitives. In this way, constructions probabilistically generalize the properties of a set of frames.

Each verb with a frame that participates in a construction has a link to that construction in the lexicon. The links are weighted by the frequency with which the verb has been seen in the participating frame, capturing the statistical usage pattern of the verb in various constructions. From the point of view of a construction, these frequencies are used in the calculation of the probabilities over the features occurring in its observed frames. Moreover, the sum of its incoming link frequencies contributes to the overall probability of a construction. Figure 2 illustrates a portion of the acquired lexicon, showing the lexical entries of verbs containing frames and their links to constructions. (See Fig. 3 for the legend of symbols used in all depictions of the lexicon throughout the article.) In the next section, we describe how the link between a frame and a construction is established by our Bayesian learning algorithm.




Fig. 2. A portion of the lexicon showing an intransitive construction, with the predominant syntactic pattern Subj(ect) V(erb); and a transitive construction, with the predominant syntactic pattern Subj(ect) V(erb) Obj(ect). *Note:* Numbers on links represent the observed frequency of the verb in a frame that participates in that construction. For example, *eat* has been seen once in an intransitive construction and 11 times in a transitive construction.

3. Acquisition of constructions

3.1. Overview of the acquisition process

The argument structure acquisition mechanism that we propose incrementally learns from the processing of each scene-utterance pair that is input to the model. As noted earlier, the system first extracts an argument structure frame from the current input; this frame is then presented to an unsupervised Bayesian clustering process. This process groups the new frame together with an existing group of frames—a construction—that probabilistically has the most similar properties to the new frame. The Bayesian approach we use is an adaptation of a model of human categorization proposed by Anderson (1991), which incrementally groups perceived items (in our case, frames) into categories of items with similar features (in our case, constructions).² It is important to note that the categories (i.e., constructions) are not predefined, but rather are determined by the similarity patterns over observed frames. A new construction is created in response to a frame when none of the existing constructions are sufficiently similar to it—that is, when no existing construction has sufficiently high probability given the properties of the new frame. An overview of the process is shown in Fig. 4.

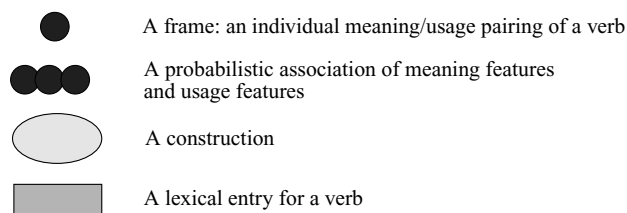


Fig. 3. The legend for symbols used in all the diagrams representing verbs and constructions throughout the article.

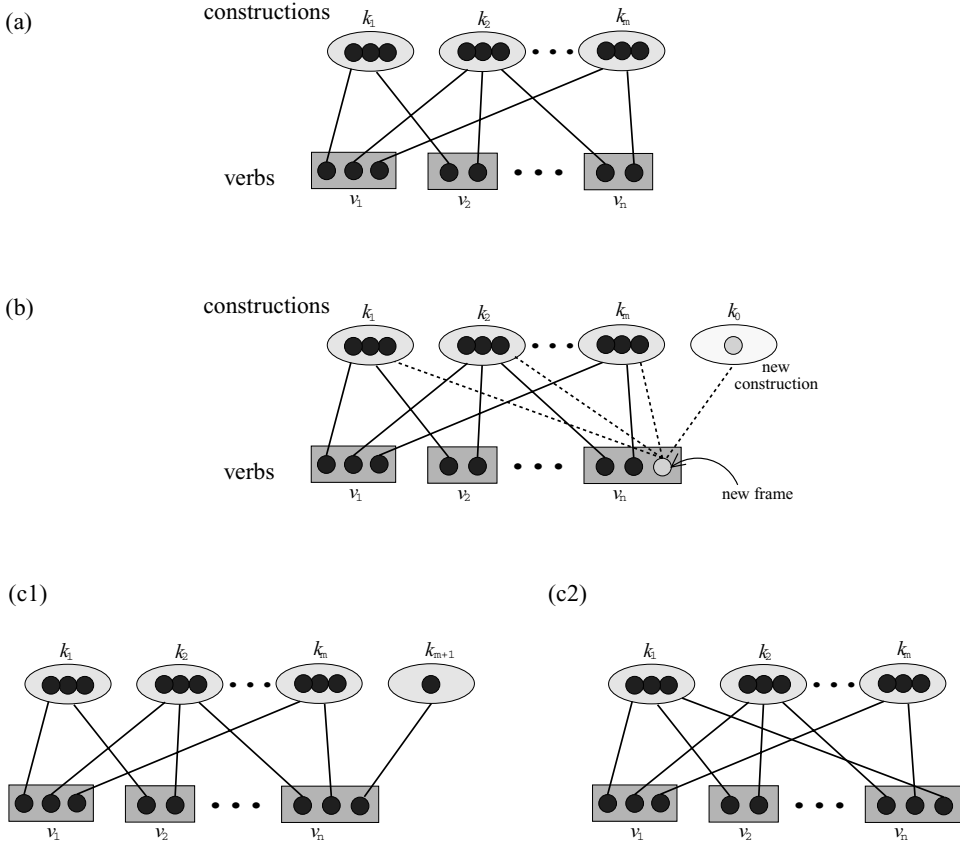


Fig. 4. Processing a new frame: (a) A snapshot of the lexicon before the new frame is processed. (b) Addition of the new frame to the lexical entry for its main verb, in this case v_n . All existing constructions and a potential new construction are considered as a match for the new frame, as indicated by the dashed lines. Two alternatives are possible, as shown in panels c1 and c2. (c1) The new construction k_0 is the best match; hence, it is added to the lexicon, and the new frame is linked to it; or (c2) one of the existing constructions (in this case k_1) is the best match; hence, the new frame is linked to it, and no new construction is created.

Grouping a new frame F with other frames participating in construction k is formalized as finding the k with the maximum probability given F :

$$\text{Best Construction}(F) = \underset{k}{\operatorname{argmax}} P(k|F) \tag{1}$$

where k ranges over the indexes of all constructions, including an index of 0 to represent a new construction. The following uses Bayes's rule, dropping $P(F)$, which is constant for all k :

$$P(k|F) = \frac{P(k)P(F|k)}{P(F)} \propto P(k)P(F|k) \tag{2}$$

The prior probability of a construction, $P(k)$, is proportional to the total frequency of frames associated with k . The conditional probability of the new frame given a construction, $P(F|k)$, is

determined by the proportion of both syntactic and semantic features it shares with the frames in the construction. Because this probability may be calculated over partial information, learning can proceed even in the face of an incomplete frame (i.e., when some of the frame features are missing).

3.2. Details of the probability model for learning

The following sections describe in detail the calculation in Equation 2 of the prior and conditional probabilities, respectively, used in learning.

3.2.1. Prior probabilities

The higher the prior probability of a construction, the more likely for a new usage of a verb to appear in it. In a usage-based framework, the probability of a construction is assumed to change over time, as the child observes additional evidence of various constructions. In particular, as a construction is used more, it becomes more entrenched and is more likely to be used in the future. To model this intuition, we calculate the prior probability of a construction based on the frequency of its frames in the observed input, rather than assigning some pre-set prior distribution across constructions. In this way, the prior probability of each construction varies over the course of learning, reflecting how entrenched that construction has become. The prior probability of construction k , for $k > 0$, is given by the following:

$$P(k) = \frac{n_k}{n + 1} \quad (3)$$

where n is the total number of observed frames, and n_k is the number of frames participating in construction k . The normalizing factor is set to $n + 1$ rather than n , to accommodate the probability mass of a potential new construction ($k = 0$), whose prior probability is given by the following:³

$$P(0) = \frac{1}{n + 1} \quad (4)$$

Thus, the prior probability for an existing construction is proportional to the frequency of its frames (Equation 3), following the intuition that it is more probable for a newly observed frame to come from a more entrenched construction. Moreover, the prior probability of a new construction is inversely proportional to the number of observed frames overall (Equation 4), capturing the intuition that the more exposure the child has to the language, the less likely it is that a brand new construction will be observed.

3.2.2. Conditional probabilities

The conditional probability of a frame F is expressed in terms of the individual probabilities of its features. To make the calculation feasible, we assume that these features are independent; thus, the conditional probability of a frame F is the product of the conditional probabilities of its component features:

$$P(F|k) = \prod_{i \in \text{FrameFeatures}} P_i(j|k) \quad (5)$$

where j is the value of the i th feature of F , and $P_i(j|k)$ is the probability of displaying value j on feature i within construction k . This probability is estimated using a smoothed maximum likelihood formulation, reflecting the emphasis on usage statistics in child language acquisition. The conditional probability estimate for each of the features conforms to the following general pattern (Lidstone's law; Manning & Schütze, 2001):⁴

$$P_i(j|k) = \frac{\mathbf{count}_i^k(j) + \lambda}{n_k + \lambda\alpha_i} \quad (6)$$

where n_k is the number of frames participating in construction k , for $k \geq 0$ (note that $n_0 = 0$), and $\mathbf{count}_i^k(j)$ is the number of those with value j for feature i . The smoothing parameter α_i is set to (an estimate of) the number of possible values that feature i can take on. Note that, for a new construction, because $\mathbf{count}_i^0(j)$ and n_0 both have the value 0, the estimated conditional probability is $1/\alpha_i$ —that is, for a new construction, all feature values are equally likely. Parameter λ in Equation 6 is set to a small constant so that the constructions that have no members with value j for feature i have a low (but non-zero) probability.

Some specific properties are worth noting for the case in which k is a well-entrenched construction—that is, when n_k is large. In this case, the conditional probability of a particular feature value is close to the relative frequency of its observed occurrence in k , $\mathbf{count}_i^k(j)/n_k$ (because λ in the formula is small). If $\mathbf{count}_i^k(j)$ is zero, the conditional probability will also approach zero. Thus, although a large n_k yields a large prior probability for the construction k (see Equation 3), if the feature value of the frame F being processed has not been seen in construction k , the effect of the large prior will be greatly reduced. If this is true for all existing constructions ($k \geq 1$), the model will likely add a new construction to accommodate the previously unseen feature value.

4. Language use as prediction

We formulate language use (production and comprehension) as a prediction process in which missing features in a frame are set to the most probable values given the available features. Thus, we model sentence production as predicting the most probable syntactic pattern, given the semantic features of a scene representation. Comprehension is conversely the prediction of (some aspects of) the intended semantics from a given utterance. For example, to produce a sentence to express the following scene semantics, we ask the model to predict the “missing” syntactic pattern from the corresponding input:

Kick_[cause,contact](TIM_(agent), BALL_(theme)).

This prediction process integrates both sources of acquired knowledge—the frames stored with each verb, along with their organization into constructions. If a usage of a verb is sufficiently complete and frequent, then it will have a strong influence on the prediction of unknown features. On the other hand, if relevant verb-based information is incomplete (or missing), constructions play an important role. Constructions are the source of generalization over observed knowledge: A prediction based on a construction generalizes over all its frames, and allows the model to extrapolate usages associated with some verbs to new usages of others.

This enables the model to produce or understand a verb in a novel (for that verb) frame, as long as semantically similar verb usages have been observed.

The value of an unobserved feature is predicted based on the match between the given partial set of observed features and the learned constructions:

$$\begin{aligned} \mathbf{Best\ Value}_i(F) &= \operatorname{argmax}_j P_i(j|F) \\ &= \operatorname{argmax}_j \sum_k P_i(j|k)P(k|F) \end{aligned} \quad (7)$$

Here again, F is a partial frame, i is a feature missing from that frame, j ranges over the possible values of feature i , and k ranges over all constructions. Intuitively, the model generalizes over the items within a construction to predict the most probable value for the missing feature across those items ($P_i(j|k)$); this is then weighted by the probability of the construction given the partial frame ($P(k|F)$). The first component of the sum, $P_i(j|k)$, is calculated as in the learning phase of the model, using Equation 6. To calculate the factor $P(k|F)$, we use Bayes's rule and drop the constant $P(F)$ term (cf. Equation 2) to yield the following:

$$P(k|F) \propto P(k)P(F|k) \quad (8)$$

Here, the prior probability, $P(k)$, and the conditional probability, $P(F|k)$, are simply determined as in our learning module (using Equations 3 and 5, respectively). Recall that one of the features in a frame is the head verb of the usage. Due to the formulation of the conditional probability, a construction that has no frames with the same head verb as the new frame F will have a very low conditional probability for the head word feature, $P_{verb}(j|k)$, compared to other constructions. This property reflects the influence of the observed usages of the verb, which gives a verb-based bias to our prediction model. On the other hand, the conditional probability of each construction is modulated by its prior probability, $P(k)$, which reflects the entrenchment of that construction. The outcome of the prediction model is mainly determined by these two important factors, which interact over the course of learning to lead to three phases of acquisition with respect to verb argument structure usage. First, we observe a period of imitation. Early on in learning, constructions generally are not highly entrenched, so the constructions that contain specific observed usages of the verb in the frame usually outweigh other constructions (even if they are not the best intuitive match for the new frame). Over time, however, constructions become more entrenched as more examples of them are observed. During this period, the knowledge of general constructions may be stronger than the knowledge of a certain usage of a particular verb. This leads to a second phase of learning, that of generalization. In some cases, this is valid generalization of the argument structure properties of the language, whereas in others this leads to over-generalization—incorrect usage of a verb according to a general pattern of the language that does not apply to that verb. With additional exposure to the language, however, the model recovers from such errors. As the frequency of the verb in its allowable usages increases, the knowledge of these usages specific to the verb again overrides the weight of a construction unseen with that verb. However, it is important

to note that such recovery from over-generalization does not prevent the model (like children) from making productive generalizations, a point we return to in our experimental results.

5. Experimental setup

As input, the model requires a corpus of child-directed natural language sentences, annotated with the meaning that each sentence conveys. This annotation must include the predicate structure of the event described by the scene, the semantic role that each argument receives by the main predicate, the semantic primitives of each predicate in the described scene, and the semantic properties of the arguments. Large scale corpora of utterances paired with such meaning representations do not currently exist, so we artificially generate the needed input based on the properties of existing child data.⁵ Artificially generating the input data enables us to simplify the syntactic and semantic properties of the input, to concentrate on those aspects of verb argument structure acquisition that are the focus of this article. Moreover, using automatically generated input allows us to test the model through a number of simulations, and evaluate it across minor variations in corpora.

We explain the basic properties of the input required by our model in section 5.1. To support our experimental simulations, we create an input-generation lexicon, which is used to randomly generate corpora on which to test our model; this lexicon is described in section 5.2. In section 5.3, we describe how the input-generation lexicon is used to create the input corpora for each simulation in our set of experiments. Finally, in section 5.4, we explain how the parameters of the model are set.

5.1. Basic properties of the input

We focus on learning the argument structure of a small group of verbs (and prepositions) whose semantic primitives are largely detectable by the child from the scene. As illustrated in Fig. 5, we use a simple logical form for representing the relevant semantics of an observed scene. Verbs and prepositions are represented as predicates (e.g., **Put**, **On**) that can take a number of arguments, each of which can be an entity (e.g., MARY, BALL, TABLE) or a predicate structure itself (e.g., **On** (TABLE)). Each argument is assigned a participant role in the event (e.g., *agent*, *theme*, *destination*), and each predicate is augmented by a set of semantic primitives that describes the characteristics of the event (e.g., [*cause*, *move*]). Currently, prepositions are associated with empty sets of primitives.

The representation of an utterance is simplified to include only words in a particular order, so that the model can extract a syntactic pattern that specifies the relative ordering among a predicate and its arguments. In the current implementation, we do not address learning of morphology; all words appear only in their root forms. This simple formalism allows us to

U(tterance): *Mary put ball on table.*
 S(cene): **Put**_[cause,move](MARY_{<agent>}, BALL_{<theme>}, **On**_[](TABLE_{<destination>})_{<destination>})

Fig. 5. A sample input pair.

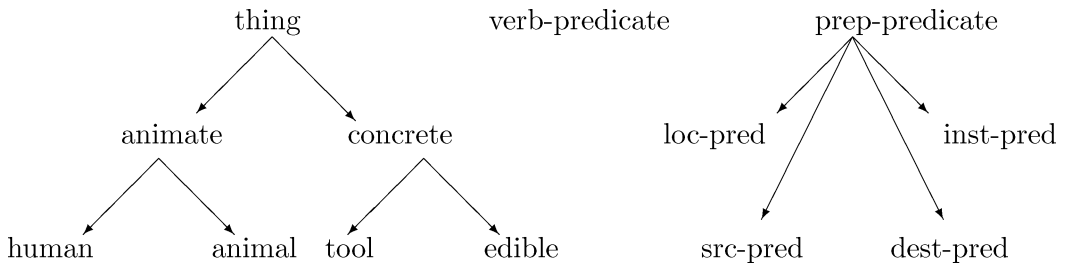


Fig. 6. The semantic hierarchy.

represent nested structures (when the argument of a predicate is another predicate). However, more sophisticated syntactic phenomena such as dependencies cannot be captured by this format.

The participant roles and semantic primitives used in the semantic representation are drawn from a finite, predetermined set that are assumed to be known to the child at the stage of learning being modeled. The participant roles we use are a list of 10 commonly accepted thematic roles from the linguistic literature (e.g., Jackendoff, 1990): *agent*, *theme*, *destination*, *source*, *beneficiary*, *stimulus*, *state*, *experiencer*, *instrument*, and *co-agent*. We enumerate a set of nine primitives for describing the coarse-level semantics of a verb: *act*, *cause*, *move*, *become*, *possess*, *change-of-state*, *perceive*, *contact*, and *manner*, drawing on linguistic proposals concerning fundamental event properties (e.g., Jackendoff, 1990; Rappaport Hovav & Levin, 1998). We augment these with features that capture the finer-grained meaning distinctions among our experimental predicates (e.g., *consume* for an eating event, *playfully* for a playing event, and *rest* for a sitting event). This allows us to compare the learning of general semantic properties of predicates to that of more specific ones.

In extracting the information for an argument structure frame from a scene representation, the model records the semantic category of all the arguments. These semantic categories are selected to reflect a simplified early ontology of a child, depicted in Fig. 6. (As noted earlier in section 2.3, these categories may be used in merging highly similar argument structure frames.) We assume that at the stage of acquisition being modeled, these and other semantic properties can be largely determined by the child from the observed scene.

5.2. The input-generation lexicon

To provide the model with input data whose distributional characteristics are similar to those of the data children receive from their parents, we analyzed a portion of the CHILDES database (MacWhinney, 1995) as follows. We extracted from CHILDES all tenses of the 20 most frequent verbs in mother's speech to each of Adam (2;3–4;10), Eve (1;6–2;3), and Sarah (2;3–5;1). Table 1 shows the list of verbs and their frequencies for each child. We selected 13 verbs from those in common across these three lists and added them to the input-generation lexicon. Each verb has associated with it in the lexicon the total frequency of its root form among the three children and a unique symbol to represent its semantic predicate (e.g., the predicate symbol **Put** for the verb *put*). We also assigned each verb a set of possible argument

Table 1
Top 20 frequent verbs for Adam, Eve, and Sarah from CHILDES database

Adam		Eve		Sarah	
do	699	put	309	do	901
put	534	do	290	get	656
go	332	get	321	go	681
get	280	go	191	put	530
look	247	eat	128	come	404
take	212	take	117	take	296
play	191	sit	90	look	278
come	140	make	89	make	218
write	117	read	86	play	189
give	114	come	75	give	175
make	106	look	as	sit	144
sit	81	play	64	eat	127
eat	70	drink	61	write	78
find	66	give	59	find	67
read	59	write	50	fix	56
hold	55	find	43	drink	55
fall	46	hold	41	fall	50
draw	37	fix	41	hold	47
drink	24	draw	13	dance	37
fit	19	fall	7	read	29

structure frames and associated frequencies, which are manually compiled by examination of all uses of a verb in all conversations of the same three children. Prepositions used in these conversations were also added to the lexicon. The selected words are shown in Table 2, and a sample lexical entry is shown in Fig. 7.

For each semantic category from Fig. 6 that appeared as an argument in the manually compiled argument structure frames, a number of nominal terms were arbitrarily selected and added to the input-generation lexicon. For example, to satisfy the semantic category (*edible*) that appears in the argument structure frames for *eat*, we added the nouns *apple* and *cookie* to the lexicon. Because the focus of the model is on learning the argument structure of predicate terms, we do not expect the frequency of nouns to affect its behavior. Therefore, instead of collecting the relative frequencies of the selected nouns from CHILDES, we assigned equal frequency to all the nouns in the lexicon (arbitrarily set to the value 1).

It is important to note that the input-generation lexicon is not used by the model in language learning and use, but only in producing the experimental corpora, as described next. The only lexical knowledge built into the model is the mapping between a word and a semantic symbol.

Table 2
Predicate terms in the lexicon

Verbs	go, put, get, make, look, take, play, come, eat, fall, sit, see, give
Prepositions	to, in, on, onto, under, from, for, at, with

WORD:	come
SYMBOL:	COME
CATEGORY:	[predicate]
FREQUENCY:	98
FRAME:	COME(\$1)
ARG-NO:	1
SPRIMITIVES:	act,move
ARGS:	\$1 is [agent]
CONDITION:	\$1 is [animate]
PATTERN:	\$1 [word]
CONST-ID:	1
FREQUENCY:	35
FRAME:	COME(\$1,\$2)
ARG-NO:	2
SPRIMITIVES:	act,move
ARGS:	\$1 is [agent]
ARGS:	\$2 is [destination]
CONDITION:	\$1 is [animate]
CONDITION:	\$2 is [dest-pred]
PATTERN:	\$1 [word] \$2
CONST-ID:	5
FREQUENCY:	58
FRAME:	COME(\$1,\$2,\$3)
ARG-NO:	3
SPRIMITIVES:	move
ARGS:	\$1 is [agent]
ARGS:	\$2 is [theme]
ARGS:	\$3 is [destination]
CONDITION:	\$1 is [animate]
CONDITION:	\$2 is [concrete]
CONDITION:	\$3 is [loc-pred]
PATTERN:	\$2 [word] \$3
CONST-ID:	8
FREQUENCY:	5

Fig. 7. A sample lexeme.

5.3. The input corpora

For each simulation in our set of experiments, an input corpus of scene-utterance pairs is automatically created from the input-generation lexicon. The input pairs are randomly generated, using the frequencies in the input-generation lexicon to determine the probabilities of selecting a particular verb and argument structure for each input. Arguments of verbs are also probabilistically selected, constrained to conform to the indicated semantic category of the argument. Arguments that are predicates (such as prepositional phrases) are constructed recursively. The syntactic pattern associated with the selected predicate is used to produce the utterance that accompanies the created scene representation.

We assume that, at the point of learning modeled by our system, the child is able to recognize the syntactic pattern of the utterance and the semantic properties of the event (see section 2.1). However, in reality, the input to children is often noisy or incomplete. For example, the child might mishear the utterance, or might not be able to extract the correct syntactic pattern from it. Similarly for the semantic information, the child might not be able to recognize the correct roles that the participants receive in a particular event, or the semantic categories that these participants belong to, or the semantic properties of the event itself. More problematic for the process of learning, children might sometimes misinterpret the perceived scene or utterance by filling the gaps based on their own (not so perfect) linguistic knowledge, which leads to noisy data.

Because the input to our model is automatically generated, it has none of the imperfections that can be observed in the input data available to children. To demonstrate that the model can function in the face of uncertainty, we need to manipulate the input corpora and introduce some degree of noise. The simulation of noise in our model, however, is to some extent arbitrary: A more accurate approach must be based on careful study of the types of noise that can be observed in child-directed data, and their relative frequency. We attempt to simulate two types of noise in our input corpora: incomplete data (where some pieces of information are missing) and misinterpreted data (where the child replaces the missing data with his or her own inference). Other types of noise, such as ungrammatical or incomplete sentences, are not currently modeled in the input. During the input-generation process, two generated input pairs out of every five have one of their features randomly removed. The removed feature can be one of the following: the syntactic pattern, the semantic category of one argument, the participant role of one argument, and all the semantic primitives of the main predicate. One of these modified input pairs is used to simulate incomplete data. This modified input pair is left as is (i.e., with one feature missing) in the generated corpus. The other modified input pair is used to simulate noise. During a simulation, the missing feature of this input pair is replaced with the most probable value predicted for it at that point in learning; the completed input pair is then used in the learning process. This corresponds to a child using his or her own inferred knowledge to fill in information missing from an observed scene-utterance pair. The resulting input pair is noisy, especially in the initial stages of learning.⁶

To investigate the effect of this partial and noisy data on the performance of the model, we run a suite of experiments both on a set of original generated input corpora (the “original” corpora), and on their corresponding versions modified as previously described (the “noisified” corpora). With this exception (which is noted explicitly in the results that follow), all experiments are over noisified corpora only.

5.4. *Setting the parameters of the model*

For the purpose of implementing the model and running simulations, we need to assign values to parameter α_i for each feature i in Equation 6. In our model, each frame has six different features: the head word, the semantic primitives of the head word, number of arguments, participant roles of the arguments, semantic categories of the arguments, and the syntactic pattern.

The parameter α_i is an estimate of the number of possible values that feature i can take on. We assign $\alpha_{head} = 25$ and $\alpha_{patt} = 15$ as an upper bound on the number of the predicates and the distinct syntactic patterns that can be observed in the input. We assume that each predicate structure can have 0 to 5 arguments, therefore $\alpha_{arg_num} = 6$. The estimated α_i value for other features is determined by the number of semantic primitives, thematic roles, and semantic categories, respectively, as defined in section 5.1: $\alpha_{prim} = 12$, $\alpha_{role} = 10$, and $\alpha_{arg_cat} = 13$.

We also need to set the parameter λ in Equation 6. For all features i in frame F , we need λ to be small enough that the resulting $P_i(j|k)$, where j differs from the value in k , penalizes construction k in the competition in Equation 1. If we wanted each construction k to contain only frames that have the same values on all of the features, we would want the posterior probability of k to be less than that of a new construction, or $P(k|F) < P(0|F)$. The extreme case would be when only one construction exists, which contains all n observed frames, every one of which has the same values as F for all features, except for feature i . By simplifying the inequality $P(k|F) < P(0|F)$ for this case, we obtain $\lambda < \Pi_i \frac{1}{\alpha_i}$. This gives us a lower bound of 10^{-7} for the value of λ . We set λ to a less strict value, 10^{-4} , to allow for constructions that have frames with similar but not identical values for some features.

In the future, we need to investigate in more detail the effect of parameter setting in the output of our model. Whereas the value for each parameter α_i is determined by the number of the possible values that feature i can take, setting λ is more arbitrary. Although we set this parameter by making minimal assumptions, it is worth examining the changes in the behavior of the model due to a different setting for this parameter. We expect a smaller value for λ to result in the formation of more fine-grained constructions (i.e., those that contain frames with more similar feature values for all the features), which might limit the generalization capabilities of the model at the earlier stages of learning. However, such effects might diminish over time as more input is processed by the model.

6. Experimental results

We report computational experiments that demonstrate the ability of our model to learn constructions representing the knowledge of argument structure regularities. We also show that the model is able to generalize this knowledge to novel situations in language use. Language use is formulated as a set of prediction tasks corresponding to comprehension (prediction of semantic properties from a partial input) and production (prediction of a syntactic form—i.e., the word order of a verb and its arguments—from the semantics to be expressed). The simulations we report all use the same parameter settings; only the generated input corpora differ.

In the experimental results reported here, we first show how general constructions emerge as the model is exposed to a range of verb usages over time. We then describe the learning curves displayed by our model for some example verbs; these show interesting variations on the well-known U-shaped learning curve in which use of (correct) imitative forms is followed by a period of (possibly incorrect) generalizations, which is eventually followed by recovery from over-generalization errors. In successive subsections, we examine in more detail some interesting stages in the generally observed U-shaped learning curve: the process of generalization

that follows use of imitative forms; the possible production of over-generalization errors, and recovery from them; and the stable process of productive generalization.

Due to the limited range of verbs in the lexicon, as well as other missing perceptual and linguistic input of importance in language development, it is not possible to directly compare the behavior of the model to that of children from different age groups. However, in the experiments reported later, we change the relative amount of input to which we expose our model to show that, over time, the model goes through different stages of learning similar to those observed for children.

6.1. Construction formation

The formation of argument structure constructions is the primary means of organizing lexico-syntactic knowledge about predicates because such constructions encode the mapping of predicate argument relations to their syntactic expression. Here, we focus on verbs because (in contrast to prepositions, the other type of predicate in our model) they exhibit a range of argument structure variations, with the potential for acquisition of interesting form-meaning mappings. In particular, we are interested in the probabilistic relations that develop in our model between verb semantic primitives (such as *act* and *cause*) and a formal pattern (the argument roles and their expression in word order).

6.1.1. Basic construction learning

We collect evidence for our view of construction learning by tracking the formation of several different constructions in five different simulations of our model. We examine the probabilities of the set of verb semantic primitives associated with transitive, intransitive, and ditransitive constructions, as follows. In each simulation, we train the model on either 50 input pairs (representing the early stage of learning) or 500 input pairs (a more advanced stage of learning). We then present the model with a typical usage of a novel verb in each construction, with the verb semantic primitives omitted. Given this partial input, the model must predict the missing semantic primitives.

Figures 8, 9, and 10 show the probabilities of primitives associated with the transitive, intransitive, and ditransitive constructions, respectively, across the five simulations. Both the transitive usage (as in *Mary made cake*) and the intransitive usage (as in *John came*) occur frequently for a number of verbs in our lexicon. The formation of the two corresponding constructions shows a similar pattern (see Fig. 8 and Fig. 9). After only 50 input pairs, the probabilistic association is unevenly distributed among a small set of primitives, and differs for each simulation. However, after processing 500 input pairs, all the simulations display similar distributions across the full set of primitives, with higher probabilities assigned to the primitives more frequently associated with the construction. For example, a closer look at the intransitive construction in Fig. 9 shows that, after processing 50 input pairs, the probabilistic distribution of semantic primitives differs across the five simulations. Moreover, each simulation shows a bias toward one or two of the verbs in the lexicon that appear in the intransitive construction: *act* and *playfully* in Simulation 1 correspond to the verb *play*; *act* and *move* in Simulations 3, 4, and 5 to the verbs *go* or *come*, with *manner* corresponding to *dance* in Simulation 5. In Simulation 2, we see a more even distribution over *act*, *move*, *playfully*,

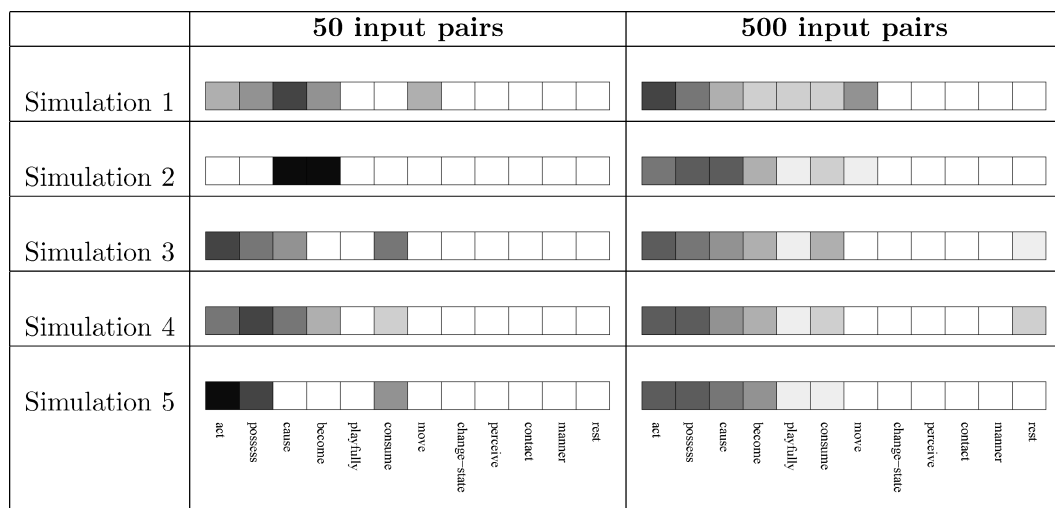


Fig. 8. The probability distribution of semantic primitives in the transitive construction at two different points in learning. *Note:* Darker squares correspond to higher probabilities.

rest, *cause*, and *become* due to receiving more intransitive usages than other simulations in the first 50 input pairs. However, after 500 input pairs, all the simulations display a similar distribution across the primitives *act*, *move*, *playfully*, *manner*, *consume*, and *rest*, with a higher probability for *act* (because it occurs with all intransitive verb usages in our lexicon). Most other primitives have a probability near zero.

An interesting trend can be observed for the ditransitive construction (as in *John gave Mary the book*). In our lexicon, the only verb that appears in this construction is *give*. An

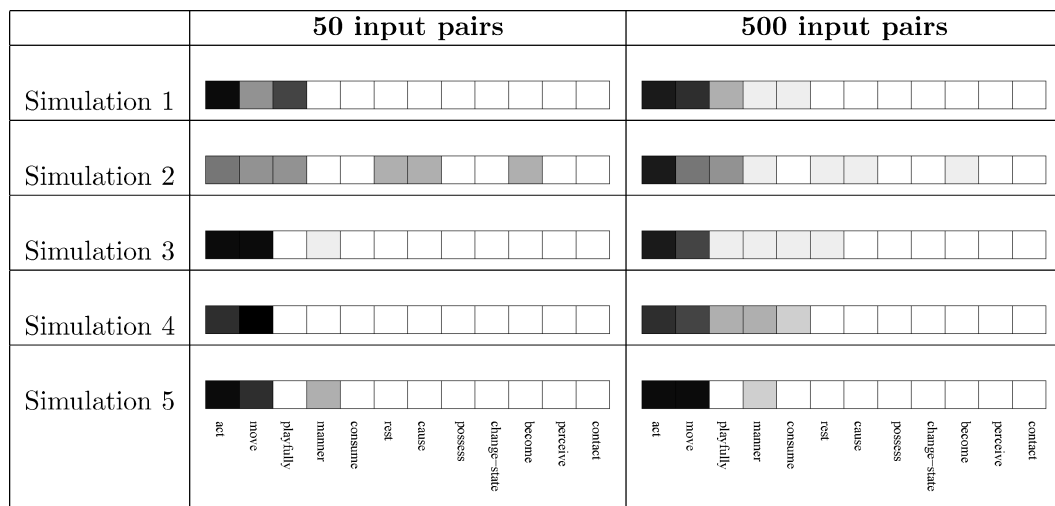


Fig. 9. The probability distribution of semantic primitives in the intransitive construction at two different points in learning. *Note:* Darker squares correspond to higher probabilities.




Fig. 10. The probability distribution of semantic primitives in the ditransitive construction at two different points in learning. *Note:* Darker squares correspond to higher probabilities.

examination of Fig. 10 reveals a very different pattern of development from that of the more common transitive and intransitive constructions. The primitives *cause* and *possess*, which are associated with the verb *give* in the input-generation lexicon, have a very high probability assigned to them from early on, and the probabilistic distribution for the construction does not change much after 500 input pairs. (In Simulations 4 and 5, after receiving 50 input pairs, the ditransitive construction is not formed yet.) Although this particular pattern of learning occurs for the ditransitive because of our limited lexicon, the emergence of highly specific constructions in the language more generally, such as ‘the way’ construction (as in *She typed her way to a promotion*), would be expected to unfold similarly.

6.1.2. Using an extended lexicon

The experiments reported above are compatible with our view of the formation of constructions as a probabilistic association between a form (a syntactic pattern) and elements of meaning (the semantic primitives of predicates). However, the small number of verbs in the lexicon, and the limited set of semantic primitives associated with them (especially the more specific ones), raise the concern that the observed patterns may not be valid for a larger and more varied lexicon.

To address this issue, we artificially expanded our input-generation lexicon by adding new verbs that differ from existing verbs only in their fine-grained semantic primitives. We chose three verbs in the original lexicon (i.e., *sit*, *eat*, and *play*) as the starting point for each of three sets of six semantically and syntactically related verbs. For each of the three selected verbs, we added five verbs to the lexicon with the same argument structure frames and frequencies as the original verb. (Note that the verb *sit* only has intransitive frames, whereas the verbs *eat* and *play* have both intransitive and transitive frames.) The five added verbs also had the same general semantic primitives as the original verb, but all differed in their fine-grained




Fig. 11. Transitive construction for the extended lexicon. *Note:* Darker squares correspond to higher probabilities.

primitives. For example, we added to the lexicon five verbs, sit_1 to sit_5 , each of which had the semantic primitive *act* associated with it, but in place of the fine-grained primitive *rest* that occurs with *sit*, each verb sit_i is associated with the primitive $rest_i$. Thus, each such group represents a class of verbs that have similar meaning and appear in the same argument structure constructions, and can in this way be thought of as a semantic verb class similar to those of Levin (1993).

Figures 11 and 12 show snapshots of the transitive and intransitive constructions, respectively, after 50 and 500 input pairs, for one sample simulation using the extended lexicon. The same effect can be observed for both constructions. After 50 input pairs, the probabilities over the semantic primitives are sporadically distributed, with a few verb-specific primitives having much higher probability than others. However, after 500 input pairs, the probability mass is more evenly distributed among primitives associated with all verbs that appear in that construction, with more general primitives (such as *act*) having a higher probability. Most

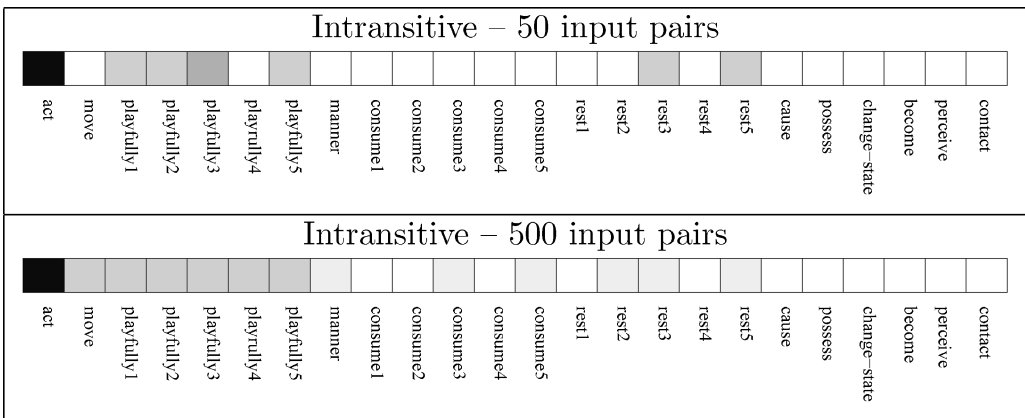


Fig. 12. Intransitive construction for the extended lexicon. *Note:* Darker squares correspond to higher probabilities.

of those primitives associated with verbs from the same semantic “class” (e.g., all the verbs *sit_i*) have very close probabilities after 500 input pairs (but not after 50 pairs). These results confirm that the general properties of a frequent construction can be the result of probabilistic association between the semantic primitives and the form. In the long run, this is determined by the frequencies of all the verbs that appear in that construction.

6.2. Learning curves

The simulations reported above present one view of learning in the model over time, focusing on the changing probabilistic associations within the constructions that are gradually learned. Here, we turn to a different view of development in the model, focusing on the learning curve that reflects the accuracy of the model over time in using this acquired knowledge, and the effect of noise on the performance of the model. To investigate this, we look at the output for sentence production (i.e., how the syntactic pattern that the model predicts based on available semantic information varies over time as the model receives more input data).

In each of the following simulations, we ran our model over 600 input pairs divided into 100 sequences. Each sequence had 5 training input pairs, followed by a 6th test input in which the utterance was removed. Our prediction model was used to find the syntactic pattern with the highest probability for the 6th (partial) test frame of each sequence. (The model does not use the test input in learning.) After each test input, we measured the cumulative accuracy of the model by counting the total number of times, up through that sequence, that the predicted pattern for a test input was exactly the same as that used in the removed utterance. (Note that this is a very conservative estimate of correctness because another syntactic pattern may express the given semantics equally acceptably.) We first performed a set of 10 such simulations to examine the learning curve for the model over all verbs. Fig. 13 shows the learning curves obtained from 4 simulations on different randomly generated input corpora. Each panel of the figure includes the learning curves for both the original input corpus and its “noisified” version (see section 5.3); this allows us to compare the effect of noisy and partial data on the performance and convergence of the model. Because the input corpora are randomly generated, the performance of the model varies across the simulations. However, we often observe early periods of fluctuation in the accuracy, followed by a more stable and accurate performance. It is evident from the graphs that the presence of noise mainly degrades performance of the model at the early stages of learning. As more input is received and processed, the model typically recovers from early mistakes, and its accuracy in the presence of noise often approaches its accuracy in the absence of noise. However, there are some cases where the order of inputs and early prediction mistakes make it impossible to completely recover from noise (see the top right panel of Fig. 13). In an actual child language situation with much more exposure to language, we expect that such scenarios would be avoided.

As a usage-based model that incorporates a generalization mechanism, we expect an overall U-shaped learning curve from our system in which early periods of more accurate performance are followed by less accurate performance, before more consistent performance is eventually achieved. However, because the graphs above report accuracies over a range of verbs and argument structures, the hypothesized U-shaped curve is only roughly evident. We also examine the individual learning curves of every verb to see whether particular verbs exhibit




Fig. 13. Cumulative accuracy of the model for predicting syntactic pattern. *Note:* Solid lines show performance with noiseless input data, and dashed lines show performance with noisy input data.

a clear U-shaped pattern. For each verb, we ran eight separate simulations using the same methodology as above, but in this case, we used only “noisified” corpora and the test input always contained the same target verb. Figure 14 shows four sample learning curves for each of the verbs *go* and *fall*. With testing of individual verbs, the variability in the performance of the model early on is more pronounced, due to variations in exposure to that verb in the random corpora. For frequent verbs with various argument structures, such as *go*, a U-shaped curve is often observed. The learning curve for *fall*, which is less frequent, shows a delay compared to more frequent verbs such as *go*, but the U-shaped curve can still be observed in most simulations.

Figures 13 and 14 show that model generally exhibits the characteristics of a U-shaped learning curve observed for children. In the following subsections, we turn to a more detailed description of the behavior of the model as it uses the knowledge it has acquired at various stages of learning along the “U”: specifically, during generalization (including an initial stage of imitation), possible over-generalization and recovery, and productive generalization.

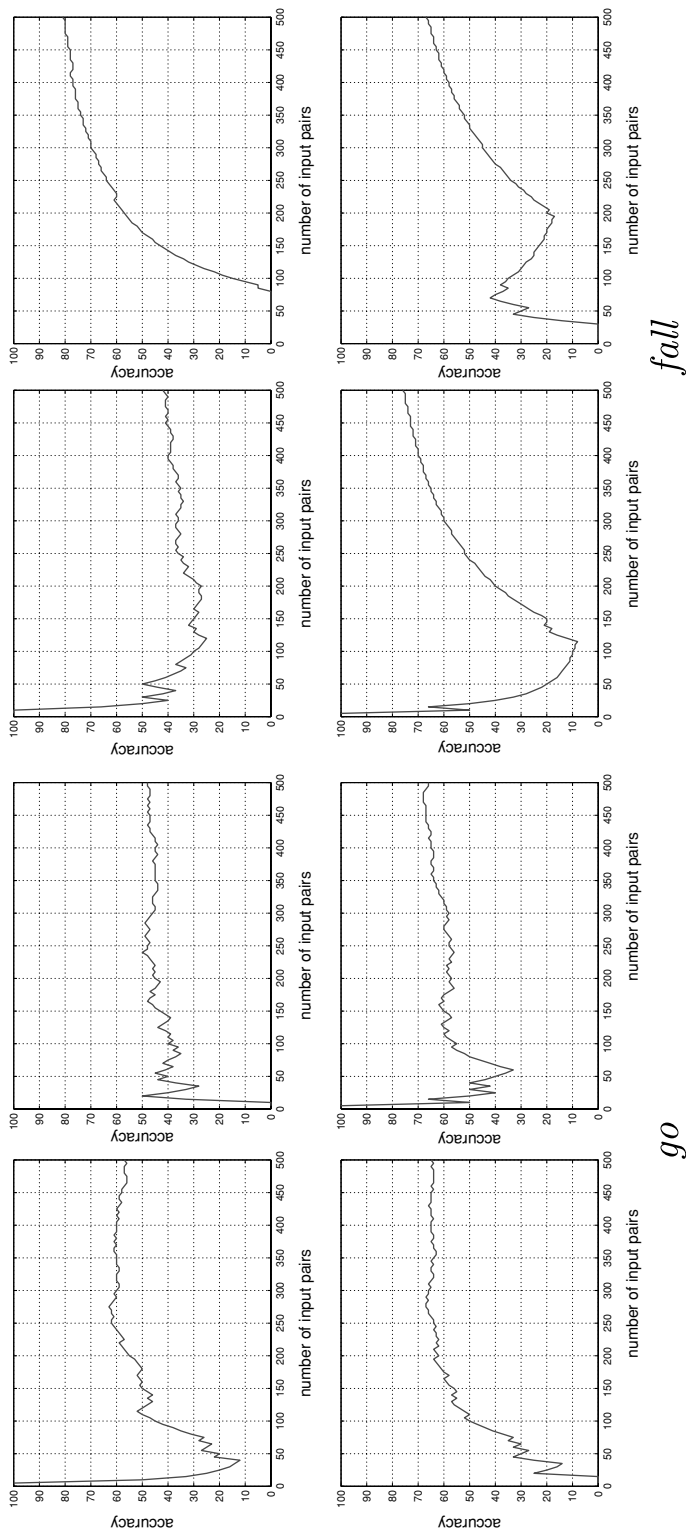


Fig. 14. Sample learning curves for different verbs, showing cumulative accuracy over time.

6.3. Generalization

Here, we investigate the ability of the model to generalize its item-specific knowledge in different tasks, drawing on the constructions it has acquired. First, we present the process of generalization in comprehension, showing that the model can use its learned constructions to recognize appropriate form-meaning pairings in the input. Next, we look at generalization in production, where the model uses the properties of a scene description to choose the best syntactic pattern with which to express the semantics. In both cases, we discuss the behavior of the model in the context of observations of actual child acquisition. We use the same corpora for both sets of simulations to compare the onset of generalization in production and comprehension, which is discussed in section 7.2.

6.3.1. Form-meaning mappings in comprehension

A number of experiments have indicated that children use the evidence of syntactic form to infer general semantic properties of a novel verb (e.g., Fisher, 2002; Naigles, 1990), a phenomenon sometimes referred to as *syntactic bootstrapping* (Gleitman, 1990). For example, in Naigles's preferential-looking experiment, children of age 2 to 3 were shown two pictures of actions, presented with an utterance with a novel (nonce) verb, and told to "find" the novel action in the pair of scenes. Naigles found that children who heard an intransitive utterance (*The bunny and duck are blinking*) were more likely to look at a picture of two characters independently performing an action (e.g., the bunny and duck each twirling their arms), whereas those who heard a similar transitive form (*The bunny is blinking the duck*) were more likely to look at a picture of one character (the bunny) performing an action on the other (the duck). This result demonstrates that the children have learned a reliable association between a syntactic form (such as the transitive) and a coarse semantics for the expressed event (i.e., one participant causally affecting another), and are able to determine the scene that is more compatible with an utterance according to this acquired knowledge.

In this section, we use our computational model to simulate Naigles's (1990) preferential looking experiment. As described in section 4, language use is modeled as choosing the best value j for a missing feature i , based on other observed features, using Equation 7. We use this approach to model sentence production in the next section. However, the same approach cannot be used to directly mimic the conditions of a preferential-looking experiment of the type used in Naigles's study, because here the problem is to choose from two possible pairings of scenes with the utterance (similar to children choosing one of the two videos being shown upon hearing the utterance). However, we can demonstrate the ability to create an association between form and meaning by showing a preference in the model for the combination of an appropriate scene paired with an utterance, compared to an inappropriate scene paired with that utterance.

First, we consider transitive and intransitive utterances, U_T and U_I respectively, analogous to those used in Naigles's (1990) experiment:

U_T : *Bunny blink duck*

U_I : *Bunny and duck blink*

Table 3
The four test input pairs used in the experiment of section 6.3.1

Label	Test Pair	Compatibility
U_T - S_{CA}	<i>Bunny blick duck.</i> $\text{Blick}_{[cause,act]}(\text{BUNNY}_{(agent)}\text{DUCK}_{(theme)})$	Matched
U_T - S_A	<i>Bunny blick duck.</i> $\text{Blick}_{[act]}(\text{And}_{[]}(\text{BUNNY.DUCK})_{(agent)})$	Unmatched
U_T - S_{CA}	<i>Bunny and duck blick.</i> $\text{Blick}_{[cause,act]}(\text{BUNNY}_{(agent)}\text{DUCK}_{(theme)})$	Unmatched
U_T - S_A	<i>Bunny and duck blick.</i> $\text{Blick}_{[act]}(\text{And}_{[]}(\text{BUNNY.DUCK})_{(agent)})$	Matched

Next, we create corresponding scene representations analogous to the pictures shown to the children in which one entity is causally affecting another, S_{CA} , or the two entities are performing independent actions, S_A :

$$S_{CA}: \mathbf{Blick}_{[cause,act]}(\text{BUNNY}_{(agent)}, \text{DUCK}_{(theme)})$$

$$S_A: \mathbf{Blick}_{[act]}(\text{And}_{[]}(\text{BUNNY}, \text{DUCK})_{(agent)})$$

These utterance and scene representations form the basis for input pairs used to test the model's acquired knowledge of the association between syntactic forms and semantic properties of a scene.

In one experimental condition, analogous to the child hearing the transitive form, each of the above scenes is combined with the transitive utterance, U_T , to form two input pairs, S_{CA} - U_T and S_A - U_T . The input pair S_{CA} - U_T corresponds to the appropriate selection of the causal action scene to go along with the perceived transitive utterance; we call this the "matched" pair. The input pair S_A - U_T corresponds to the inappropriate selection of the independent action scene for the utterance; we call this the "unmatched" pair. The other experimental condition uses analogous pairings with the intransitive form, U_I . The four resulting input pairs are shown in Table 3.

We then (separately) present each of these four input pairs to our model. The model will extract a frame F from the presented input pair, and will determine the best construction k for that frame F . The model records the value of $P(k|F)$ (Equation 1), for the best construction k given F , as its response to each input. If the input pair yields a frame that is compatible with an existing construction, this indicates that the scene-utterance combination corresponds to a form-meaning association in the model, represented as an existing construction. In this case, the value of $P(k|F)$ will be relatively high. If the extracted frame is not compatible with an existing construction, then the best construction for the input pair is a new one, which will have low prior probability. Thus, when comparing the values recorded in response to the matched and unmatched pairing for each utterance, a higher value of $P(k|F)$ corresponds to the child "recognizing" the appropriate scene for the utterance.

Table 4 shows the value of log $P(k|F)$ across the conditions, after varying amounts of learning (10, 50, 100, 150, and 500 training input pairs; averaged over 10 simulations). (We report the log probabilities because the actual probability values are very small.) The sizable

Table 4
Log $P(k|F)$ values for matched and unmatched scene-utterance pairs

Test Input Pair	# Training Input Pairs				
	10	50	100	150	500
U_T-S_{CA} (matched)	-7	-6	-6	-6	-6
U_T-S_A (unmatched)	-7	-7	-7	-8	-8
U_T-S_{CA} (unmatched)	-10	-10	-10	-10	-11
U_T-S_A (matched)	-4	-4	-3	-3	-3

difference between the two (matched and unmatched) pairs for each utterance type mimics the child's ability to pick out the appropriate scene for an utterance based on learned argument structure regularities. Even when the log values are the same (i.e., after 10 input pairs in the transitive condition), the actual score value for the matched pair (7.5×10^{-7}) is almost four times higher than the one for the unmatched pair (1.9×10^{-7}). Although this does not directly show that the model can infer aspects of the meaning of a verb from its syntactic usage, as syntactic bootstrapping proponents have suggested, we do demonstrate that ability in the later experiments on productive generalizations (section 6.5).

6.3.2. Generalization in sentence production

The above simulations show that the model can use its knowledge of constructions in comprehension. In this section, we turn to the use of constructions in production. Specifically, we examine the behavior of the model when presented with a novel verb for which it must produce a sentence. Children can, after a certain age, use a novel verb in a construction that has not been observed for that verb, when the construction is one to which there has been frequent exposure (Maratsos, Gudeman, Gerard-Hgo, & DeHart, 1987; Pinker, Lebeaux, & Frost, 1987; Tomasello, Akhtar, Dodson, & Rekau, 1997). In some experiments, a novel verb is modeled in some construction and then children are expected to produce it in a related construction (Maratsos et al., 1987; Pinker et al., 1987). In others, the verb is explicitly introduced through minimal syntactic contexts, such as *Look! Gorp!*, so that only the relation between the verb form and the semantics of the action is observed by the child (Tomasello et al., 1997).

We show a similar ability in our model through the following simulation. After training the model, we present it with the following input pair in which the scene representation is as given, and the utterance is left empty:

$$S_{G1} : \mathbf{Gorp}_{[cause, act]} (\text{KITTY}_{(agent)}, \text{DOGGY}_{(theme)})$$

$$U_{G1} : ?$$

The main predicate, **Gorp**, corresponds to a verb, *gorp*, which has not been previously seen by the model in an utterance. We build in the knowledge that the predicate **Gorp** is the meaning of *gorp*, analogous to children being told that "this [demonstrated action] is gorping," as in the Tomasello et al. (1997) experiments. Given the above utterance-less input, the model will extract a frame that has no syntactic (word order) pattern. We then monitor the syntactic pattern that the model predicts in response. Because the scene representation resembles the

Table 5
Average probability of the predicted patterns

Generated Pattern	# Input Pairs				
	10	50	100	150	500
Observed (SOV)	0.90	0.70	0.40	0.00	0.00
“Corrected” (SVO)	0.10	0.30	0.60	1.00	1.00

Note. SOV = Subject-Object-Verb; SVO = Subject-Verb-Object.

meaning of a typical transitive usage, we expect the model to predict the transitive syntactic pattern, “arg1 *verb* arg2,” despite not having seen usages of this particular verb. To observe the pattern of responses over the course of acquisition, we test the model after varying amounts of training data. Averaging over 10 simulations on different input corpora, the model predicts the expected transitive word order pattern with 100% probability after processing 50 input pairs. These results show that our model, like children, can start generalizing ubiquitous constructions to new verbs from early on.

More interesting is the varying influence of specific (verb-based) and general (construction-based) knowledge in production over the course of acquisition, which is a key feature of our model. To explore this, we mimic the conditions of an experiment by Akhtar (1999) in which English-speaking children aged 2 to 4 years old were taught novel verbs used in non-standard English word orders. (These were Subject-Object-Verb [SOV] or Verb-Subject-Object [VSO], rather than the standard Subject-Verb-Object [SVO]). In both spontaneous and elicited productions of these verbs, 2- and 3-year-olds imitated the observed patterns roughly half the time and “corrected” the order to the English SVO pattern roughly half the time. On the other hand, the 4-year-olds rarely imitated the observed order, almost always correcting to SVO order. Thus, the children show an early stage of imitation, followed by a later stage in which they generalize the pervasive word order pattern of their linguistic experience to a new verb, even when that verb has been observed in a different pattern.

We trained our model using different numbers of input pairs to simulate differing amounts of exposure to the language, roughly corresponding to different “ages.” We then provided the model the following training input pair with a novel (previously unseen) verb used in an utterance with a non-standard word order:

$S_{G2} : \mathbf{Gorp}_{[cause, act]} (\text{KITTY}_{(agent)}, \text{DOGGY}_{(theme)})$

$U_{G2} : \text{kitty doggy gorp}$

Note that the scene representation, S_{G2} , resembles observed scenes with transitive (SVO) utterances, but U_{G2} has a different (SOV) word order. Next, we presented the same input pair to the model, but without the utterance, as in S_{G1} - U_{G1} earlier, and recorded the syntactic pattern predicted by the model.

We performed 10 simulations for each number of training input pairs, and averaged the probability of predicting each of the SOV (observed) and SVO (“corrected”) patterns. As can be seen in Table 5, when 100 input pairs have been processed, the model predicts the observed

pattern a little more than half the time, and the corrected pattern a little less than half the time. This resembles the 2- and 3-year-old age group, who imitate and correct about equally often. However, when the number of input pairs is small, the model is more likely to imitate in production the single usage of the new verb that it has seen. When the number of input pairs goes up to 150 or more, the model predicts the corrected pattern almost all the time, which is similar to what the 4-year-old children do. These results show that the model, like children, shows a shift with more exposure to the language, from imitation, in which even a form that is incompatible with the ambient language is repeated, to generalization, where the ubiquitous patterns of the language predominate in production.

6.4. Over-generalization and recovery

In addition to enabling effective generalization of acquired knowledge, as shown earlier, the use of general knowledge sometimes leads children to over-generalize. However, children eventually recover with only additional positive evidence. A typical over-generalization in the area of argument structure is when a non-causative (intransitive-only) verb is used causatively (in a transitive structure). This occurs in an example such as *Don't you fall me down* (Bowerman, 1982) in which the intransitive verb *fall* is used transitively, on analogy to a verb like *drop*, which can be used in both forms. In the simulations reported here, we tracked the usage of *fall* by our model to see if we can detect a pattern of over-generalization and recovery similar to that in children.

The entry for *fall* in the input-generation lexicon allows only an intransitive syntactic pattern (as in *The blocks fell*). However, the scene representation for a use of *fall* may include the agent who caused the falling (e.g., if Adam pushed the blocks over) because this may be part of the observation of the child. Each use of *fall* in these simulations has an intransitive syntactic pattern with 1.00 probability, but includes a causal agent in the scene description with a 0.50 probability. We therefore expect the semantic similarity of the scene to that of a transitive construction to sometimes lead to over-generalization—that is, the prediction of a transitive syntactic pattern with *fall* in a scene with a causal agent, although this pattern has never been observed with *fall*.

In these simulations, we test the behavior of the model in producing a syntactic pattern for *fall* over the course of acquisition. Every 5 training input pairs was followed by a test pair in which the utterance was left empty, and the scene representation had **Fall** as its main predicate. We then record the prediction of the model for the missing syntactic pattern. (Note that the model does not learn from this test pair.) Over 10 such simulations, the average probability of receiving a training pair containing *fall* was .01, and the average probability of receiving an instance of the transitive construction was .16. In all 10 simulations, the model showed a pattern of over-generalization and recovery—using the transitive syntax for *fall* at some point, and eventually producing only correct intransitive forms as it received more examples of *fall* in the training input.

Figure 15 shows the first 8 uses of *fall* for Adam in CHILDES (at 27 months), together with the first 8 sentences generated by our model in one of the simulations. The output of the model illustrates the mix of the two (transitive and intransitive) syntactic patterns at this stage.

Adam	Our model
<i>go fall!</i>	<i>John fall ball</i>
<i>no no fall no!</i>	<i>toy fall</i>
<i>no fall!</i>	<i>Mary fall book</i>
<i>oh Adam fall.</i>	<i>toy fall</i>
<i>Adam fall toy.</i>	<i>cookie fall</i>
<i>Adam fall toy.</i>	<i>kitty fall</i>
<i>oh fall.</i>	<i>spoon fall</i>
<i>I not fall.</i>	<i>ball and toy fall</i>

Fig. 15. Comparing the first eight uses of *fall* for Adam with the first eight sentences generated by our model.

(Note that a number of the patterns shown in Adam's speech are unknown to our model, such as negatives and imperatives.)

6.5. Productive generalization

We noted that children sometimes mistakenly over-generalize, but eventually recover from these errors only by receiving additional positive evidence. However, this ability to converge on appropriate argument structures for each verb does not prevent a speaker from making productive generalizations in which a verb may be used in a construction that is "unusual" for it to convey particular semantic properties (Goldberg, 1995). For instance, in *The fly buzzed into the room*, the sound emission verb *buzz* adopts a manner of motion interpretation due to the directional movement construction it occurs in. We showed above that phases of over-generalization and recovery, similar to those of children, can be observed in the output of our model. In this section, we demonstrate that the model can handle instances of productive generalization as well. We test our model with a verb appearing in an unusual (for that verb) construction, and show that the model can determine appropriate semantic properties of the usage based on what it has learned for that construction when used with other verbs. We also show that the model can produce an appropriate syntactic form for a predicate used with semantic properties that are unusual for it.

We add a new verb *dance* to the input generation lexicon, with a single intransitive frame, as shown in Fig. 16. We train the model on 500 training input pairs, an amount corresponding to an advanced stage of learning. In all training pairs in which *dance* appears, it is used

WORD:	dance
SYMBOL:	DANCE
CATEGORY:	[predicate]
FRAME:	DANCE(\$1)
ARG-NO:	1
PRIMITIVES:	act,manner
ARGS:	\$1 is [agent]
CONDITION:	\$1 is [animate]
PATTERN:	\$1 [word]

Fig. 16. The lexical entry for *dance* in the input-generation lexicon.

WORD:	put
SYMBOL:	PUT
CATEGORY:	[predicate]
FRAME:	PUT(\$1,\$2,\$3)
ARG-NO:	3
SPRIMITIVES:	cause, move
ARGS:	\$1 is [agent]
ARGS:	\$2 is [theme]
ARGS:	\$3 is [destination]
CONDITION:	\$1 is [animate]
CONDITION:	\$2 is [concrete]
CONDITION:	\$3 is [dest-pred]
PATTERN:	\$1 [word] \$2 \$3

Fig. 17. The lexical entry for *put* in the input-generation lexicon.

intransitively. We then present the model with the following input pair in which *dance* is used in a different argument structure, and the semantic primitives of the predicate and the semantic roles of the arguments are omitted:

$$S_{P1} : \mathbf{Dance}_{[??]}(\text{KITTY}_{(?)}, \text{DOGGY}_{(?)}, \mathbf{Under}_{\square}(\text{TABLE}_{(?)})_{(?)})$$

$$U_{P1} : \text{kitty dance doggy under table}$$

The partial scene representation corresponds to a child hearing the verb *dance* in this novel syntactic form, and being able to determine the mapping of syntactic arguments to semantic argument positions (i.e., we have built in the output of that process into the scene representation because our current implementation cannot support the prediction of an entire scene representation from an utterance alone). Given this partial input, the model must predict the multiple missing semantic features, based on the utterance. Averaged across 10 simulations, the model predicts novel semantic primitives for the predicate **Dance** in this usage with a probability of .46 for *cause*, .33 for *move*, .19 for *possess*, and negligible probabilities for other primitives. The roles predicted for the three arguments, with associated probabilities, are *agent* (.99), *theme* (.99), and *destination* (.60), respectively. (The probability of the predicted role for the last argument is lower than the other two. Because other prepositional phrases may appear as the third argument in this syntactic form, different role assignments are possible for that argument.) Thus, the model assigns semantic properties that capture the causative movement intended by this novel usage of *dance*. Like humans, the model can successfully comprehend the meaning of a situation described by a sentence with a verb used in an unusual construction. It accomplishes this by generalizing the feature values of a construction associated with other verbs, in this particular case, one corresponding to the usage of a verb such as *put*, shown in Fig. 17.

We also test the ability of the model to use its learned constructions to correctly produce a syntactic form for a predicate used with an “unusual” semantics. Here, we present the model with an input pair having the completed scene representation corresponding to S_{P1} and the utterance missing, as in the following:

$$S_{P2} : \mathbf{Dance}_{[cause, act]}(\text{KITTY}_{(agent)}, \text{DOGGY}_{(theme)},$$

$$\mathbf{Under}_{\square}(\text{TABLE}_{(destination)})_{(destination)})$$

$$U_{P2} : ?$$

Table 6
 Probabilities of semantic primitives associated with a typical and “unusual” usage of the verb *dance*

Verb Semantic Primitives	# Input Pairs	
	50	500
Typical usage		
act	.44	.41
move	.34	.31
playfully	.18	.14
manner	.04	.09
consume	10^{-4}	.02
rest	10^{-4}	.01
cause	10^{-4}	10^{-3}
change-state	10^{-4}	10^{-5}
possess	10^{-4}	10^{-5}
become	10^{-4}	10^{-5}
perceive	10^{-4}	10^{-5}
contact	10^{-4}	10^{-5}
“Unusual” usage		
cause	.41	.46
move	.39	.33
possess	.04	.19
act	.02	10^{-2}
become	10^{-3}	10^{-3}
change-state	.09	10^{-4}
playfully	10^{-3}	10^{-5}
consume	10^{-3}	10^{-5}
rest	10^{-3}	10^{-6}
perceive	10^{-3}	10^{-6}
manner	10^{-3}	10^{-6}
contact	10^{-3}	10^{-6}

We then record the model’s prediction of the most probable syntactic pattern. The appropriate sentence, *kitty dance doggy under table*, is produced as expected.

Table 6 shows the probabilistic association of verb semantic primitives with the verb *dance* in a typical usage (e.g., *kitty dance*), and the “unusual” one (as in U_{P1}), after processing 50 and 500 training input pairs. It can be seen that the semantic primitives predicted for the unusual usage are as strong as those in a typical usage. This shows that the model is as powerful in making productive generalizations as in applying the item-based information in a typical usage of a verb.

7. Discussion of the results

The onset of abstract knowledge in children has been extensively studied, with the goal of determining when children internalize and use general patterns in their language. This is an important step in language acquisition, as it allows the learner to creatively and productively use their language. The issues surrounding this ability are particularly interesting in the domain of argument structure knowledge, which is a very complex aspect of language for children to master. The computational experiments presented here show that knowledge of argument structure constructions can be represented as probabilistic form-meaning associations that reflect generalizations over specific verb usages. Such constructions emerge in the model over time, reflecting stages that match well with those of children: early imitation of observed forms, later generalization of such observed knowledge to new situations, mistakes reflecting over-generalization and eventual recovery from such errors, and attainment of established knowledge that allows for productive generalization. Here, we discuss the reported behavior and underlying mechanisms of our model in the context of alternative accounts of argument structure acquisition and early use of such knowledge.

7.1. *The emergence of constructions*

There are a number of differing views on how the knowledge of constructions is acquired. One perspective is that argument structure patterns are initially learned on a verb-by-verb basis; then, through a process of categorization and generalization over the input, constructions are formed by associating certain meaning components with common syntactic patterns (Goldberg, 1999; Tomasello, 2000). In Goldberg's (1999) account, the generalization of constructional meaning is based largely on the meanings of highly frequent and semantically general "light" verbs such as *go*, *do*, *make*, *give*, and *put*. Goldberg (1995) showed that the usages of such light verbs correspond closely to the meanings associated with argument structure constructions. For example, the intransitive construction, in which the pattern "Subject_X Verb Oblique_Y" is paired with the meaning "X moves Y," corresponds to a core usage of the light verb *go*.⁷ Goldberg (1995) suggested that children record a correlation between a certain formal pattern and the meaning of a specific verb that is used earliest and most frequently in that pattern. Because light verbs are more frequent than other verbs and are learned early, they tend to be the ones around which constructional meaning centers. Thus, constructions associate the semantically general meanings of light verbs with basic syntactic forms.

However, Tomasello (2003) noted that the evidence from actual child language data is not so clearcut, and that a child may use a number of different verbs in a construction very early on in his or her learning, depending on the verbs he or she has been exposed to most frequently. This accords with our model in which constructions arise regardless of the generality of any particular verb's semantic properties. Because a construction in our model is comprised of a usage-based probabilistic association between syntactic and semantic properties, all verb usages with similar properties contribute to the establishment of the construction. Our experiments show that, although an increasing number of semantic primitives may, over time, become weakly associated with a construction given exposure to a greater variety of verbs, the most general semantic primitives will become strongly associated with the construction

because they will be observed across many more frames. For example, we showed that the general primitive *act* has a very high probability given a transitive form, compared to more specific primitives like *cause* and *become*, and even more so compared to very fine-grained primitives like *consume*. Although constructions may be likely to reflect light verb usages due to the frequency of such verbs, the emergence of a construction need not be dependent on the presence of a light verb as its core. Thus, this same probabilistic mechanism can explain both the emergence of semantically general constructions, such as the transitive or intransitive, as well as specific constructions (as in *He poisoned/wangled/hit/pitched/battled his way to a promotion*).⁸

7.2. Generalization in comprehension and production

The development of abstract argument structure knowledge in children appears to be reflected differentially in comprehension and production processes. Recent psycholinguistic evidence shows that very young children use verbs conservatively and only in those constructions in which they have heard them used before (Tomasello, 2000). It seems that children do not produce utterances in accord with general (verb-independent) regularities in the mapping of form to meaning. In contrast to such production evidence, however, experiments on language comprehension using preferential-looking techniques indicate that children show signs of being aware of such argument structure regularities much earlier (Fisher, 2002; Naigles, 1990). To explain this, Tomasello and Abbot-Smith (2002) suggested that children begin by learning weak constructions that enable only certain linguistic operations, such as comprehension. Over time, children develop a more robust construction representation that supports use in production.

This timing of construction use in different tasks can be observed in the behavior of our model without recourse to a notion of qualitatively different construction representations. Because our model seamlessly integrates specific verb-based and more general construction-based knowledge, it shows an interplay of imitation and generalization over the timecourse of learning that can account for the observed effects. Less-entrenched constructions emerge very early on (after 10 training input pairs, for a common construction), enabling the model to successfully recognize (comprehend) established form-meaning mappings, as shown in our simulation of the preferential-looking (comprehension) experiments of Naigles (1990). Appropriate use of such general knowledge in production is not established in the model until at least after 100 input pairs, as shown in our experiment that mimics that of Tomasello et al. (1997). Moreover, such general knowledge cannot consistently override contrasting verb-specific information until later (after 150 input pairs), as shown in our simulation of Akhtar's (1999) elicitation experiments. In this way, our model naturally accounts for the child acquisition data, without having to posit different mechanisms of construction representation and use over time.

7.3. Mechanisms for recovery from over-generalization errors

Although generalization over the observed input is an important means for a language learner to attain productive linguistic competence, it can also lead to mistaken application

of observed patterns. How children recover from such over-generalization errors has been the subject of much debate. There are a number of reasonable explanations of how children learn the productive patterns and constructions of their language, and why they over-generalize them at times. However, accounting for how children recover from making over-generalization errors is a challenge, especially given the general agreement that children are not being guided in the recovery process by negative evidence, in the form of explicit corrective feedback (Marcus, 1993). The complexity of this problem has made it a focal point in theories of argument structure acquisition.

Researchers who support an item-based approach to language acquisition have proposed a number of learning mechanisms as factors in recovery from over-generalization including entrenchment (Braine & Brooks, 1995; Goldberg, 1995), competition (MacWhinney, 1987), and cue construction (MacWhinney, 2004, drawing on Bowerman, 1982 and Pinker, 1984). However, none of these mechanisms alone can explain all aspects of the timecourse of development of the productive use of language. We briefly present each in turn.

Degree of entrenchment (Braine & Brooks, 1995) simply refers to the token frequency of an item. Once an argument structure pattern has been learned for a particular verb, that pattern tends to block the creative use of the verb in any other argument structure pattern, unless the competing pattern is also witnessed in the input. The role of entrenchment is more articulated in the Competition Model (MacWhinney, 1987). There the productive use of general linguistic patterns is viewed as a process governed by two types of pressures: the underlying analogic force that produces the generalization and the growth in the rote episodic auditory representation of an observed form. The entrenchment of the rote representation over time ensures the recovery from errors when the analogic process yields over-generalizations. Both of these proposals, however, focus on morphological and syntactic generalizations, dealing with such errors as *He falled the ball*. However, these accounts fail to explain errors that stem from semantic over-generalization of a construction. This occurs in situations such as the application of a resultative construction in **The child played the game finished*, on analogy to *The child hammered the putty flat*. Here, the over-generalized verb *played* is not semantically appropriate for this situation due to a lack of causation of a change of state (in contrast to *hammer*). To address this type of error, cue construction suggests that the process of recovery from over-generalization leads the child to construct new semantic features to block the application of the analogic pattern. Given the complementarity of these various mechanisms, MacWhinney (2004) suggested that all of them must be used by the learner, in order to account for different types of over-generalization errors and recovery from them. He also proposed that the learner estimates the “typical” rate of generalization for each form, whose distribution serves as “indirect negative evidence” in recovery from over-generalization. Goldberg (2006) similarly suggested combining multiple cues in a formal way to model productivity: Each cue is represented as a “weak hypothesis”; a boosting algorithm evaluates the prediction of each hypothesis against training items and changes the weight of that hypothesis accordingly. The output of the algorithm is a weighted sum of the predictions made by all the hypotheses. However, evaluating the hypotheses depends on the availability of some sort of feedback or negative evidence. Goldberg suggested that statistical preemption be used to estimate such feedback. (The process of statistical preemption is discussed further in the next section.)

The Bayesian approach we propose here uses a single core mechanism, the prediction model, to govern all aspects of the productivity of acquired constructions. This mechanism displays all the effects that the above diverse mechanisms collectively describe. For example, when faced with a language use task such as sentence production, the prediction model chooses the best syntactic form based on the probability of each of the possible forms. In calculating these probabilities, both of the “pressures” suggested by MacWhinney (1987) are involved: The underlying analogic force is represented in the high value for the posterior probabilities of the constructions that match the situation in hand, and the rote episodic memory of a correct form is reflected in the verb-specific knowledge learned by the model. Although in MacWhinney’s (2004) view this competition process handles syntactic generalization and recovery, he posited that cue construction is required for recovery from semantic over-generalization. In our model, however, we learn and store syntactic and semantic knowledge as a single unit—the argument structure frame—whose probabilistic associations are used to make both syntactic and semantic predictions. Therefore, what is called “cue construction” happens in our model analogously to the competition between syntactic forms: More fine-grained semantic features of each construction are learned and highlighted over time by processing more input data.

Both verb-specific and construction-based sources of predictions in our model are weighted by their corresponding frequencies in the input data (i.e., their degree of entrenchment). Higher frequency of a construction results in a higher over-generalization rate, whereas receiving and processing more instances of a correct specific form leads to recovery from over-generalization. This intuition is supported by Theakston (2004), who showed that human participants rate sentences with argument structure errors containing low frequency verbs more grammatical than those containing high frequency verbs. Because not seeing a particular usage in the input means both lower probability for that usage as well as higher probabilities for competing ones, indirect negative evidence in our model simply results from the differential frequencies of different usages in the input data, without any explicit calculation of generalization rates.

7.4. Comparing recovery and productive generalization

The process of entrenchment, or hearing a pattern with sufficient frequency, plays a key role in constraining over-generalization. However, despite entrenchment of observed forms, it is still the case that verbs can sometimes be used creatively in new argument structure patterns, without any trace of ill-formedness. On closer inspection, effects that have been ascribed solely to entrenchment may be better described as the outcome of statistical preemption, a more encompassing process that critically involves the role of semantic or pragmatic contrast (Goldberg, 2006). In this view, over-generalizations are constrained by the fact that more specific knowledge preempts general knowledge in determining a usage in production, given sufficient statistical support. However, this preemption can only occur if the more specific usage satisfies the functional demands of the context as well as the usage supported by the general knowledge. On the other hand, if a more general form better fits the semantic and pragmatic constraints of the situation, it may be felicitously used even in the presence of differing entrenched knowledge of the specific verb.

Entrenchment and statistical preemption are integrated into a single mechanism in our prediction model. The probability of each usage, as mentioned before, is a function of its

frequency. Therefore, the more entrenched (i.e., the more frequent) a usage is, the more weight it carries in the probabilistic comparison. Importantly, we compare frames, not just syntactic forms, so that the probabilities take into account the satisfaction of both syntactic and semantic properties of the situation. As a result, when a verb-specific frame becomes entrenched enough, it generally wins the competition with construction-based predictions, resulting in a decline over time in over-generalization errors. However, if the model encounters a situation for which none of the well-entrenched frames for a verb are suitable enough, the model uses a matching general construction that is compatible with the scene.

Unlike over-generalization errors, the ability of the model to generate novel utterances for unusual situations (or comprehend the meaning of the unusual utterances) does not fade over time by processing more input. Crucially, in these “unusual” cases, the model has not learned a verb-specific frame that sufficiently matches the partial frame to be processed. Therefore, the only reliable knowledge source is a matching construction (if such a construction exists). This property of the model embodies the interaction of entrenchment and statistical preemption in construction use suggested by Goldberg (2006).

A question that arises here is, what are the limitations of using a verb in a novel construction? Although many of the innovative uses of verbs in unusual constructions are acceptable, many others are not. The distinction seems to come from the fundamental semantic properties of the verb, and their compatibility with the scene to be described. For example, one can say *The child hammered the putty flat* but not **The child played the game finished* because *hammer* itself encodes a potential change of state to the object, whereas *play* does not. This is a distinction that our model cannot make in its current form. Addressing this problem in our model will require a more sophisticated, fine-grained semantic representation for verbs and scenes, augmented with appropriate semantic constraints.

8. Related computational models

Learning the meanings and appropriate usages of verbs is a challenge for children, as well as for computational approaches trying to model this aspect of human language acquisition. As in learning nouns (the words for objects), the learner must determine the meaning primitives that are associated with the verb in the face of referential uncertainty (e.g., as in Siskind, 1996; Yu et al., 2005). However, verbs pose the added complexity of having to determine the participants in the event described by the verb, how the participants relate to the event and to each other, and how they can be expressed in a syntactic form. Computational models of verb acquisition have tackled various aspects of this difficult learning problem.

Fern, Givan, and Siskind (2002) addressed the issue of acquiring the set of logical formulas that describe an event expressed by a verb, using a set of annotated videos of simple actions in a supervised learning setting. Fleischman and Roy (2005) instead focused on determining which aspect of an event is being described by an utterance, in forming the word-to-meaning mappings in an expectation-maximization framework. For example, their system determines whether an utterance is describing an action such as “get axe,” or a subcomponent of that action, such as “find axe.” Li, Farkas, and MacWhinney (2004) used a self-organizing neural network to model early lexical development. Their model shows the emergence of linguistic

categories, and lexical confusion and age-of-acquisition effects, but does not deal with any kind of ambiguity regarding the correct mapping between a form and a meaning representation. In our work, we simplify the assumptions along all of these lines, assuming a simple logical form representation for an action, and presupposing that the learner has determined the correct level of interpretation of the event. Future work in our framework will need to incorporate more extensive semantic representations, as well as the means for disambiguating the utterance with respect to the intended semantics.

A number of connectionist models have been proposed in recent years that attempt to learn syntactic structure from sequences of words. The most influential approach of this kind is due to Elman (1990, 1991), who trained a simple recurrent network to predict the next input word, for sentences generated by a small context-free grammar. The network learns an abstract representation of grammatical relations through learning lexical categories of the input words such as nouns and verbs (with further subcategorization of nouns as animate–inanimate, human–non-human, etc.), and the permitted combination of these categories in a sentence. Elman (1991) extended the previous model by using stimuli in which there are underlying hierarchical and recursive relations (e.g., subject nouns agree with their verbs; verbs have different argument structures; recursive relative clauses are permitted). Elman’s model showed very well the possibility of categorizing and learning abstract structures from strings of words. However, the artificial language it learns is purely syntactic, whereas natural language learning is crucially an attempt to discover the relation between meaning and linguistic form. In an attempt to fill the gap between the learning of syntax and semantics in connectionist models, McClelland and Kawamoto (1986) presented a connectionist model for assigning roles to constituents of sentences. They train the model on two sets of input: the surface structure of the sentence, and the thematic roles assigned to each constituent. Their model can learn the linking between the syntactic constituents and the thematic roles that it has seen in the input, and can also generalize these linking patterns to unseen data. However, the learning process is fully supervised, and the only form of language use performed by this model is limited comprehension.

Here, we focus particularly on learning the argument structure of a verb—the roles assigned by a verb to its participants, and the word order patterns in which those roles are expressed. Our work thus follows that of other recent research that has paid close attention to the relation between expressed meaning and syntactic form, and the importance of the syntax-semantics mapping in verb learning. Niyogi’s (2002) model, like ours, uses a Bayesian framework, and shows effects of both syntactic and semantic bootstrapping (i.e., the use of syntax to aid in inducing the semantics of a verb, and the use of semantics to narrow down possible syntactic forms in which a verb can be expressed). In contrast to our model, Niyogi’s approach relies on extensive prior knowledge in the form of a Bayesian hypothesis space and the probabilities over it. Desai’s (2002) work also demonstrated the role of syntactic and semantic bootstrapping, but does so using a very restricted set of semantic features. Moreover, the model shows only limited generalization abilities, where a major focus of our work lies in showing how verb-specific knowledge can be generalized in a way that supports acquisition and exhibits child-like stages of behavior. Allen (1997) proposed a connectionist system that focuses particularly on learning the mapping between roles and the entities that play those roles in an event. Due to the use of a distributed “microfeature” representation of the semantics, this model is able

to make interesting generalizations over argument structure syntax and semantics. However, the acquired knowledge is only used for limited comprehension, the potential of the model in performing more varied tasks of language use is not investigated, and a more explicit analysis of the formation of the argument structure constructions in the model is unavailable.

Other computational models are more similar to ours in honing in on the form-meaning mapping that must be acquired for verbs, specifically looking at the acquisition of constructions. Chang (2004) presented a model for learning lexically specific multiword constructions from annotated child-directed transcript data. The goal of the model is to learn associations between form relations (typically word order) and meaning relations (typically participant role-filler bindings), and to use them in language comprehension. However, in contrast to our approach, this model can only acquire verb-specific constructions that cannot be generalized to other verbs. Another approach for learning grammatical constructions by pairing form and meaning is presented by Dominey (2003) and Dominey and Inui (2004), in a system that learns from video. Although the model has some ability to generalize its knowledge to new verbs, its learning is highly dependent on the unrealistic assumption of having each syntactic form uniquely identify the associated meaning (i.e., forms and meanings are in a one-to-one mapping). Both of these models are thus limited in their ability to acquire knowledge of general verb argument structure patterns and to use that knowledge robustly when faced with new verbs or novel semantic combinations.

Many computational systems have instead addressed the problem of the acquisition of general grammatical knowledge (e.g., Clark, 2001a, 2001b; Gobet, Freudenthal, & Pine, 2004; Jones, Gobet, & Pine, 2000; Solan, Horn, Ruppin, & Edelman, 2004). These approaches extract general syntactic rules from sentential data, in contrast to our goal of learning form-meaning mappings, especially those that encapsulate verb argument structure. The latter relies crucially on the interaction between lexically-specific knowledge and general patterns over them. It is interesting to note that Onnis, Roberts, and Chater (2002) studied the problem of recovering from over-generalization in grammar learning in the context of an artificial language. They show that a grammatical representation that incorporates both general rules and lexical exceptions provides a simpler encoding as the learner is exposed to larger amounts of data. Although very different from our model in its intent, this work lends support to our assumption that it is important to balance verb-specific knowledge with that of general constructions for effective learning.

9. Conclusions

We have described a Bayesian model for the representation, acquisition and use of argument structure constructions in a usage-based framework. The proposed model suggests a novel description for the nature of constructions. Instead of formulating constructions as pairs of form and meaning, we view each construction as a probabilistic association between syntactic and semantic properties of verbs and their arguments. This probabilistic association emerges over time through a Bayesian acquisition process in which similar verb usages are detected and grouped together to form general constructions, based on their syntactic and semantic properties. The model demonstrates the feasibility of learning a range of abstract constructions from examples of verb usage even in the presence of noisy or incomplete input data.

Although the model reflects the verb-based nature of argument structure acquisition, a main advantage of our model is its explicit formulation of the interaction between verb-based and construction-based knowledge in comprehension and production, which contributes to the stages of learning exhibited by children. The suggested Bayesian mechanism not only accounts for the correct use of verbs in their appropriate argument structures and the gradual cessation of over-generalization errors, it also shows the possibility of productive generalization of constructions in the adult language facility. Many factors have been suggested in the linguistic and psycholinguistic literature for explaining these effects, such as entrenchment, competition, cue construction, and statistical preemption. Our proposed prediction model of language use unifies these factors, demonstrating their smooth integration in the processes of generalization and of recovery from over-generalization.

Despite the success of the model in accounting for observed behaviors over the timecourse of child argument structure acquisition, a number of limitations must be addressed in future work. For example, although we have not built in the knowledge of possible constructions, or innate principles of syntax-semantics mappings, we do assume that the child can figure out the appropriate participant roles for observed entities, drawing on a predefined set of possible roles. In an extended version of this work (Alishahi & Stevenson, 2007b), we directly address the learning and use of semantic roles. The modified model associates each argument position of a predicate with a probability distribution over a set of semantic properties of the arguments that appear in that position. We call this distribution over semantic properties a *semantic profile*. The preliminary results show that initially the semantic profiles of an argument position yield verb-specific conceptualizations of the role associated with that position. As the model is exposed to more input, these verb-based roles gradually transform into more abstract representations that reflect the general properties of arguments across the observed verbs, as has been suggested in the literature (e.g., Tomasello, 2000).

As noted earlier, we also assume that the model has learned the mapping between each word and its meaning (i.e., semantic symbol) before argument structure acquisition begins. This is an unrealistic assumption because the two processes are more likely to be interleaved in learning. Extensions to the model will require integrating a word-to-meaning mapping algorithm, similar to that of Siskind (1996) or Fleischman and Roy (2005). This approach may allow the model to use its acquired argument structure knowledge to facilitate learning the meaning of new words, which is a natural extension of the syntactic bootstrapping behavior demonstrated here.

Currently, our model does not record the actual words that participate in the acquired argument structure frames, instead replacing words by their semantic categories. These categories may later be generalized to their superordinates in the semantic hierarchy, thus enabling the acquisition of general patterns of argument types from instances of verb usages. Although this approach supports learning of general constructions, it entails that constructions that depend on specific word forms cannot be acquired. This includes constructions in the adult repertoire such as “to not *verb*₁ let alone *verb*₂,” which depends on the words “let alone.” Moreover, some researchers have hypothesized that children’s early constructions may depend on specific words or morphemes, as in *I’m Ving it* or *It Ved* (Childers & Tomasello, 2001; Lieven, Pine, & Rowland, 1998). We are currently incorporating into our model of

early verb learning a mechanism for word-specific constructions to provide a fuller account of construction representation, acquisition, and use.

Notes

1. For example, Gleitman (1990) argued that paired verbs such as *chase* and *flee* can always be used to describe the same event. Therefore, just by watching a chasing scene, it is impossible for the child to decide whether the verb describing the scene means “chase” or “flee.”
2. In Alishahi and Stevenson (2005a), we referred to these categories as *classes*; here, we use the term *construction* to refer to a group of similar frames.
3. The formulas used in Equation 3 and Equation 4 are similar to those used by Anderson (1991), $P(k) = \frac{cn_k}{(1-c)+cn}$ and $P(0) = \frac{(1-c)}{(1-c)+cn}$, with his “coupling probability” c set to the mid-value of .5. This value determines the degree to which items “prefer” to be grouped together. The mid-value makes the fewest assumptions, by avoiding a strong inclination to either prefer or disprefer grouping an item into an existing category.
4. Smoothing techniques, introduced in Good (1953), have been extensively used to enable statistical models of language to effectively deal with sparse data (Pereira, 2000). We make no claims regarding the psychological validity of this particular means for smoothing.
5. In Alishahi and Stevenson (2007a), we automatically reconstruct the logical representation of the meaning of each sentence based on its parse tree from a processed corpus, and extract the semantic properties of words from WordNet. However, the semantic roles of the arguments and the event-specific semantic primitives of the predicates cannot be automatically added.
6. In predicting a syntactic pattern for an incomplete frame, it is possible that the pattern will have place holders for more arguments than are present in the scene representation. In cases such as these, when creating the corresponding utterance, the excess argument slots in the predicted pattern are simply left blank.
7. Here, *Oblique* refers to a prepositional or adverbial argument indicating a direction or goal, as in *Hans goes to school* or *Ingrid went home*.
8. Examples adapted from Google searches on “his way to a promotion” (retrieved June 19, 2006).

Acknowledgments

In this work, we have greatly benefited from the interaction with the members of the Computational Linguistics group at the Department of Computer Science, University of Toronto. Many thanks to them and to the three anonymous reviewers for their insightful comments and suggestions, which helped us improve the article. We are also grateful for the financial support from the Natural Sciences and Engineering Research Council of Canada of Canada and the University of Toronto.

References

- Akhtar, N. (1999). Acquiring basic word order: Evidence for data-driven learning of syntactic structure. *Journal of Child Language*, 26, 339–356.
- Alishahi, A., & Stevenson, S. (2005a). The acquisition and use of argument structure constructions: A Bayesian model. In *Proceedings of the 2nd workshop on Psycho-Computational Models of Human Language Acquisition, ACL-2005* (pp. 82–90). Ann Arbor, MI: Association for Computational Linguistics.
- Alishahi, A., & Stevenson, S. (2005b). A probabilistic model of early argument structure acquisition. In *Proceedings of the 27th annual conference of the Cognitive Science Society* (pp. 97–102). Stresa, Italy.
- Alishahi, A., & Stevenson, S. (2007a). A cognitive model for the representation and acquisition of verb selectional preferences. In *Proceedings of the ACL-2007 workshop on Cognitive Aspects of Computational Language Acquisition* (pp. 41–48). Prague, Czech Republic: Association for Computational Linguistics.
- Alishahi, A., & Stevenson, S. (2007b). A computational usage-based model for learning general properties of semantic roles. In *Proceedings of the 2nd European Cognitive Science Conference* (pp. 425–430). Delphi, Greece.
- Allen, J. (1997). Probabilistic constraints in acquisition. In A. Sorace, C. Heycock, & R. Shillcock (Eds.), *Proceedings of the GALA '97 conference on Language Acquisition* (pp. 300–305). Edinburgh, Scotland: HCRC.
- Anderson, J. R. (1991). The adaptive nature of human categorization. *Psychological Review*, 98, 409–429.
- Bowerman, M. (1982). Evaluating competing linguistic models with language acquisition data: Implications of developmental errors with causative verbs. *Quaderni di Semantica*, 3, 5–66.
- Bowerman, M. (1988). The “no negative evidence” problem. In J. Hawkins (Ed.), *Explaining language universals* (pp. 73–101). Oxford, England: Blackwell.
- Bowerman, M. (1996). Argument structure and learnability: Is a solution in sight? In *Proceedings of the Berkeley Linguistics Society* (Vol. 22, pp. 454–468).
- Braine, M. D. S., & Brooks, P. (1995). Verb argument structure and the problem of avoiding an overgeneral grammar. In M. Tomasello & W. E. Merriman (Eds.), *Beyond names for things: Young children's acquisition of verbs* (pp. 353–376). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.
- Chang, N. (2004). Putting meaning into grammar learning. In *Proceedings of the 1st workshop on Psycho-Computational Models of Human Language Acquisition* (pp. 17–24). Association for Computational Linguistics.
- Childers, J. B., & Tomasello, M. (2001). The role of pronouns in young children's acquisition of the English transitive construction. *Developmental Psychology*, 37, 739–748.
- Clark, A. (2001a). Unsupervised induction of stochastic context free grammars with distributional clustering. In *Proceedings of Computational Natural Language Learning* (pp. 105–112). Toulouse, France: Association for Computational Linguistics.
- Clark, A. (2001b). *Unsupervised language acquisition: Theory and practice*. Unpublished doctoral thesis, University of Sussex.
- Desai, R. (2002). Bootstrapping in miniature language acquisition. *Cognitive Systems Research*, 3, 15–23.
- Dominey, P. F. (2003). Learning grammatical constructions in a miniature language from narrated video events. In *Proceedings of the 25nd annual conference of the Cognitive Science Society* (pp. 354–359). Boston: Cognitive Science Society.
- 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. In *Proceedings of the 1st workshop on Psycho-Computational Models of Human Language Acquisition* (pp. 33–40). Association for Computational Linguistics.
- Elman, J. L. (1990). Finding structure in time. *Cognitive Science*, 14, 179–211.
- Elman, J. L. (1991). Distributed representation, simple recurrent networks, and grammatical structure. *Machine Learning*, 7, 195–225.
- Fern, A., Givan, R., & Siskind, J. M. (2002). Specific-to-general learning for temporal events with application to learning event definitions from video. *Journal of Artificial Intelligence Research*, 17, 379–449.

- Fillmore, C., Kay, P., & O'Connor, M. K. (1988). Regularity and idiomatcity in grammatical constructions: The case of *let alone*. *Language*, 64, 501–538.
- Fisher, C. (2002). Structural limits on verb mapping: The role of abstract structure in 2.5-year-olds' interpretations of novel verbs. *Developmental Science*, 5, 55–64.
- Fisher, C., Hall, G., Rakowitz, S., & Gleitman, L. (1994). When it is better to receive than to give: Syntactic and conceptual constraints on vocabulary growth. *Lingua*, 92, 333–375.
- Fleischman, M., & Roy, D. (2005). Intentional context in situated language learning. In *Proceedings of the 9th conference on Computational Natural Language Learning* (pp. 104–111). Association for Computational Linguistics.
- Gleitman, L. (1990). The structural sources of verb meanings. *Language Acquisition*, 1, 135–176.
- Gobet, F., Freudenthal, D., & Pine, J. M. (2004). Modelling syntactic development in a crosslinguistic context. In *Proceedings of the 1st workshop on Psycho-Computational Models of Human Language Acquisition* (pp. 53–60). Association for Computational Linguistics.
- Goldberg, A. (1995). *Constructions: A construction grammar approach to argument structure*. Chicago: University of Chicago Press.
- Goldberg, A. (1999). The emergence of the semantics of argument structure constructions. In B. MacWhinney (Ed.), *Emergence of language* (pp. 197–212). Mahwah, NJ: Lawrence Erlbaum Associates, Inc.
- Goldberg, A. (2006). *Constructions at work: The nature of generalization in language*. New York: Oxford University Press.
- Good, I. J. (1953). The population frequencies of species and the estimation of population parameters. *Biometrika*, 40, 237–264.
- Grimshaw, J. (1990). *Argument structure*. Cambridge, MA: MIT Press.
- Jackendoff, R. (1990). *Semantic structures*. Cambridge, MA: MIT Press.
- Jones, G., Gobet, F., & Pine, J. M. (2000). A process model of children's early verb use. In L. R. Gleitman & A. K. Joshi (Eds.), *Proceedings of the 22th annual conference of the Cognitive Science Society* (pp. 723–728).
- Lakoff, G. (1987). *Women, fire and dangerous things: What categories reveal about the mind*. Chicago: University of Chicago Press.
- Langacker, R. (1999). *Grammar and conceptualization*. Berlin: de Gruyter.
- Levin, B. (1993). *English verb classes and alternations: A preliminary investigation*. Chicago: University of Chicago Press.
- Li, P., Farkas, I., & MacWhinney, B. (2004). Early lexical development in a self-organizing neural network. *Neural Networks*, 17, 1345–1362.
- Lieven, E., Pine, J., & Rowland, G. (1998). Comparing different models of the development of the English verb category. *Linguistics*, 36, 807–830.
- MacWhinney, B. (1982). Basic syntactic processes. In S. Kuczaj (Ed.), *Language development: Syntax and semantics* (Vol. 1, pp. 73–136). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.
- MacWhinney, B. (1987). The competition model. In B. MacWhinney (Ed.), *Mechanisms of language acquisition* (pp. 249–308). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.
- MacWhinney, B. (1995). *The CHILDES project: Tools for analyzing talk* (2nd ed.). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.
- MacWhinney, B. (2004). A multiple process solution to the logical problem of language acquisition. *Journal of Child Language*, 31, 883–914.
- Manning, C. D., & Schütze, H. (2001). *Foundations of statistical natural language processing*. Cambridge, MA: MIT Press.
- Maratsos, M., Gudeman, R., Gerard-Ngo, P., & DeHart, G. (1987). A study in novel word learning: The productivity of the causative. In B. MacWhinney (Ed.), *Mechanisms of language acquisition* (pp. 89–113). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.
- Marcus, G. F. (1993). Negative evidence in language acquisition. *Cognition*, 46, 53–85.
- Marcus, G. F., Pinker, S., Ullman, M., Hollander, M., Rosen, T. J., & Xu, F. (1992). Overregularization in language acquisition. *Monographs of the Society for Research in Child Development*, 57(4, Serial No. 228).

- McClelland, J. L., & Kawamoto, A. H. (1986). *Mechanisms of sentence processing: Assigning roles to constituents of sentences*. Cambridge, MA: MIT Press.
- Naigles, L. (1990). Children use syntax to learn verb meanings. *Journal of Child Language*, 17, 357–374.
- Niyogi, S. (2002). Bayesian learning at the syntax-semantics interface. In W. D. Gray & C. Schunn (Eds.), *Proceedings of the 24th annual conference of the Cognitive Science Society* (pp. 697–702).
- Onnis, L., Roberts, M., & Chater, N. (2002). Simplicity: A cure for overgeneralizations in language acquisition? In W. D. Gray & C. Schunn (Eds.), *Proceedings of the 24th annual conference of the Cognitive Science Society* (pp. 720–725).
- Pereira, F. (2000). *Formal grammar and information theory: Together again?* London: Royal Society.
- Pinker, S. (1984). *Language learnability and language development*. Cambridge, MA: Harvard University Press.
- Pinker, S. (1989). *Learnability and cognition: The acquisition of argument structure*. Cambridge, MA: MIT Press.
- Pinker, S., Lebeaux, D. S., & Frost, L. A. (1987). Productivity and constraints in the acquisition of the passive. *Cognition*, 26, 195–267.
- Quine, W. (1960). *Word and object*. Cambridge, MA: MIT Press.
- Rappaport Hovav, M., & Levin, B. (1998). Building verb meanings. In M. Butt & W. Gender (Eds.), *The projection of arguments: Lexical and computational factors* (pp. 97–134). Stanford, CA: CSLI Publications.
- Siskind, J. M. (1996). A computational study of cross-situational techniques for learning word-to-meaning mappings. *Cognition*, 61, 39–91.
- Solan, Z., Horn, D., Ruppin, F., & Edelman, S. (2004). Unsupervised context sensitive language acquisition from a large corpus. In L. Saul (Ed.), *Advances in neural information processing* (Vol. 16, pp. 961–968). Cambridge, MA: MIT Press.
- Theakston, A. L. (2004). The role of entrenchment in children's and adults' performance on grammaticality judgment tasks. *Cognitive Development*, 19, 15–34.
- Tomasello, M. (2000). Do young children have adult syntactic competence? *Cognition*, 74, 209–253.
- Tomasello, M. (2003). *Constructing a language: A usage-based theory of language acquisition*. Cambridge, MA: Harvard University Press.
- Tomasello, M., & Abbot-Smith, K. (2002). A tale of two theories: Response to Fisher. *Cognition*, 83, 207–214.
- Tomasello, M., Akhtar, N., Dodson, K., & Rekau, L. (1997). Differential productivity in young children's use of nouns and verbs. *Journal of Child Language*, 24, 373–387.
- Yu, C., Ballard, D. H., & Aslin, R. N. (2005). The role of embodied intention in early lexical acquisition. *Cognitive Science*, 29, 961–1005.