Towards the emergence of verb-general constructions and early representations for verb entries: Insights from a computational model

Judith Gaspers (jgaspers@cit-ec.uni-bielefeld.de)
Semantic Computing Group, CITEC, Bielefeld University

Anouschka Foltz (anouschka.foltz@uni-bielefeld.de)
Emergentist Semantics Group, CITEC, Bielefeld University

Philipp Cimiano (cimiano@cit-ec.uni-bielefeld.de)
Semantic Computing Group, CITEC, Bielefeld University

Abstract

Recent findings suggest that i) children can build initial verb entries on the basis of syntactic information alone without any additional information provided by a visual context, and ii) that the early representation of verbs encompasses statistical information on the co-occurrence of these verbs with their potential meanings/referents, enabling children to infer verb meanings under referential uncertainty. In this paper we present a computational model that acquires verb-general constructions under referential uncertainty. The model stores linguistic knowledge in line with construction grammar in the form of an interrelated network of constructions. Learning proceeds in line with usage-based theories in an item-based fashion. Computational results show that the model can account for the above-mentioned findings: The model produced patterns similar to those observed in these studies. Our findings hence shed light on the potential mechanisms involved in the emergence of early verb entries and verb-general constructions as well as the representation and refinement of verb entries.

Introduction

Unlike nouns, verbs describe actions that involve a number of participants who play certain (thematic) roles in the event. Hence, sentence structure, i.e. syntactic frames, may serve as a "zoom lens" to guide the child's attention to relevant aspects of verb meaning, in particular to thematic relations during verb learning (e.g. Gleitman & Fisher, 2005). In line with this assumption, Arunachalam & Waxman (2010) showed that 27-month-old children can create an initial verb entry based on information from the syntactic context and without any visual information, and retrieve this information when encountering the verb later on.

Scott & Fisher (2012) further provide evidence that 2.5-year-old children are able to use cross-situational statistics to infer verb meanings under *referential uncertainty* (Quine, 1960). This mechanism is also called *cross-situational learning* (e.g. Siskind, 1996) and has typically been investigated in the context of mapping nouns/single words to objects (e.g. Smith & Yu, 2008). Scott & Fisher (2012) showed that children can abstract across different actors and objects, suggesting that they can attach information about possible referents to novel verb entries along with their co-occurrence statistics and refine these entries over time.

However, what remains unclear is i) how verb-general constructions emerge and how they are represented, ii) how they can guide attention to establish verb entries based on syntactic information alone, iii) how information about possible referents and co-occurrence statistics might be stored with verb entries, and iv) how this information is updated incrementally over time, thus allowing for learning of verb meanings across situations.

In this paper we present a computational model that acquires verb-general constructions under referential uncertainty in order to shed light on the potential learning mechanisms involved in early verb acquisition. Specifically, we extend a previous usage-based computational model that can learn verb-specific constructions (Gaspers & Cimiano, in press) to also acquire verb-general constructions, which can explain the empirical findings of the above mentioned studies. In line with construction grammar (Goldberg & Suttle, 2010), the model represents linguistic knowledge in the form of an interrelated network of constructions containing both item-based knowledge and generalizations. Linguistic knowledge evolves over time in an item-based or bottom-up fashion. Abstraction from the observed input occurs as proposed in usage-based approaches (Behrens, 2009). The existing model first establishes mappings for lexical units. Only once some lexical knowledge has been learned with sufficient confidence, verb-specific constructions are learned in a bottom-up fashion by bootstrapping on the mappings for single lexical units. Thus, learning occurs through a generalization process which searches for syntactic variation corresponding to semantic variation.

The extended model presented here builds on these basic principles: Verb-general constructions are learned bottom-up from verb-specific constructions only once verb-specific knowledge has been derived with sufficient confidence. Again, generalization occurs in an item-based fashion (albeit with respect to more complex structures/mappings) by searching for variation at the linguistic layer which has corresponding variation at the meaning layer. We present empirical results replicating Arunachalam & Waxman's (2010) and Scott & Fisher's (2012) studies with the model. Depending on its "age", the model behaves very similarly to the children in these studies. The results suggest possible learning mechanisms implicated in the early acquisition and representation of verbs and verb-general constructions.
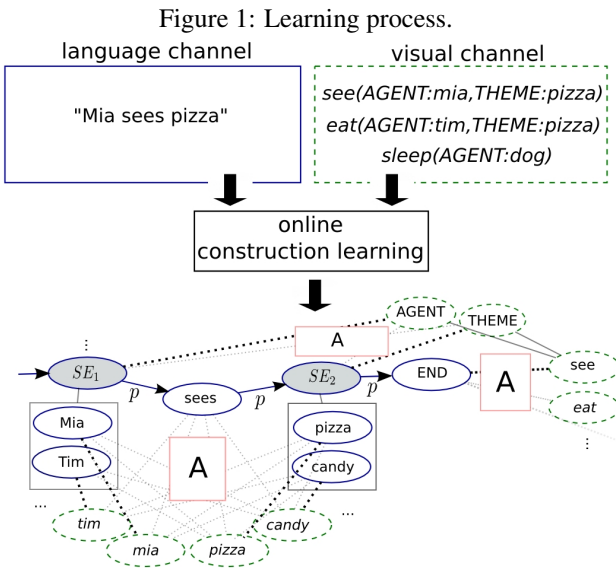
The computational model

In the following, we will first briefly describe the existing model that can learn verb-specific constructions (Gaspers & Cimiano, in press). Then, we present extensions to this model that allow it to acquire verb-general constructions.

Background

The existing computational model relies on an interrelated network of constructions, which are acquired incrementally on the basis of observed input. It encodes a construction grammar as a set of nodes and (weighted) edges. The model's input is a natural language (*NL*) sentence, i.e. a sequence of words, coupled with a symbolic description of the visual context in the form of meaning representations (*MR*) expressed by way of predicate logic. Each action $mr_i \in MR$ consists of a predicate ξ along with a set of thematic relations. It is important to note that the model can learn under referential uncertainty, i.e. given a situational context with several different actions out of which at most one corresponds to the utterance. The model input thus consists of two (temporally paired) channels: a language channel and a channel with information about the visual context. The learning process and an example of a verb-specific construction stored in the network are illustrated in Fig. 1.

The learned network consists of two subnetworks, one representing lexical and one representing syntactic constructions. The syntactic subnetwork builds on the lexical subnetwork and is divided into two sublayers: a slot-and-frame (S&F) pattern layer and a mapping layer. The **lexical subnetwork** encodes simple phrases, i.e. (short sequences of) words, as nodes together with the associated semantic referents, e.g. the word "Tim" and the corresponding semantic referent *tim* in Fig 1. The **S&F pattern layer** represents syntactic constructions as sequences of nodes that together constitute a path. Paths can contain variable nodes that represent slots in the

Figure 1: Learning process.

syntactic pattern. These slots can be filled with elements contained in specific groupings. This layer also encodes the associated semantic frames. For instance, in Fig. 1 a syntactic construction is represented as a path p which expresses a pattern "SE_1 sees SE_2", where SE_1 and SE_2 represent syntactic slots in the pattern, which can be filled with groupings of elements such as "Mia" and "Tim" in the case of SE_1 or "pizza" and "candy" in the case of SE_2. The semantic frame associated with the pattern is *see(AGENT,THEME)*. The **mapping layer** contains networks representing construction-specific argument mappings between syntactic patterns and semantic frames together with mappings of the syntactic arguments to semantic arguments. For example, in Fig. 1 an individual mapping network captures the correspondences between SE_1 and *AGENT* as well as SE_2 and *THEME*.

The network contains nodes of two types: Nodes representing **linguistic entities** include i) simple phrases including noun phrases, e.g. "Tim" or "the cat", ii) syntactic patterns represented as paths with slots, e.g. "SE_1 sees SE_2", and iii) syntactic slots of constructions represented as sets of elements containing all the simple phrases that can fill the slot, e.g. SE_1 = [Mia → *mia*, Tim → *tim*]. Nodes representing **semantic entities** include i) simple semantic referents, i.e. individuals such as *tim*, ii) semantic frames, e.g. *see(AGENT,THEME)*, and iii) argument slots of frames, e.g. *AGENT*.

Our computational model applies Hebbian-style cross-situational learning to establish connections between linguistic and semantic nodes that are activated concurrently in a certain linguistic and visual context. The model thus learns correspondences between linguistic and semantic nodes, i.e. the semantics of linguistic constructions. In all cases, we apply associative networks to capture co-occurrence frequencies (cf. Rojas, 1993).

During learning, input examples are processed one-by-one, causing immediate changes in the network structure. Learning is roughly divided into two learning steps: i) update of the lexical layer, where connections between lexical units and semantic referents are established, and ii) update of the construction layer, where the model mainly attempts to merge paths, and thus generalizes. The model implements a usage-based approach to generalize over observed examples and paths contained in the network. In particular, it exploits type variations that have a semantic implication to generalize observed *NL* sentences and (partially generalized) patterns to more abstract patterns. Consider the following example: A learner hears "Mia eats" and "Peter eats" in the above-mentioned visual context. To learn across situations, the model would use its knowledge that the linguistic phrase "Mia" refers to the semantic entity *mia* and that the phrase "Peter" refers to the semantic entity *peter* to bootstrap that the type variation in the sentences' first position ("Mia" vs. "Peter") reflects the meaning difference in the *AGENT* role of *eat*. The model would use its knowledge to acquire the more general pattern shown in (1), where SE_1 = [Mia → *mia*, Peter → *peter*].

(1)	Syntactic pattern	SE_1 eats
	Semantic frame	*eat(AGENT)*
	Mapping	$SE_1 \rightarrow AGENT$

Given an input *NL* sentence, the model finds a meaning by searching the network for corresponding paths/lexical nodes and ranking all possible meanings based on the weights stored in the associative networks. An *NL* sentence is parsed by first replacing units contained in groupings expressing syntactic slots (e.g. Mia) by the set (e.g. SE_1). Then, the model determines the semantic frame that corresponds to the path in the graph, if such a path exists. Finally, the model retrieves the meanings of lexical units at positions of syntactic slots from the lexical network. It uses the construction's mapping, i.e. the mapping specifying that SE_1 is the *AGENT*, to insert these meanings into the corresponding argument slots in the semantic frame. For details, please see Gaspers & Cimiano (in press). The important aspect to bear in mind is that the same Hebbian-style learning approach is used to train all layers of the network, in particular to learn the correspondences between linguistic and semantic units/nodes. Also note that the model incorporates a disambiguation bias (Merriman & Bowman, 1989) in the lexical subnetwork: Weights for new connections are initialized such that new lexical units are preferably associated with referents which have not yet been associated with other lexical units.

Extension to learning verb-general constructions

The model's extension relies on the same learning and representational devices in order to learn verb-general constructions that abstract form particular verbs and thus represent a generic signature of verbs, e.g. transitive, intransitive, ditransitive etc. verbs. It also detects cross-situational type variance and uses Hebbian-style reinforcement of edges to connect linguistic and semantic entities that are observed "together". To illustrate how the model learns verb-general constructions, consider the verb-specific construction given above in example (1) and the additional construction in (2), where SE_1 = [Mia $\rightarrow$ *mia*, Peter $\rightarrow$ *peter*]. From these verb-specific constructions, we would like to acquire the verb-general construction in (3), where SE_1 = [Mia $\rightarrow$ *mia*, Peter $\rightarrow$ *peter*], VE_1 = [eats $\rightarrow$ *eat*, sleeps $\rightarrow$ *sleep*], and VE_1 maps to the *ACTION* predicate.

(2)	Syntactic pattern	SE_1 sleeps
	Semantic frame	*sleep(AGENT)*
	Mapping	$SE_1 \rightarrow AGENT$

(3)	Syntactic pattern	$SE_1 \ VE_1$
	Semantic frame	*ACTION(AGENT)*
	Mapping	$SE_1 \rightarrow AGENT$

Since verbs map to action frames taking arguments expressing thematic roles of participants involved in the action, the model should also capture cross-situational statistics of verbs and syntactic frames, i.e. associate verbs with syntactic frames (and hence with possible argument structures).
In the network, we model sets of elements mapping to an *ACTION* predicate analogously to sets of elements expressing slots in syntactic patterns, namely, as nodes which group

sets of lexical elements. An additional associative network A_V captures the co-occurrence of the lexical units contained in these sets with syntactic patterns. A_V is included into the network structure and models associations between specific verbs and syntactic frames.
An additional third learning step is executed while processing input examples: The S&F layer is updated to yield verb-general constructions. In particular, the model searches for paths which show minimal variation in the surface structure and where exchanging elements at the position that varies yields a corresponding change in the associated meaning with respect to predicates. To illustrate the underlying intuition, consider the two verb-specific constructions in (4) and (5). These two examples can be merged into the verb-general construction shown in (6), assuming that "see" and "take" mean *see* and *take*, respectively.

(4)	Syntactic pattern	SE_1 sees SE_2
	Semantic frame	*see(AGENT,THEME)*
	Mapping	$SE_1 \rightarrow AGENT, SE_2 \rightarrow THEME$

(5)	Syntactic pattern	SE_1 takes SE_2
	Semantic frame	*take(AGENT,THEME)*
	Mapping	$SE_1 \rightarrow AGENT, SE_2 \rightarrow THEME$

(6)	Syntactic pattern	$SE_1 \ VE_1 \ SE_2$
	Semantic frame	*ACTION(AGENT,THEME)*
	Mapping	$SE_1 \rightarrow AGENT, SE_2 \rightarrow THEME$

More precisely, we define that two paths encoding syntactic constructions are mergeable if both differ in at most one position. Moreover, both paths must already have a learned (verb-specific or verb-general) meaning (see Gaspers & Cimiano, in press, for a definition) that includes the same mapping between sets of elements and thematic relations. Each element at a variable position must either have a learned meaning which corresponds to the predicate of the associated semantic frame or to a set of elements mapping to an *ACTION* predicate. The latter possibility allows the model to directly merge verb-specific with verb-general paths. All mergeable paths – if any – are then merged into a single path and elements at a variable position are replaced by a node that expresses a set of elements that maps to the *ACTION* predicate in an associated frame. Note that merging paths involves summing up weights in the corresponding associative networks (see Gaspers & Cimiano, in press).
Subsequently, A_V is updated, regardless of whether or not verb-general merging is possible. This update involves searching for paths with learned verb-general meanings which differ from the modified *NL* sentence only in the position expressing an *ACTION* predicate, e.g. a path p = "SE_1 VE_1 SE_2" in case of the modified *NL* sentence "SE_1 eats SE_2". The associative network A_V then learns that the lexical unit "eats" can occur within the syntactic pattern represented by the path (e.g., "eat" with path p). A meaning for an *NL* sentence that corresponds to a verb-general construction can be determined in the same way as for verb-specific constructions. Given a verb (lexical unit), an associated syntactic frame can be retrieved from A_V.

Experimental results and discussion

Our model requires initial linguistic knowledge to evaluate it with respect to the psycholinguistic studies with children. This section describes how input data were generated, and then presents how data from our experiments compare to the psycholinguistic findings.

Input data

Input data were generated similarly to Alishahi & Stevenson (2008) using the Eve corpus from the CHILDES database (Brown, 1973), which contains transcriptions of interactions with the child Eve. As input we used utterances spoken by Eve's mother. Since Arunachalam & Waxman (2010) and Scott & Fisher (2012) only consider transitive and intransitive structures (in one case including conjoined subjects), we extracted all patterns of the form "AGENT verb" and "AGENT verb THEME" from the corpus. We considered the same 13 verbs as Alishahi & Stevenson (2008). Since two of the verbs did not occur in the considered forms, we included the following 11 verbs in our experiments: *come*, *eat*, *fall*, *get*, *go*, *look*, *make*, *put*, *see*, *sit* and *take*.

The input generation lexicon contained all patterns along with their occurrence frequencies as well as the concrete nouns appearing at the AGENT and THEME positions of each verb along with their occurrence frequencies. Two nouns conjoined by "and" were also included. "Me", "you", and "we" were treated as "Eve", "Mom" and "Mom and Eve", respectively. We created *NL* examples from the input lexicon by randomly selecting patterns and referents according to their distribution in the lexicon/dataset. Semantic representations *mr* were created automatically using words appearing in generated *NL* sentences to denote the corresponding semantic referents. For example, "Mom sees" is represented as *see(AGENT:mom)*. Semantic referents are only arbitrary symbols to the model: It still needs to establish connections between words and referents. Two referents conjoined by "and" were treated as separate arguments having the same thematic relation. For instance, "mom and eve see" was represented as *see(AGENT1:mom,AGENT2:eve)*. In this paper, we do not address learning morphology. Hence, all words appear in their root form only. Ten different simulations containing 500 examples of the form (*NL,mr*) were generated and used for the following experiments. Presented results are averaged over the ten simulations. Model parameters were optimized on an independent data set.

Cross-situational verb learning

As mentioned above, Scott & Fisher (2012) investigated cross-situational verb learning and found that 2.5-year-old children can use cross-situational statistics to infer verb meanings under referential uncertainty, even if this requires abstraction across different actors and objects. This suggests that children can attach information about possible referents to novel verb entries along with their co-occurrence statistics and refine this information across trials. The study investigated learning of both transitive and intransitive verbs.

During each of the 12 experimental trials, children heard two intransitive (such as "she's pimming") or transitive (such as "she's pimming her toy") sentences, each containing a different novel verb, while watching two videos showing two different actors, each performing a novel action. In the transitive condition the action was performed with different objects. Children in the intransitive condition were significantly above chance in choosing the target actions over the distractor actions. Performance in the transitive condition depended on children's vocabulary size: Only children with large vocabularies performed significantly above chance.

We tested whether our model can infer meanings for novel verbs without receiving unambiguous label trials for any of the verbs. Thus, we tested whether the model can set up verb entries that contain information about possible referents and update co-occurrence frequencies over time. Note that since we used symbolic input, we cannot investigate the influence of abstraction over different actors and objects at the visual level. We used the same verbs, i.e. "pim", "nade", "rivv", and "tazz", and pairings of verbs as Scott & Fisher (2012). Referents for verbs were selected from the input generation lexicon (i.e., "mom" and "celery"). Since the model processes one sentence at a time, each input example contained one sentence and two possible *mr*s. For example, the first two intransitive input examples (which correspond to one trial in the study) were *NL:* "mom pim"; mr_1 *:pim(AGENT:mom)*; mr_2: *nade(AGENT:mom)* and *NL:* "mom nade"; mr_1 *:pim(AGENT:mom)*; mr_2: *nade(AGENT:mom)*. After receiving the examples, the model was asked to retrieve the semantic representations for the novel verbs, e.g. for "mom nade", and counted how often the model returned the correct representation, e.g. *nade(AGENT:mom)*. We performed the experiment with different numbers of examples observed prior to the experimental trials, corresponding to different "ages" of the model. Fig. 2 shows the results. In line with the children in the experi-

Figure 2: Proportion of the model's choice of the correct semantic representation for the novel verbs in the transitive and intransitive sentences.
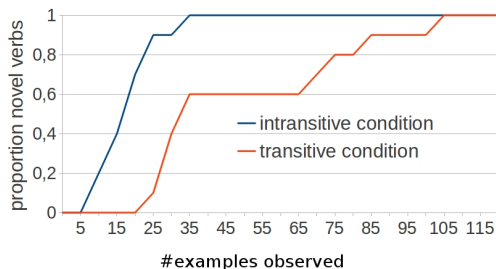
ments, the model can solve the task i) in both conditions from a certain "age" on and ii) earlier in the intransitive condition compared to the transitive condition, which was also more difficult for children. In particular, several children in Scott & Fisher (2012) failed in the transitive condition, while even 12-

to-14-month-old children typically master such a task when it involves mapping nouns to objects (Smith & Yu, 2008). This led Scott & Fisher (2012) to conclude that in cross-situational learning the same learning mechanisms may not apply uniformly for words of different categories. However, our model shows that applying the same learning mechanism for tracking co-occurrence statistics at different levels can yield behavior similar to that observed in psycholinguistic studies.

In our model, sentence-/verb-to-action mapping lags behind word-to-object mapping because it involves more complex structures whose acquisition depends on the prior acquisition of less complex structures, i.e. nouns. In particular, in order to establish a mapping for a verb "pim" in a sentence "mom pim celery", an *NL* pattern like "SE_1 pim SE_2" must previously have been derived, and "mom" and "celery" must be contained in the groupings SE_1 and SE_2, respectively. Furthermore, a necessary condition for deriving the pattern is that meanings for "mom" and "celery" must have been learned. Thus, similar to the children, the model's ability to solve the task depends on vocabulary, though not on the absolute vocabulary size, but rather on whether the meanings for the words observed at argument positions have already been learned (though, of course, the probability that the needed lexical units have already been learned may be higher for larger vocabularies).

The model learns faster in the intransitive compared to the transitive condition because it must have acquired only one word instead of two words for referents. In addition, patterns containing fewer groupings are in general learned earlier because the model generalizes based on type variation observed in one position. Notice, however, that we do not claim that children learn "mom" or "celery" at a specific age; these words were chosen arbitrarily for our experiments because they appear in our input data. Notice further that in contrast to the following experiment, the model can solve the above task even without the proposed extension.

Syntax as a zooming lens into semantics

As mentioned above, Arunachalam & Waxman (2010) showed that 27-month-old children can use the syntactic context to set up an initial verb entry and retrieve this entry when encountering the verb later on. The study comprised two training trials involving known verbs (to familiarize children with the task) and four experimental trials involving different novel verbs. During each trial of the study, toddlers first heard verbs presented within a dialogue without any relevant referent scenes. Each verb was presented eight times, either in transitive (e.g., "the lady mooped my brother") or conjoined-subject intransitive sentences (e.g., "the lady and my brother mooped"), without accompanying visual information. Toddlers then viewed two different scenes side-by-side depicting the same two participants: a synchronous scene and a causative scene. Toddlers were then instructed to find the scene that corresponded to the syntactic structure in which the verb had been presented. This instruction (e.g. "find moop") provided no syntactic information. Toddlers in the

transitive condition were significantly above chance in choosing the causative scene. In contrast, toddlers in the intransitive condition performed at chance level.

We tested the model in a similar manner both in a transitive and intransitive condition. Training trials were omitted since there was no need to make the model familiar with the task. Thus, there were four experimental trials, each featuring a different novel verb. Each verb was presented to the model eight times in either a transitive or subject-conjoined intransitive sentence (depending on the experimental condition). Referents for verbs were chosen from the input data, and the same referents were used in both conditions. The model was trained using these sentences (without accompanying *mr*s). Then, for each trial the model was asked to "find *new-verb*" in the presence of two *mr*s, a causative and a synchronous one. Since toddlers do not stop learning during experimental test periods, each test input (e.g. *NL:* "find moop"; mr_1: *moop(AGENT1:mom,AGENT2:eve)*; mr_2: *moop(AGENT:mom,THEME:eve)*) included a learning step. We then asked the model to retrieve the *mr* associated with the syntactic frame linked to the novel verb. Again, results were computed for different numbers of examples observed. Fig. 3 shows that, similarly to the children, at a certain "age" the model picks the causative scene (significantly) above chance in the transitive condition, but performs at chance level in the conjoined-subject intransitive condition. In line with the psycholinguistic results, the model can create an initial verb entry based on syntactic information alone, and it can retrieve this information when encountering the verb later on, if a suitable verb-general syntactic frame has been learned prior to the experimental trials. The model performs at chance level in the conjoined-subject condition because it learns the syntactic frame for this condition after that of the transitive condition: Further experiments revealed that with a greater number of input examples, the model also performs above chance in the conjoined-subject condition (i.e. it chooses the causative scene significantly less often). Thus, the model associates conjoined-subject intransitives with non-causal events at a later "age" than transitives with causal events, i.e. it learns the corresponding verb-general construction later. Similarly,

Figure 3: Proportion of the model's choice of the causative scene in the transitive and conjoined-subject intransitive conditions.
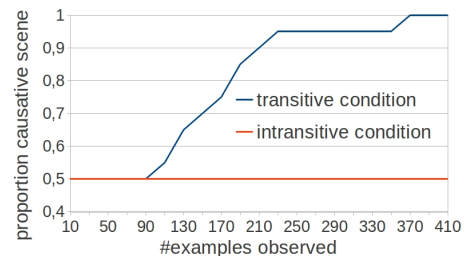
children do not succeed in the conjoined-subject intransitive task until the age of 3;4, while they can succeed in the tran-

sitive task at the age of 2 (Nobel, Rowland, & Pine, 2011). Since the model acquires both types of verb-general constructions in the same manner, the same proposed learning mechanisms can account for the different results. In the model, this result is due to the input data: The model acquires the conjoined-subject intransitive later because conjoined-subject intransitive sentences appear in the data much less frequently than transitive sentences.

General discussion and conclusions

We have presented a computational model for the acquisition of verb-general constructions under referential uncertainty. The model establishes form-meaning mappings under referential uncertainty by relying on cross-situational learning at different levels. Several models that acquire constructions have been proposed (see Gaspers & Cimiano, in press), including models that acquire verb-general constructions (e.g. Alishahi & Stevenson, 2008). However, most models assume that words or lexical mappings are already learned and/or do not address learning under referential uncertainty. However, such learning is relevant for the experiments simulated here since they address the acquisition of verb entries, including the establishment of lexical mappings under referential uncertainty. Several computational models can also use cross-situational learning to establish form-meaning mappings under referential uncertainty (e.g. Frank, Goodman, & Tenenbaum, 2007). However, these models have mainly focused on establishing mappings between words and referents. In contrast, our model applies the same cross-situational learning mechanism consistently to establish correspondences between form and meaning beyond simple word-referent mappings, in particular, between *NL* patterns/syntactic frames and actions, including thematic relations. Hence, our model can represent verb entries in the framework of these *NL* patterns, allowing it to store additional information about possible referents with verb entries.

We have presented empirical results that show how the model can establish verb meanings under referential uncertainty. Moreover, we have shown how the model can learn verb-general constructions, and how it can use this knowledge to create initial verb entries based on syntactic information alone. In line with usage-based approaches, the model's behavior depends on the input data, taking into account both token frequency and type variation. Overall, our results suggest that enough suitable input data in combination with the model's learning mechanisms can yield the behavior observed in children, and the model hence provides one possible formal explanation for the observed behavior. While several models that acquire constructions and/or word-to-meaning mappings have been proposed (see Gaspers & Cimiano, in press), we are not aware of any other model that describes all the specific learning mechanisms and representations that our model explores with respect to early verb learning. Experiments with children which test the model's predictions may establish whether or not children indeed apply learning mechanisms that are similar to those implemented in the model. For instance, cross-situational verb learning can be explored through more detailed analyses of children's vocabularies and by testing children with novel vs. known nouns as referents for verbs.

Acknowledgments

References

Alishahi, A., & Stevenson, S. (2008). A computational model of early argument structure acquisition. *Cognitive Science*, *32*(5), 789–834.

Arunachalam, S., & Waxman, S. R. (2010). Meaning from syntax: Evidence from 2-year-olds. *Cognition*, *114*, 442-446.

Behrens, H. (2009). Usage-based and emergentist approaches to language acquisition. *Linguistics*, *47*(2), 383–411.

Brown, R. (1973). *A first language: the early stages*. Harvard University Press, Cambridge MA.

Frank, M., Goodman, N. D., & Tenenbaum, J. B. (2007). A bayesian framework for cross-situational word-learning. In J. C. Platt, D. Koller, Y. Singer, & S. T. Roweis (Eds.), *Nips.* Curran Associates, Inc.

Gaspers, J., & Cimiano, P. (in press). A computational model for the item-based induction of construction networks. *Cognitive Science*.

Gleitman, L. R., & Fisher, C. (2005). Universal aspects of word learning. In J. A. McGilvray (Ed.), *The Cambridge companion to Chomsky.* Cambridge University Press.

Goldberg, A., & Suttle, L. (2010). Construction grammar. *Wiley Interdisciplinary Reviews: Cognitive Science*, *1*(4), 468–477.

Merriman, W., & Bowman, L. (1989). The mutual exclusivity bias in children's word learning. *Monographs of the Society for Research in Child Development*, *54*(3–4), 1–129.

Nobel, C. H., Rowland, C. F., & Pine, J. M. (2011). Comprehension of argument structure and semantic roles: Evidence from english-learning children and the forced-choice pointing paradigm. *Cognitive Science*, *35*, 963–982.

Quine, W. V. O. (1960). *Word and object*. MIT Press.

Rojas, R. (1993). *Theorie der neuronalen Netze*. Springer Verlag.

Scott, R. M., & Fisher, C. (2012). 2.5-year-olds use cross-situational consistency to learn verbs under referential uncertainty. *Cognition*, *122*(2), 163–180.

Siskind, J. M. (1996). A computational study of cross-situational techniques for learning word-to-meaning mappings. *Cognition*, *61*(1–2), 39-91.

Smith, L., & Yu, C. (2008). Infants rapidly learn word-referent mappings via cross-situational statistics. *Cognition*, *106*(3), 1558–1568.