

Barend Beekhuizen, Rens Bod, and Arie Verhagen

3 Acquiring relational meaning from the situational context: What linguists can learn from analyzing videotaped interaction

1 The difficulties of acquiring relational meaning¹

The acquisition of linguistic expressions of relational meaning is typically regarded as a difficult task (Gentner 1978; Gleitman 1990).² What hampers this process, as compared to the acquisition of expressions of object reference, is not so much the complexity of relational concepts per se, which children seem to understand well before their first birthdays for many types of relational concepts (Casasola, Bhagwat, and Ferguson 2006; Mandler 2006; Sootman-Buresh, Woodward, and Brune 2006), but the question which subset of the large set of relations perceivable in situational context should be mapped to a linguistic expression. Furthermore, as especially Gleitman (1990) argues, not all relations expressed are directly perceivable, as some pertain to mental states of intentional agents and others are simply not present in the here-and-now of the speech situation. As Gleitman succinctly puts it (p. 5): “[T]here is *not enough* information in the whole world to learn the meaning of even simple verbs, or . . . there is *too much* information in the world to learn the meaning of . . . verbs.”

In this chapter, we evaluate Gleitman’s double problem. Is the context indeed too poor or too rich for the learner to find situational, contextual correlates for the linguistic items he perceives that can help him bootstrap the meaning of

1 Corresponding author: Barend Beekhuizen, Leiden University. Email address: barendbeekhuizen@gmail.com. The work reported here is part of the first author’s Ph.D. project *Constructions Emerging*, funded by NWO. The second and third authors have made significant contributions to the development of the described method. The authors would like to thank Marian Bakermans and Rien van IJzendoorn of the department of Child Studies at Leiden University for allowing us to use their data, as well as the two anonymous reviewers for their valuable comments.

2 As a working definition of *relational*, we adopt Gentner and Kurtz’s (2005: 153) explanation that “[e]ntity categories can be thought of as first-order partitions of the world (Gentner 1982) and relational categories second-order ways of organizing and linking those first-order partitions”.

Barend Beekhuizen, University of Toronto

Rens Bod, University of Amsterdam

Arie Verhagen, Leiden University

those items? We address this question for a broad class of word types, both relational (verbs and prepositions) and non-relational (nouns and adjectives) ones.³ We do so by using a corpus of child-directed language that contains precise information on the behavior of the participants and several states in the world. This information is derived from annotations of the behavior of the caregiver and the child, based on videotaped interaction in the setting of a simple game.

Note that our focus is mainly on the situation rather than on cognition. Work by Gentner and colleagues on relational categories (see e.g. Gentner and Kurtz 2005) has provided us with much insight in the cognitive mechanisms at work when acquiring relational meanings of words. This chapter looks at the other side of the same coin: what information is available in naturalistic situations on which these mechanisms operate. Is there any relation between two objects that could be construed as “containment” present in the environment in which the language-learning child operates when the word *in* is uttered?

We believe the development and use of methods like these, in combination with computational modeling techniques, is paramount for the development of a more comprehensive usage-based theory. If we assume that the interaction between caregivers and children is the locus of the development of the acquisition of things like construal and symbolic mappings, we must at least have a decent understanding of what happens in those situations. Without analyzing actual behavioral data, we will not know what the information available is, and any form of theorizing about the mechanisms used in the acquisition of symbolic mappings can effectively not be evaluated, as the interpretive step from a controlled experiment tapping into the hypothesized mechanisms and the situation “in the wild” is unwarranted. It is exactly this understanding that is currently too weakly present in acquisitionist research.

This precise method has, to the best of our knowledge, not been used before in language acquisition research, but similar work has been done on object reference using video data in which the present objects and social cues such as eye gaze and pointing were coded (Frank, Goodman, and Tenenbaum 2009; Frank, Tenenbaum, and Fernald 2013). Because of this, we discuss the considerations in developing the data in some detail first (Section 2). In Section 3 we then address the question of the usefulness of the context for acquiring relational meaning by looking at a quantitative measure of association between linguistic elements and aspects of the situation in which these elements are produced. We

³ We realize that it is the case that there are verbs that have non-relational meanings, as well as nouns and adjectives that have relational ones. The particular items we study do not belong to these types.

limit ourselves to verbs, and other relational elements that encode conceptualizations of physical relations and actions.⁴

2 Developing a corpus of contextualized language

If we want to understand how informative the situations are for the language-learning child, we need to have material that captures the situational context of every utterance in detail and that is situated in a relatively ecologically valid setting. Regular corpora contain little information on the context, and deriving information about the context from the language in order to find associations between the two, which is often done in computational modeling studies of word-meaning acquisition, leads to circularity.

Because of this lack of data, we decided to develop a dataset on the basis of videotaped interactions in Dutch between children and their mothers, which was recorded for other purposes. The child-directed language was transcribed by the first author, and the situation was described according to a formalized coding scheme by two assistants. In the following sections, we discuss the choices and the justification for them, as well as the coding procedure and evaluation in some detail, so as to provide other researchers with a starting point for doing similar research.

2.1 The nature of the underlying data

Our material consists of 32 videotaped fragments, each containing a unique mother-daughter dyad playing a full game of putting blocks in holes of a toy. The mothers are all middle-class, native speakers of Dutch who were living in or close to the Dutch town of Leiden at the time. The fragments are on average 4 minutes and 54 seconds long, ranging from 3 minutes and 12 seconds to 7 minutes and 5 seconds. The game starts with the observer giving the toy to the mother, and ends with the mother returning the toy to the observer. The observer

⁴ This is not to say that verbs encoding mental states, intentions and dispositions have no ground in the child's situational understanding independently of language. From a very young age, children have a basic understanding of intentions and emotions (Behrend and Scofield 2006; Tomasello 2003) that might become associated with certain verbs. Barak, Fazly, and Stevenson (2012) provide an interesting first step towards modelling the acquisition of such meanings.




Figure 1: The toy and the twelve blocks.

does not take part in the interaction, so the mother and the daughter are jointly focused on playing the game.

The game-related objects are displayed in Figure 1. The game typically takes place on the table or on the floor, and there are few non-game objects that are manipulated during the game. A typical game consists of the mother or daughter opening the bucket and getting all blocks out. Then the daughter tries to fit each of them back in, sometimes with the help of the mother, but always with verbal comments, suggestions and questions by the mother. Several children lose interest after a while and wander off, while others are so excited about their success that the game is played two or three times during the fragment.

All children were around sixteen months at the time of the recording. This means that their ages are within a desirable time window for this type of observational study. Before the age of one, we cannot expect children to understand the symbolic nature of language. When children's early intentional understanding enters the picture, they are able to start learning linguistic symbols (Tomasello 2003: 21–28). However, as we get closer to the child's second birthday, chances increase of the child being able to talk it-self, in which case the child is often directing the conversation by means of single-word utterances triggering responses from the caregiver. The age of sixteen months squarely falls within the period in which children understand the symbolic, communicative nature of language, but are just beginning to connect word forms and meanings. Because

of this, we can make the assumption that this is actually the kind of situated child-directed language input that children receive.

2.2 Developing a coding scheme

What aspects of the situational context are worth taking into account? The participants perform various actions related to the game, and hence the spatial states of the game objects change continuously. The most straightforward aspects of the scenes we could describe thus are changes in spatial relations between objects and the behavior of participants leading to these changes. We found that with nine predicates, we were able to cover almost all of the participants' behavior in the fragments. The top nine predicates, given in Table 1 and described as English verbs, reflect the building blocks of object manipulation: there are descriptions of bringing an object under manual control (**reach** and **grab**), and of letting go of this control (**let go**). In between, one can do all sorts of things with objects: **move** them from one location to another, position them vis-à-vis a location, exert **force** upon them and letting them go (as in pushing, throwing) and showing them to another participant. Furthermore, we can **point** to non-held objects. Situations that do not fit into this set of categories are assigned the label **other**, and in case of doubt between two existing codes, the label **unclear** could be used. Finally, when one of the participants' behavior is not visible, because that participant is outside of the camera frame or behind some occluding object, the predicate **out-of-view** is assigned. For a more detailed description of these predicates, we point to the coding manual.⁵

Each of these predicates dictates a number of roles together with which they form a semantic predicate-argument structure. The predicate **grab**, meaning to bring something under one's manual control, has a grabber and a grabbed object, and possibly an instrument of grabbing other than the hands (a spoon, the mouth), **move** has a mover, a moved object and a source and goal location. The second column of Table 1 describes the valency of the predicates. As to the fillers of these roles, all roles marked with subscript *O* require objects to fill them, whereas the *S* subscripts describe a role to be filled with a spatial predicate. The objects either come from a closed class of descriptions of game-related objects (Table 2) or are assigned freely but consequently by the coder, in the case of objects that are not part of this closed set.

⁵ Available at https://github.com/dnr/b/publications/blob/master/annotation_guidelines.pdf

Table 1: Behavioral predicates and their roles.

predicate	arguments
reach	(theme _{o/s} , [instrument _o])
grab	(theme _o , [instrument _o])
show	(theme _o , recipient _o , [instrument _o])
position	(theme _o , location _s , [instrument _o])
move	(theme _o , source _s , goal _s , [instrument _s])
hit	(theme _o , [instrument _o])
let_go	(theme _o , [instrument _o])
force	(theme _o , source _s , goal _{o/s} [instrument _o])
point	(theme _o , recipient _o , [instrument _o])
other	n.a.
unclear	n.a.
out-of-view	n.a.

Table 2: Object labels for game-related arguments.*

predicate	arguments
to	toy
bu	bucket
ha	handle
li	lid
ho (ci sq st tr) ₁	hole in the lid with a certain shape ₁
b (bl re gr ye) ₁ (ci sq st tr) ₂	block with a certain color ₁ and shape ₂
mo	mother
ch	child
ob	observer
table	table
floor	floor
air	air

* ci = circle, sq = square, st = star, tr = triangle, bl = blue, re = red, gr = green, and ye = yellow.

Independent spatial predicates are used in the case of some arguments (sources, goals, locations) and in order to describe salient events of the game. The spatial relations are given as English prepositions, but, again, rely upon descriptions that the coders used when coding the states. The four relations that are coded, are **in** for containment, **on** for horizontal support, **at** for all other forms of physical contact between two objects and **near** for all forms of non-contact but a salient proximity between two objects. Two pairs of game-related objects are independently described because of their relation to the success of the participants. The first is the lid being on or not-on (**off**) the bucket, and the second

that of each of the blocks being in or not-in (**out**) the bucket. These relations are only coded if there is a change in situation.

When developing a set of predicates and objects like this, two important questions arise. First: does the handcrafted set have the right degree of granularity? Second, and perhaps more important: does the set reflect perceivable relations and not leave out perceivable relations. Regarding the first question, the set of primitives was developed so as to be as specific as possible, while still being codable by an instructed coder. In evaluating the coding procedure, there are some demarcation issues between predicates that can be solved by decomposing them further, but doing so would involve methods beyond coding video data to obtain the information.

As to the question whether the coding scheme is not too specific: our aim was to code primarily the physical behavior, and as little intention as possible. Actions at a larger scale, such as getting something out of something, often consist of parts that are simple predicates like **grab** and **move**, but are as a whole intentional, goal-directed actions. Inferring these from the perceived simple action schemas is a future direction we are currently looking into (see also Section 2.5).

Turning to the perceptual availability of the predicates, we believe the behavioral predicates reflect very basic, primitive sensory-motor experiences that are available from at least 12 months of age (Meltzoff and Moore 1995). These may then be grouped into culture-specific gestalts, combining properties of the shape of the manipulated object, the presence and nature of instruments etc. (see for an approach to the way languages encode these types of features Narasimhan and Kopecka 2012), but as such, we expect the basic building blocks of these analytically complex concepts to be perceptually available independent of culture.

The same holds, *mutatis mutandis*, for the spatial relations and properties and category labels of objects. Especially in the case of the former, the work by Baillargeon and Wang (2002) shows that notions like ‘containment’, ‘nearness’ and ‘physical contact’ have developed well before the child’s first birthday. The predicate **on**, denoting horizontal support, is arguably more problematic. Choi (2006) presents evidence that children understand support relations later than containment relations, a fact that she ascribes to the many ways in which there can be a support relation, such that a common conceptual core is harder to abstract. Also, some potential perceptual sources of spatial relations that are encoded in different languages, such as pointwise vs. surface contact, tight fit vs. loose fit, are not coded in the data. Again, decomposing spatial relations into more primitive notions of relations (e.g. contact, surroundedness of figure by ground on both the horizontal and vertical axes, similarity in shape) combined

with properties of the figure and the ground (mass-like or discrete, animacy) can provide a promising avenue (see e.g. Feist 2000 for such an approach), but coding observational data with such fine-grained distinctions (as opposed to setting up experimental situations) might prove too difficult for coders. Summarizing: the coded predicates are all available, but might not reflect all perceptually available relations. Nevertheless, we believe that we can still address the question whether the context is not too rich, as we code perhaps not all, but at least many predicates that function as noise for learning a word.

2.3 The coding procedure

The predicate-argument structures described above were coded using ELAN, a piece of software for video annotation.⁶ Regular intervals of three seconds were created, so that the coders had to describe the events that took place within that window of time. Predicates taking place in two subsequent windows were only coded in the window in which they started. The behavior of the mother and the child was coded on different tiers, as were the spatial states of the game objects.

The coders, a second-year and a fourth-year student at Leiden University, were paid as research assistants and worked in total 56 and 64 hours respectively on the project and were supervised by the first author. They received six hours of training. The first two hours were spent on an in-depth explanation of the predicates, arguments and objects, after which the coders and the first author jointly coded a fragment in four hours. After this, the coders reported to feel confident about the task. Apart from the training, they received a manual with a reference sheet to be used during the coding itself. This document contained detailed descriptions about the predicates, decision trees for anticipated demarcation problems and a general instruction about the workflow using ELAN. The coders and the supervisor had contact at every working day about unclear cases and more general issues.

2.4 Coding the behavioral data: insights and figures

The coders worked for 56 and 64 hours on coding fragments. In this time, they coded 23 and 19 fragments, respectively. Each coder did one fragment twice and there was an overlap of three fragments between the coders. Five fragments

⁶ Developed by The Language Archive at the Max Planck Institute for Psycholinguistics, Nijmegen, and publicly available from <http://tla.mpi.nl/tools/tla-tools/elan/>. For a description, see Brugman and Russel (2004).

were discarded due to a low visibility of the actions (three cases) or children who failed to play the game for most of the time (two cases). On average, coding a fragment took the coders two hours and thirty-four minutes, with barely any difference between the coders (2 hours and 31 minutes, and 2 hours and 38 minutes respectively). This resulted in 175 minutes of coded material, out of which 157 minutes were useful, unique fragments. The average rate of coding thus was one minute of video per approximately 37 minutes of coding time.

The coders reported few problems with the procedure. Recurrent difficulties were the difference between **force** (i.e. moving without grabbing) and **hit**, which indeed is a distinction along a continuum, viz. that of the duration of the impact, which is more pointwise in hitting and more durative in forcing. This can be overcome if the dimensions of the manual contact are split out: one has to decide upon a cut-off point for the length of the impact to demarcate **force** from **hit**.

Also, the difference between **position** and **move** was felt to be unclear at times. Again, these predicates seem to form a continuum, with the amount of motion with respect to the ground being low for the former and high for the latter predicate. Truly distinguishing them would require the interpretation of intentions of the agents: one positions something in order to establish a different relation between the undergoing object and the location, whereas one moves something to change the object's location from a source to a goal location.

In order to evaluate the reliability of the coding, we can calculate the inter- and intra-coder agreement. Three games were coded by both coders and each coder did one fragment twice. We measure the agreement by checking, for each predicate and for each time frame of three seconds whether at least one instance of that predicate was coded. A good measure for agreement is Cohen's kappa (Cohen 1960), which looks at the extent to which two codings of the same situation deviate (either two codings of one situation by the same coder or two codings of one situation by two different coders). A κ of 1 means perfect agreement, and a κ of 0 complete disagreement. This measure is often used in studies using human coders (e.g. Carletta 1996). Table 3 lists the Cohen's kappa scores for each predicate. A $\kappa \geq .80$ is generally taken to be the standard for very reliable coding (though often lower scores, above $\kappa \geq .60$ are taken to be acceptable as well).

We can see that several predicates are coded very reliably: **show**, **move**, **let go**, **point** and **grab**. The figures for **hit** are reasonably good, but the predicate is very infrequent, and all scores of 1.00 for this predicate are due to single predicates. **Reach** was coded unreliable across coders, but reliably within. This was due to the fact that one coder regarded them to be 'implied' by other predicates

(**grab**, **hit** and sometimes **force**) and only coded **reach** when these other predicates did not apply.

Table 3: Cohen's kappa per fragment and coder.

predicate	agent	Inter-coder agreement per fragment			Intra-coder agreement	
		1	2	3	coder 1	coder 2
reach	mother	.00	.29	.50	.80	.82
	daughter	.57	.35	.75	.66	.78
show	mother	1.00	1.00	.91	.91	1.00
	daughter	.80	1.00	.91	.91	1.00
move	mother	.77	.92	.85	.81	.92
	daughter	.79	.93	.86	.85	.91
position	mother	.70	.00	.69	.80	.00
	daughter	.59	.38	.69	.67	.39
hit	mother	.00	1.00	1.00	1.00	1.00
	daughter	.66	1.00	1.00	.75	.00
let_go	mother	.96	.84	.89	.90	.88
	daughter	.85	.83	.79	.92	.84
force	mother	1.00	1.00	1.00	1.00	1.00
	daughter	.00	.75	1.00	.27	.00
point	mother	.56	.89	1.00	.89	.89
	daughter	.59	.89	1.00	.90	.89
grab	mother	1.00	.90	.87	.81	.83
	daughter	.85	.82	.85	.85	.87
out		.00	1.00	1.00	1.00	.79
off		1.00	.83	1.00	1.00	1.00
in		.49	.92	1.00	.91	.87
on		.83	.91	.49	1.00	1.00

Position, finally, was coded variably between coders and within coders. This may have to do with the fact that, as the coders reported, position and move were found difficult to discriminate. How to create a more transparent coding scheme for the motion-positioning continuum remains to be seen. For now, we will retain the predicate, but in future work, this action will need to be demarcated clearer for the coders. The space predicates, finally, were coded reliably overall, except for two low scores, one for **out** and one for **in**. In the first case, one coder coded **out** in the time window after the one in which the other coder coded it, so effectively, they agree upon the predicate, but placed in a different time frame. Part of the explanation of the low score for **in** is the disagreement on positioning: one coder used more **position** labels (as opposed to non-actions) and in many cases the positioning took place in the air.

2.5 Deriving information

The codings resulting from this process were paired with all utterances that take place within each time window, to give us a corpus of child-directed language enriched with detailed situational information. Table 4 gives us a sample of this corpus.

Table 4: A sample of the coding and transcribed language paired.

time	type	action/utterance*
0'00	situation	hold(mo,li)
	language	een. nou jij een.
	gloss	one. now you one.
0'03	situation	position(mo,to,on-floor) grab(ch,byetr) move(ch,byetr,on-floor,near-hoci)
	language	nee daar.
	gloss	no there.
0'06	situation	point(mo,hotr,ch) hold(mo,li) position(ch,byetr,near-hoci)
	language	nee lieverd hier past ie niet.
	gloss	no sweetheart here fits he not.
0'09	situation	point(mo,hotr,ch) letgo(mo,lid) move(mo,byetr,near-hoci,near-hotr) grab(ch,bblst) move(ch,bblst,on-floor,in-air) grab(mo,byetr) letgo(ch,byetr)
	language	hier in. kijk e(en)s. een twee.
	gloss	here in. look once. one two.

* to = toy, li = lid, ho = hole, b = block, bl = blue, ye = yellow, st = star, ci = circle, tr = triangle, mo = mother, ch = child.

Given the fine-grained descriptions of the behavior of the participants, it is possible to derive more information about the situation without having to code it manually. An important aspect of understanding a situation is understanding the intentions and goals of the participants. Although intentions are not directly perceivable, they can be inferred from the actions the participants partake in (cf. Fleischman and Roy 2005). It is likely that a person grabbing a block has some goals involving that block. Suppose we assume that the participants understand the global outline of the game, we could say that one participant's grabbing of a block can make the other participant infer that the first participant has (or: should have) the goal of putting that block in the bucket through the hole. We are currently exploring this direction further.

On a lower, less intentional level, it is interesting to know whether there is a match in shape or not between a block and a hole when a participant moves and positions a block next to a hole. Furthermore, the blocks can be split out

in such a way that **grab (mother, blue-triangular-block)** is given as **grab(e_i, e_j)**, **blue(e_j)**, **triangular(e_j)**, **block(e_j)**, **mother(e_i)**.

2.6 The language

The child-directed language used in the 157 minutes of useable material consists of 7,842 word tokens in 2,492 utterances, on average 3.14 words per utterance. Many utterances consist of merely ja ‘yes’ or goed zo! ‘well done!’. The word tokens contain 480 types and 355 word lemmas (i.e. lemmas without inflectional and diminutive morphological marking).⁷ Table 5 gives an overview of lemma statistics split over some categories. Nouns and adjectives are lumped, because many adjectives are used nominally. We can see many game-related lemmas in the top of the frequency distributions of the verbs, nouns and adjectives, and prepositions and spatial particles.

3 A first exploration

3.1 Data

To see if the situational context can be informative for the child trying to assign a meaning to a word, we investigate the association between coded aspects of the situation and certain lemmas. If we can find a strong association between certain lemmas and aspects of the situation related to their meaning, we can establish that the situational context is rich enough and not too rich to learn the meaning, or at least point the child in the right direction.

The starting point is a corpus of pairs of utterances (U) and situations (S), based on the 157 minutes of coded and transcribed interactional data discussed in the previous section. An utterance consists of the string of all lemmas $w_1 \dots w_n$, and the situations are represented as a set of features $f_1 \dots f_n$ that are present within the three-second time window in which the utterance takes place. So, every datapoint d is a pair U, S , where $U = w_1 \dots w_n$ and $S = \{f_1 \dots f_n\}$.

⁷ We take a type to be a unique transcribed representation of the speech, including diminutives and different sorts of inflection. A lemma, on the other hand, is taken to be the core word (whether it is a compound, another derivation or a mono-morphemic element). Examples of different types belonging to the same lemma are *ster* ‘star’, *sterren*, ‘star.PL’ and *sterretje*, ‘star.DIM’, that are all types related to the lemma *ster* ‘star’. No methodological problems were found in determining the types and lemmas.

Table 5: Lemma frequencies split out over part of speech type and ranked, and total number of lemmas and tokens per part of speech type.*

verbs		nouns/ adjectives		prepositions/ spatial particles		other	
lemma	<i>n</i>	lemma	<i>n</i>	lemma	<i>n</i>	lemma	<i>n</i>
<i>doen</i> 'do'	260	<i>goed</i> 'good'	197	<i>in</i> 'in'	287	<i>ja</i> 'yes'	361
<i>gaan</i> 'go'	205	<i>mooi</i> 'beautiful'	58	<i>op</i> 'on'	54	<i>die</i> 'that'	306
<i>zijn</i> 'be'	185	<i>ander</i> 'other'	49	<i>uit</i> 'out'	43	<i>zo</i> 'like that'	288
<i>kijken</i> 'look'	178	<i>moeilijk</i> 'difficult'	46	<i>met</i> 'with'	25	<i>maar</i> 'but'	259
<i>moeten</i> 'should'	127	<i>mama</i> 'mom'	36	<i>af</i> 'off'	15	<i>d'r</i> 'there (red.)'	191
<i>proberen</i> 'try'	70	<i>keer</i> 'time'	26	<i>door</i> 'through'	15	<i>hier</i> 'here'	168
<i>komen</i> 'come'	60	<i>deksel</i> 'lid'	25	<i>van</i> 'of'	10	<i>eens</i> 'PRT'	162
<i>kunnen</i> 'can'	51	<i>blok</i> 'block'	24	<i>voor</i> 'for'	10	<i>nee</i> 'no'	158
<i>pakken</i> 'grab'	40	<i>rond</i> 'round'	23	<i>aan</i> 'on, at'	9	<i>een</i> 'a'	147
<i>passen</i> 'fit'	38	<i>leuk</i> 'nice'	19	<i>naar</i> 'towards'	8	<i>je</i> 'you (red.)'	143
<i>zullen</i> 'shall'	35	<i>ster</i> 'star'	17	<i>bij</i> 'at'	6	<i>hij</i> 'he'	138
<i>zitten</i> 'sit'	33	<i>makkelijk</i> 'easy'	15	<i>om</i> 'in order to'	4	<i>dat</i> 'that'	132
<i>halen</i> 'get'	32	<i>puzzel</i> 'puzzle'	13	<i>na</i> 'after'	3	<i>he</i> 'huh'	115
<i>hebben</i> 'have'	30	<i>gat</i> 'hole'	11	<i>neer</i> 'down'	5	<i>nog</i> 'still'	112
<i>maken</i> 'make'	28	<i>spel</i> 'game'	11	<i>heen</i> 'ither'	5	<i>niet</i> 'not'	106
<i>draaien</i> 'turn'	26	<i>stuk</i> 'piece'	10	<i>naast</i> 'next to'	2	<i>we</i> 'we (red.)'	105
<i>stoppen</i> 'stick'	18	<i>vierkant</i> 'square'	9	<i>achter</i> 'behind'	1	<i>nou</i> 'PRT'	98
<i>zien</i> 'see'	13	<i>emmer</i> 'bucket'	8	<i>boven</i> 'above'	1	<i>'el'en</i> 'one'	96
<i>weten</i> 'know'	12	<i>meid</i> 'girl'	8	<i>buiten</i> 'outside'	1	<i>daar</i> 'there'	86
<i>willen</i> 'want'	12	<i>groen</i> 'green'	7	<i>dichtbij</i> 'close'	1	<i>deze</i> 'this'	85
74 lemmas		103 lemmas		20 lemmas		158 lemmas	
1640 tokens		765 tokens		505 tokens		4932 tokens	

* The gloss PRT is given when the word primarily function as a discourse particle for which no translation equivalent exists in English. (red.) means that this is a conventional, reduced form that is morpho-syntactically distinct from the full form.

For this exploration, we simplified the representation of the situation. Whereas we coded the information as predicate-argument structures, we discarded all arguments and used the predicates as atomic features. The predicate-argument structure **grab** (**mother**, **bucket**) and **on** (**lid**, **bucket**) thus become **grab** and **on** respectively. Apart from the behavioral predicates and spatial predicates that take place at a certain point in time, the following features are also included:

- properties of blocks and holes present among the arguments of the predicates coded at that moment in time (e.g. **red**, **star**);
- category labels for all objects that are arguments of those predicates (e.g., **mother**, **block**, **table**);
- two derived predicates: **match** and **mismatch**, that apply if there is a coded spatial relation between a block and a hole in the time window and there is either a match in the shape of both or a lack thereof, respectively.

In total, this procedure gives us 89 feature types, with an average of 12.2 features per situation.

To give a fuller example, consider the utterance in (1), which is paired with the situation in (2).

- (1) *allemaal even d'r uit halen*
 all PARTICLE there out remove.INFINITIVE
 ‘Let’s get them all out.’
- (2) reach (mother, bucket) grab (mother, bucket) move (mother, bucket,
 on-floor, in-air) position (mother, bucket, in-air) out (blocks, bucket)

From the predicate-argument structures, we extract for the purpose of this study the features listed above. This gives us the situation S given in example (4). The lemmas of the words in the utterance in (1) are given in (3).

- (3) $U = \{\text{allemaal, even, d'r, uit, halen, gewoon, zo}\}$
- (4) $S = \{\text{reach, grab, move, position, air, block, bucket, floor, mother, blue, green yellow, red, in, on, out, round, square, star, triangle}\}$

3.2 Associating language and situation

Suppose we are interested in seeing whether the lemma *ster* ‘star’ is associated with the (non-relational) feature **star**. We check in our dataset of utterance-situation pairs for every sentence whether the lemma *ster* is used in the utterance ($ster \in U$) and whether the feature **star** is present in the situation ($\mathbf{star} \in S$). This gives us the 2×2 table given in Table 6.

Table 6: 2×2 table for the lemma *ster* and the feature **star**.

	$\mathbf{star} \in S$	$\mathbf{star} \notin S$	row sum
$ster \in U$	16	1	17
$ster \notin U$	748	1727	2475
column sum	764	1728	2492

What we can see here is that the lemma *ster* is used almost exclusively in situations in which the feature **star** is present, that is: in situations in which a star-shaped object is an argument to an action or space predicate taking place at that time. However, in an overwhelming majority of cases in which the feature **star** is

present, the lemma *ster* is not used. Despite this fact, we can see that the lemma and the feature are clearly associated: there is a higher chance of finding the feature if the lemma is present than if the lemma is not present (94.1% vs. 30.2%) and a slightly higher chance of finding the lemma if the feature is present than if it is not present (2.1% vs. 0.1%).

This association is reflected in a highly significant association using the Fisher exact test ($p = 6.6399e-8$). Note that the association between lemmas and features resembles that between syntactic environments and words in certain slots in collocation analysis (Stefanowitsch and Gries 2003). Where Stefanowitsch and Gries establish behavioral profiles for constructions, we use the same statistics to show what the situational profiles are for lemmas. Following Stefanowitsch and Gries (2005), we take the negative natural logarithm of the Fisher exact p -value as a more readily interpretable association measure. This effectively means that the lower the probability of this and more extreme distributions, the higher the association value. For the feature **star** and the lemma *ster*, the association value is $-\ln 6.6399e-8 = 16.53$. Note that, as we use the natural logarithm, association values greater than 2.995 and 4.605 reflect associations significant at $\alpha = .05$, resp. $\alpha = .01$.

Importantly, the Fisher exact test and derived association value will inform us about both poverty and excessive richness of the context. If the intended meaning is not grounded in some perceivable feature (i.e. the context is too poor) the lemma will be associated with other features that make little sense from an adult point of view. On the other hand, if there are too many features consequently present in the situational context of a lemma, the lemma will be associated about equally with all of them. A high association between a lemma and a feature thus means that that feature is found more often with that lemma than with other lemmas and that that lemma is found more often with that feature than with other features.

Now, we can calculate the association values between the lemma at hand (i.e. *ster*) and all possible features. Doing so and sorting the features on their association values with *ster*, from high to low, we can determine the situational profile for the lemma (Table 7).

What we see is that the lemma is most strongly associated with what we conceive to be its meaning, namely the feature **star**. Apart from that, we find noise, with the second strongest association being between the lemma and the feature **air**. We can repeat this procedure for a range of lemmas. Table 8 gives the top three strongest associated features for the four words referring to the colors of the blocks, and Table 9 does so for the words referring to the shape of the blocks.

What we can see from these tables, is that in five out of eight cases, the arguably correct feature is the one with the strongest association with the lemma. The cases in which the correct feature is not associated with the lemma are either very low in frequency (*blauw* ($r = 3$), *driehoek* ($r = 4$)), or have an ‘unfortunate’ split in time windows (three out of nine cases of *groen*). This means that the child grabbed a green block, and a second later, just within the next time window the mother comments on the block, using the lemma *groen*. Because the child was not *doing* anything with the block in the time window in which the utterance was produced, there is no predicate in which it is involved and hence **green** is not included in *S*.

Table 7: Features with the highest association values with the lemma *star*.

feature	association value
star	16.53
air	7.35
in	6.04
move	3.59
ball	3.39

Table 8: Lemmas for color and features most strongly associated with them.

<i>rood</i> ‘red’		<i>groen</i> ‘green’		<i>blauw</i> ‘blue’		<i>geel</i> ‘yellow’	
feature	assoc.	feature	assoc.	feature	assoc.	feature	assoc.
red	3.99	move	3.11	circle	3.43	yellow	4.49
let_go	2.06	lid	2.79	out	2.53	let_go	1.65
in	1.91	green	2.6	mother	1.91	mother	1.45

Table 9: Lemmas for shape and features most strongly associated with them.

<i>vierkant</i> ‘square’		<i>driehoek</i> ‘triangle’		<i>rond</i> ‘round’		<i>ster</i> ‘star’	
feature	assoc.	feature	assoc.	feature	assoc.	feature	assoc.
square	5.17	yellow	3.23	round	19.35	star	16.53
force	4.42	point	2.39	floor	6.25	air	7.35
match	3.44	mother	1.65	mismatch	4.26	in	6.04

3.3 Relational terms and their situational contexts

In this section, we investigate to what extent the situation can guide a learner towards the meaning of relational terms. We consider eight lemmas that relate to aspects of the game. Linguistic items that do relate to the limited world of the game are action verbs, spatial prepositions and terms referring to properties of the blocks and holes (colors and shapes, as we have discussed in the previous section). Let us turn to the former two categories now.

The utterances contain many verbs pertaining to game-related actions. We find verbs of placement, like *stoppen* ‘put in’, and *halen* ‘remove, get’. Then there are verbs of manipulation, such as *duwen* ‘push’, *pakken* ‘grab’ and *draaien* ‘turn, rotate’. Finally, of interest are verbs denoting relations between blocks and holes, such as *horen* ‘belong’ and *passen* ‘fit’. For all of these relations, a range of modals and light verbs is also used, especially in combination with certain argument-structure constructions and spatial particles. We will not go into these in this exploration. Here, we discuss four verbs, namely *halen*, *pakken*, *passen* and *proberen* ‘try’. The former three are members of the categories of placement verbs, manipulation verbs and verbs pertaining to stative relations. The latter is a frequently used verb in the game situation but has no clear correlate in the coded predicates. We discuss it to see what sort of situational profile it displays.

The prepositions and verbal particles also form an interesting category. The caregivers often talk about the spatial relations of the game objects, especially in combination with the three types of verbs discussed above (someone pushing something in something, someone getting something off something, something fitting in something). Four prepositions we will discuss are *in* ‘in’, *op* ‘on’, *uit* ‘out’ and *af* ‘off’.

What are the situational profiles for the eight items we selected? Tables 10 and 11 present the top five most strongly associated features per lemma.

Table 10: Four verbs and the features most strongly associated with them.

<i>passen</i> ‘fit’		<i>pakken</i> ‘grab’		<i>halen</i> ‘get’		<i>proberen</i> ‘try’	
feature	assoc.	feature	assoc.	feature	assoc.	feature	assoc.
mismatch	35.69	hit	4.87	bucket	24.82	square	10.95
hole	23.37	unknown1	4.13	out	19.18	block	9.02
near	19.56	unknown2	3.45	on	7.02	mismatch	7.11
block	11.92	cheek_child	2.77	off	6.07	child	5.53
circle	11.04	observer	2.23	floor	4.95	show	5.29

Table 11: Four prepositions/particles and the features most strongly associated with them.

<i>in</i> 'in'		<i>op</i> 'on'		<i>af</i> 'off'		<i>uit</i> 'out'	
feature	assoc.	feature	assoc.	feature	assoc.	feature	assoc.
square	5.17	yellow	3.23	round	19.35	star	16.53
force	4.42	point	2.39	floor	6.25	air	7.35
match	3.44	mother	1.65	mismatch	4.26	in	6.04

Turning to the verbs first, we find two verbs that have readily interpretable situational profiles. The four most strongly associated features for *passen* 'fit' are all elements of the meaning we would like to assign to the verb. Used mostly with negation, *passen* is about the match or mismatch ('doesn't fit') between a figure (i.c. the block) and a ground (i.c. the hole) that is salient because the two are in some relevant relation to each other (i.c. nearness, as the child is trying to fit the block in the hole). The fact that the precise figure and ground are associated is obviously due to the nature of the game: the only things that fit or don't fit are blocks and holes.

The lemma *halen* 'remove, get' is in our corpus always used with either *af* 'off' or *uit* 'out'. The things that are typically removed are the lid from the bucket and the blocks from the bucket. Now, the verb *halen* is indeed associated strongly with features related to these two situations: the bucket is a pivotal object out of which or off of which objects are removed. These objects then enter into a new spatial relationship with another object (typically a support relation with the floor, hence on and floor). As the mode of removal varies (grabbing and moving by hand, positioning the bucket in such a way that the blocks fall out), no specific feature pertaining to the behavior (e.g., move, position) is associated with *halen*.

It may be argued that if a child acts in a certain way when a certain word is used, the child understands that word. Suppose the mother says *Go take the lid off*, and the child subsequently takes the lid off. It does not follow logically that the child does so by having an understanding of the particle verb *take off*, the auxiliary *go* and the noun *lid*. In fact, in situations like the ones occurring in this corpus, it is clear what is expected from the participants. An openable and closeable container with blocks that fit in certain holes and not others needs hardly any instruction to be played with. When the toy is given to the mother and the mother shows the child the toy, and says *Go take the lid off*, the child will just as ready start taking the lid off as when the mother makes a comment about the toy (e.g. *That's a nice puzzle!*) or produces another adhortative utterances (e.g. *Let's do this puzzle*).

For the other two verbs, the situational context is not so clear. This is to be expected from a verb denoting an intention, such as *proberen*, for which encoding the intentional states of the agents might provide a way out, but not so much for a highly concrete verb like *pakken*. We expect *pakken* to become associated with the feature **grab**, but in only in 57.5% of utterances containing *pakken*, that feature is present, whereas it is also present in 52.2% of all utterances not containing *pakken*.

What then, is the situational action *pakken* relates to? Out of the forty cases of utterances in which the lemma occurs, thirty-one have a grabbing event that can be identified with the utterance within a time frame of two three-second windows before and after the window in which the utterance was produced. An example is the utterance in example (5), produced within the fifteen-second fragment of situational context in Table 12. The mother comments on the choice of the daughter to grab the red, round block and does so a few seconds after the action referred to takes place.

- (5) *pak je weer de zelfde*
 grab you again the same
 'You picked the same one again!'

Table 12: The situational context of the utterance *Pak je weer dezelfde*.

<i>time</i>	<i>Predicates*</i>
-2	reach(child,bgrsq)
-1	grab(child,breci) move(child,breci,on-floor,in-air)
0	move(child,breci,in-air,at-hoci)
1	letgo(child,breci) in(breci,bucket)
2	reach(child,bgrsq) grab(child,bgrsq)

* ho = hole, b = block, re = red, gr = green, sq = square, ci = circle.

Widening the window thus seems to make the situation available in which the referred action takes place. Nevertheless, if a learner used a window this wide, the feature **grab** would be present in 93.9% of the utterance-situation pairs. This presents a great amount of noise for the learner, and so if a child is to learn the meaning of *pakken*, other cues must be used. Now, it falls outside of the scope of this exploration to go much deeper into this matter, but as other cues we can imagine bootstrapping the meaning from pre-established object reference: if we know that the pronoun *die* 'that' refers to a certain block, it should be more likely for the lemma *pakken* to be associated with predicates in which that specific block is an argument. The use of these advanced cues can only start when

the child has a basic understanding of other referential expressions in the utterance, and hence cannot be the initial drive behind verb learning. This position resembles that of the Emergentist Coalition Model (Hollich, Hirsh-Pasek, and Golinkoff 2000; Poulin-Dubois and Forbes 2006), in which the validity of different cues for acquiring word meanings changes over ontogenetic time. Another route to pursue would be to look at argument-structure constructions, see if those can be associated with certain clusters of meanings and then use that information to bootstrap the meaning of the verb. This is essentially the idea of syntactic bootstrapping (Gleitman 1990; Naigles 1990), but then from an emergentist perspective: both argument-structure constructions and verbs are acquired through use and mutually reinforce each other's acquisition.

For the prepositions, we see a similar pattern: some have relatively strong associations with the (correct) spatial features and the typical figure and ground objects in those relations (*af* and *uit*). For *op*, the highest four features are object categories (three of which are typical objects found in a support relation), and the fifth is the correct meaning of *op*. Although the feature **on** only comes fifth, it ranks higher than all other spatial, or in general relational features, and we can hence say that the situational context does provide some evidence for the acquisition of the preposition *op*.

The most frequent ($r = 286$) preposition, *in*, however, shows hardly any interesting associations. The reason why the lemma is not associated with the feature **in** is the same as for *pakken*: the feature is highly frequent in general, and hence there are many occurrences of the feature in which the lemma is not used. Furthermore, in only 51.7% of all instances of *in*, a change of state resulting in a containment relation took place. This might be due to the fact that mothers often encourage their child by saying *doe 'm daar maar in* 'go put it in there', when the actual putting-in only takes place half a minute later. This can be seen from the fact that the accompanying pointing gesture is highly associated with the lemma *in*. What cues there are that might guide the child towards the intended meaning of *in*, remains to be seen. In any event, the learner would have to be able to suppress a lot of noise in the data and perhaps ignore several uninformative instances of the lemma. Perhaps the a priori salience of the containment relation (Choi 2006) helps finding the intended meaning as well.

4 Conclusion and directions

It seems that some relational terms are used in a way that an ideal learner can associate them with the correct features. Others, like *pakken*, *proberen* and *in* do

not have any clearly associated features that we would recognize to be valid meaning representations from an adult point of view. Interestingly, *pakken* is a highly concrete verb, denoting a single, atomic predicate in our representation (i.e. **grab**), but there is too much noise for the feature and lemma to be associated. We suggested that other cues might help understand the meaning of *pakken*. In the case of a preposition with a relatively accessible intended meaning like *in*, the situational context provides little help in acquiring that meaning, at least if we see learning as naïve association. How a preposition like *in* is acquired from the situational context, remains to be seen.

It thus seems that Gleitman's disjunction is only partly right. Some relational meanings do have clear correlates in the perceivable situation and are highly associated with the lemmas denoting these meanings. We discussed eight lemmas, five of which were strongly associated with features related to the meaning they were intended to convey. This should be enough to get a learner (at least) started learning terms with relational meanings. One of the other three, *pakken* 'grab', also has grabbing events in its context, only taken a little wider (six seconds before and six seconds after the time frame of the utterance). Surprisingly, *in* 'in' has no good association with the feature denoting a containment relation. For the lemmas for which the situational context does not provide reliable grounding, Gleitman's argument still holds: additional cues are needed to explain their acquisition.

In this study, we presented a method for studying the situational context of child-directed language in more detail. We discussed the decisions and the process, as this is an often overlooked aspect of the study of child language. The main reason for developing such a dataset is to evaluate claims about the excessive richness (size of the hypothesis space) and the poverty of the situational context. We re-evaluated the statement that the situational context contains both too much information to find the right correlate for a word and too little – with the intended relation or action not being present. On the basis of the developed dataset, we concluded that for several relational meanings, the situational context does contain features that are very significantly associated with the word that refers to the relation captured in that feature. Other predicates and prepositions are less strongly associated with features related to their meaning. For these words, it is indeed the case that the hypothesis space is often too big, in the sense that the feature is often present when utterances not containing the word are produced, so that the association becomes weaker. Nevertheless, the association can form a starting point for word learning, which is later supplemented with additional cues (other words in the sentence, syntactic frames, social cues, better understanding of the task), a position that has been worked out in more detail in the Emergentist Coalition Model (Hollich et al. 2000).

What this study shows, is that it is possible to *get started* on learning relational semantics on the basis of mere association. The relevance of this analysis for usage-based theorizing is that domain-specific biases for word learning are not needed to start developing a lexicon. However, learning the symbolic mappings remains difficult, and there often is too much information in the situation. One challenge for usage-based approaches is to develop accounts of structured learning (as opposed to merely associative learning) where constructions (for example) play a role in bootstrapping word meanings, and to apply this in a structured way to naturalistic data like the dataset presented in this chapter.

Using the atomic-feature representation, we plan to work with actual word learning models (e.g. Fazly, Alishahi, and Stevenson 2010) in order to see if the meaning can indeed be learned by a cognitively realistic incremental model. A next step would be to see if it is possible to have a computational learner induce a simple, (semi-)productive grammar that can be used to interpret unseen utterances.

In our exploration we left aside interesting, but more complex issues, such as the relation between the nature of the child-directed language and the output. For prepositions, it would be interesting to see if there is a correlation between the availability of situational context guiding the learner in the correct direction and the age of acquisition or the overextension patterns of certain lemmas (e.g. the preposition *af* ‘off of’ and *uit* ‘out of’ [Bowerman 1995] or the set of *op* ‘canonical support’, *aan* ‘support by hanging and attachment’ and *in* ‘containment’ and the overextension of *op* to *aan*-situations [Bowerman 1993]).

Modeling known phenomena, whether they were established by experimental method or by observation, gives us a greater insight in the cognitive mechanisms underlying them. However, in order to model the acquisition of meaning, a detailed understanding of the context is needed, especially given claims about the nature of that context. With the development of a dataset, we hope to encourage research in language acquisition that takes the situatedness of actual child-directed language into account in a systematic way.

References

- Baillargeon, Renee, and Su-Hua Wang. 2002. Event categorization in infancy. *Trends in Cognitive Sciences* 6 (2): 85–93.
- Barak, Libby, Afsaneh Fazly, and Suzanne Stevenson. 2012. Modeling the acquisition of mental state verbs. In *Proceedings of the NAACL-HLT Workshop on Cognitive Modeling and Computational Linguistics (CMCL)*.

- Behrend, Douglas A., and Jason M. Scofeld. 2006. Verbs, actions, and intentions. In *Action Meets Word: How Children Learn Verbs*, Kathryn Hirsh-Pasek, and Roberta M. Golinkoff (eds.), 286–307. Oxford, UK: Oxford University Press.
- Bowerman, Melissa. 1993. Typological Perspective on Language Acquisition: Do Crosslinguistic Patterns Predict Development. In *Proceedings of the Twenty-fifth Annual Child Language Research Forum*, Eve V. Clark (ed.), 7–15. Stanford, CA: CSLI Publications. 1995. Learning how to structure space for language: A crosslinguistic perspective. In *Language and Space*, Paul Bloom, Mary. A. Peterson, Lynn Nadel, and Merrill F. Garrett (eds.), 385–436. Cambridge, MA: MIT Press.
- Brugman, Hennie, and Albert Russel. 2004. Annotating multimedia/multi-modal resources with ELAN. In *Proceedings of the Fourth International Conference on Language Resources and Evaluation (LREC)*, 2065–2068. Paris: European Language Resources Association.
- Carletta, Jean. 1996. Assessing agreement on classification tasks: The kappa statistic. *Computational Linguistics* 22 (2): 249–254.
- Casasola, Marianella, Jui Bhagwat, and Kim T. Ferguson. 2006. Precursors to verb learning: Infant's understanding of motion events. In *Action Meets Word: How Children Learn Verbs*, Kathryn Hirsh-Pasek, and Roberta M. Golinkoff (eds.), 160–190. Oxford, UK: Oxford University Press.
- Choi, Soonja. 2006. Preverbal spatial cognition and language-specific input: Categories of containment and support. In *Action Meets Word: How Children Learn Verbs*, Kathryn Hirsh-Pasek, and Roberta M. Golinkoff (eds.), 191–207. Oxford, UK: Oxford University Press.
- Cohen, Jacob. 1960. A coefficient of agreement for nominal scales. *Educational and Psychological Measurement* 20 (1): 37–46.
- Fazly, Afsaneh, Afra Alishahi, and Suzanne Stevenson. 2010. A probabilistic computational model of cross-situational word learning. *Cognitive Science: A Multidisciplinary Journal* 34 (6): 1017–1063.
- Feist, Michele I. 2000. On *in* and *on*. An investigation into the linguistic encoding of spatial scenes. Ph.D. diss., Northwestern University.
- Fleischman, Michael, and Deb K. Roy. 2005. Why verbs are harder to learn than nouns: Initial insights from a computational model of intention recognition in situated word learning, In *Proceedings of the 27th Annual Meeting of the Cognitive Science Society*.
- Frank, Michael C., Noah D. Goodman, and Joshua B. Tenenbaum. 2009. Using speakers' referential intentions to model early cross-situational word learning. *Psychological Science* 20 (5): 578–585.
- Frank, Michael C., Joshua B. Tenenbaum, and Anne Fernald. 2013. Social and discourse contributions to the determination of reference in cross-situational word learning. *Language, Learning, and Development* 9: 1–24.
- Gentner, Dedre. 1978. On relational meaning: The acquisition of verb meaning. *Child Development* 49: 988–998. 1982. Why nouns are learned before verbs: Linguistic relativity versus natural partitioning. In *Language Development*. Vol. 2. *Language, Thought, and Culture*, Stan A. Kuczaj II (ed.), 301–334. Hillsdale, New Jersey: Lawrence Erlbaum Associates.
- Gentner, Dedre, and Kurtz Kenneth. J. 2005. Relational categories. In *Categorization Inside and Outside the Lab*, Woo-Kyoung Ahn, Robert L. Goldstone, Bradley C. Love, Arthur B. Markman, and Phillip Wolff (eds.), 151–175. Washington, DC: Amer Psychological Association.
- Gleitman, Lila. 1990. Sources of verb meanings. *Language Acquisition* 1 (1): 3–55.
- Hollich, George. J., Kathryn Hirsh-Pasek, and Roberta M. Golinkoff. 2000. Breaking the language barrier: An emergentist coalition model for the origins of word learning. *Monographs of the Society for Research in Child Development* 65 (3): 1–135.

- Mandler, Jean M. 2006. Actions organize the infant's world. In *Action Meets Word: How Children Learn Verbs*, Kathryn Hirsh-Pasek, and Roberta M. Golinkoff (eds.), 111–133. Oxford, UK: Oxford University Press.
- Meltzoff, Andrew N. and M. Keith Moore. 1995. Infants' understanding of people and things: From body imitation to folk psychology. In *The Body and the Self*, José L. Bermúdez, Anthony Marcel, and Naomi Eilan (eds.), 43–69. Cambridge, MA: MIT Press.
- Naigles, Letitia R. 1990. Children use syntax to learn verb meanings. *Journal of Child Language* 17 (2): 357–374.
- Narasimhan, Bhuvana, and Anetta Kopecka. (eds.). 2012. *Events of putting and taking: A cross-linguistic perspective*. Amsterdam: John Benjamins.
- Poulin-Dubois, Diane, and James N. Forbes. 2006. Word, intention, and action: A two-tiered model of action word learning. In *Action Meets Word: How Children Learn Verbs*, Kathryn Hirsh-Pasek, and Roberta M. Golinkoff (eds.), 262–285. Oxford, UK: Oxford University Press.
- Sootman-Buresh, Jennifer, Amanda L. Woodward, and Camille W. Brune. 2006. The roots of verbs in prelinguistic action knowledge. In *Action Meets Word: How Children Learn Verbs*, Kathryn Hirsh-Pasek, and Roberta M. Golinkoff (eds.), 208–227. Oxford, UK: Oxford University Press.
- Stefanowitsch, Anatol, and Stefan T. Gries. 2003. Collocations: Investigating the interaction of words and constructions, *International Journal of Corpus Linguistics* 8(2): 209–243.
2005. Covarying collexemes. *Corpus Linguistics and Linguistic Theory* 1 (1): 1–43.
- Tomasello, Michael. 2003. *Constructing a language: A Usage-Based Theory of Language Acquisition*, Cambridge, MA: Harvard University Press.