Available online at www.sciencedirect.com

SCIENCE DIRECT

COGNITIVE DEVELOPMENT

Cognitive Development 21 (2006) 174–193

The distributed learning effect for children's acquisition of an abstract syntactic construction

Ben Ambridge[a,*], Anna L. Theakston[b], Elena V.M. Lieven[c], Michael Tomasello[c]

[a] *Max Planck Child Study Centre, University of Manchester, Manchester, UK*
[b] *University of Manchester, Manchester, UK*
[c] *Max Planck Institute for Evolutionary Anthropology, Leipzig, Germany*

Abstract

In many cognitive domains, learning is more effective when exemplars are distributed over a number of sessions than when they are all presented within one session. The present study investigated this *distributed learning effect* with respect to English-speaking children's acquisition of a complex grammatical construction. Forty-eight children aged 3;6–5;10 (Experiment 1) and 72 children aged 4;0–5;0 (Experiment 2) were given 10 exposures to the construction all in one session (massed), or on a schedule of two trials per day for 5 days (distributed-pairs), or one trial per day for 10 days (distributed). Children in both the distributed-pairs and distributed conditions learnt the construction better than children in the massed condition, as evidenced by productive use of this construction with a verb that had not been presented during training. Methodological and theoretical implications of this finding are discussed, with particular reference to single-process accounts of language acquisition.

Keywords: Syntax acquisition; Construction learning; Distributed learning; Spacing effect

Children's initial multiword utterances mostly derive from lexically specific schemas, such as *I want [X], I'm [X]ing it*, and so forth (Braine, 1976; Tomasello, 1992). They then generalize across these to form more abstract constructions,[1] such as *[SUBJECT] [VERB] [OBJECT]*. In usage-based accounts of language acquisition, children form these abstractions via processes

* Corresponding author at: Department of Psychology, University of Liverpool, Bedford Street South, Liverpool L69 7ZA, UK. Tel.: +44 151 794 1109; fax: +44 151 794 6937.
E-mail address: Ben.Ambridge@liverpool.ac.uk (B. Ambridge).

[1] In this paper, the term *construction* is used in its most common sense of *verb argument structure construction*. We would argue, however, that the phenomena discussed here are equally applicable to other types of construction (e.g., morphological constructions, such as the past tense VERB+ed construction).

0885-2014/$ – see front matter
doi:10.1016/j.cogdev.2005.09.003

of schematization and analogy (Tomasello, 2003). Importantly, schematization and analogy are *general* learning processes that have their conceptual origins in non-linguistic domains. Piaget (1952) discusses the development of generalized sensory-motor schemas for acting on objects in infancy, and Gentner and colleagues have investigated analogy formation in various domains of later conceptual development (e.g., Gentner, Loewenstein, & Thompson, 2003; Kotovsky & Gentner, 1996; Lowenstein & Gentner, 2001).

Under their *structure mapping* theory, Markman and Gentner (1993) argue that the basis of analogy is relational similarity: the existence of similar internal relations between elements of the two structures to be mapped (see also Goldstone, 1996; Holyoak & Koh, 1987; Holyoak & Thagard, 1989). Providing the two structures share this relational similarity, *object commonality* – similarity between a particular element of one structure and the corresponding element of the other structure – is not required for an analogy to be made. As an experimental demonstration, children might be shown two pictures: one of a truck towing a car and the other of an identical car towing a boat. When asked to indicate the item in the second picture that is the "best match" for the car in the first picture, children tend to choose not the car, but the boat. Children are able to align the *tow-er/tow-ee* structure of the two pictures, and form an analogy between the two *tow-ees*.

As an analogous linguistic example, consider the hypothetical utterances *I kiss Mummy* and *Daddy threw the ball*. Although the two utterances have no morphemes in common, they share relational similarity such that the agent–action relation between *I* and *kiss* parallels the relation between *Daddy* and *threw*, whilst the action–patient relation between *kiss* and *Mummy* parallels the relation between *throw* and *the ball*. Thus, this relational similarity allows the child to form an analogy between the two utterances and move towards a wholly abstract *SVO* construction schema.

Although the formation of abstract constructions (via schematization/analogy) is a crucial feature of any constructivist acquisition theory, very little experimental research has investigated the specific details of this process. Factors that have been hypothesized to influence this process include (1) the token frequency of the schema in the child's input (Cameron-Faulkner, Lieven, & Tomasello, 2003; Rowland & Pine, 2000; Theakston, Lieven, Pine, & Rowland, 2001, 2002); (2) the variability or type frequency of variable elements in the construction (e.g., of the VERB in the *I'm [VERB]ing it* construction; Bybee, 1995; Gomez, 2002; Onnis, Monaghan, Christiansen, & Chater, 2004); (3) the presence in the construction of invariant elements, such as pronouns and inflectional morphemes (Childers & Tomasello, 2001; Pine, Lieven, & Rowland, 1998); and (4) prior knowledge of related constructions (Abbot-Smith & Behrens, in press; Croft, 2001; Elman et al., 1996; Goldberg, 1995; Langacker, 1988, 1991, 2000; Ruhland, Wijnen, & van Geert, 1995).

One factor that has not been studied is the temporal distribution of exemplars of the construction in the input. There are three plausible hypotheses. First, based on previous learning research, we might hypothesize that construction learning will show the well-known *distributed learning* or *spacing* effect: given a certain number of exposures to a stimulus, or a certain amount of training, learning is always better when exposures or training trials are distributed over several sessions than when they are massed into one session. This finding is extremely robust in many domains of human cognition. For example, in a meta-analysis of 97 studies, Janiszewki, Noel, and Sawyer (2003) found a distributed learning effect for meaningful and meaningless stimuli (real words versus nonce), familiar and novel stimuli, isolated and embedded stimuli (single word alone versus target word in a sentence), verbal and pictorial stimuli, and for simple stimuli (single words), structurally complex stimuli (sentences) and semantically complex stimuli (homographs). The effect was observed using both intentional and incidental learning paradigms, with or without simple or complex intervening material between stimuli, and at many different interval lags ranging from minutes to months, assessed with tests of both recall (spontaneous and cued) and recognition. The distributed learning effect is not confined to verbal memory, purely cognitive tasks, or even to humans (see Dempster, 1996; Underwood, 1961 for reviews).

Perhaps of special importance in the current context, two studies have found a distributed learning effect for the learning of words in young children (Childers & Tomasello, 2002; Schwartz & Terrell, 1983). Most dramatic is Childers and Tomasello's (2002) finding that, for both nouns and verbs, four presentations on different days were more effective than eight presentations on a single day. As in the usage-based account, word learning and construction learning are conceptualized as a single process, simply operating on a different scale, it would seem plausible that the distributed learning effect would also apply to construction learning (the *distributed-advantage* hypothesis).

However, there is an important difference between word learning and construction learning: Word learning involves repeated presentation of identical lexical items, whereas construction learning involves different exemplars, so that the creation of the construction requires an abstraction across exemplars. To our knowledge, no study to date has investigated the possibility of a distributed learning effect for the abstraction process using linguistic stimuli. In fact, some experimental investigations of structure mapping have shown an advantage for massed over distributed presentation in non-linguistic domains. For example, in a study conducted by Loewenstein and Gentner (2001), 3-year-old children were shown the location of a toy in two structurally identical but perceptually different *hiding rooms*, then searched for the toy in third room, structurally identical to the first two. Children who were shown the two hiding rooms simultaneously (massed presentation) were more successful at locating the toy in the third room than children who were shown the two hiding rooms one at a time (distributed presentation), presumably because simultaneous presentation of the two rooms facilitated the comparison process (see also Gentner, Loewenstein, & Thompson, 2003; Kotovsky & Gentner, 1996).

A second, competing hypothesis in the present study, therefore, is that acquisition of a partially abstract construction will be facilitated by massed as opposed to distributed presentation (the *massed-advantage* hypothesis) because the relational similarity between the utterances will be more apparent when those utterances are temporally contiguous than when they are more widely distributed in time.

A third hypothesis is that the optimum presentation schedule is one which involves distributed presentation of pairs (or small numbers) of utterances that instantiate the target construction (the *distributed-pairs-advantage* hypothesis). Gentner (personal communication, May 20th, 2004) suggests that this schedule combines the advantages of massed and distributed presentation as the presentation of two instantiations of the construction together "will help [the child] extract and encode the relation ... and having those pairs spaced will help him consolidate future access to that relation".

The aim of the present study, then, was to investigate which of these three hypotheses, if any, most accurately describes early construction learning. Toward this end, in the present study 3- to 5-year-old children were presented with 10 different instantiations of the highly infrequent object-cleft construction all at once (massed) or of a single instantiation per day for 10 days (distributed) or on a schedule of one pair per day for 5 days (distributed-pairs). A control group was given single instantiation. Children's learning of this construction was then assessed using an elicited-production test, in which they were required to use this construction with a verb that was not presented in this construction during training.

1. Experiment 1: construction learning—massed versus distributed-pairs

Experiment 1 was designed to compare the *massed-advantage* and *distributed-pairs-advantage* hypotheses.

1.1. Method

1.1.1. Participants

Participants were 36 children aged 3;6–4;6 (M = 4;3) and 36 children aged 4:10–5:10 (M = 5;4) consisting of an approximately equal number of males and females. All participants had English as their first language and attended primary schools in Manchester, England. In total, 117 children were tested, of whom 45 were subsequently excluded. Children were excluded if they were unable to complete the warm-up session (10 children), displayed uncooperative behaviour having begun the study (5 children), were absent from school for one or more training sessions (11 children), or, having begun the study, failed to repeat at least four target utterances during training (12 children) or (for the control groups) the sole target utterance (7 children). It must be acknowledged that the drop-out rates for both this study and Experiment 2 were high. However, if we do not count children that were excluded before the training phase of the study began (a figure which is often unreported), and those who were absent for one or more sessions, then the drop-out rates are 20% and 24%, respectively. That these rates are still relatively high reflects the need to exclude children who were unable to successfully repeat E2 during training. Pilot testing revealed that children who did not repeat at least four target utterances over 10 trials during training typically produced no scorable responses at test.

1.1.2. Materials

Over the course of the study, five different animal puppets performed actions on 10 familiar inanimate objects (such as a cup, a tree and a cake). During the training and test sessions the experimenter used 14 monosyllabic, transitive, English verbs to describe the various actions performed by the animals on the objects (*bite, hold, touch, take, punch, hide, choose, grab, rub, pull, move, kick, drop, find*). The verbs used were selected as the 14 most frequent monosyllabic, transitive verbs from the spoken-texts section of the British National Corpus that could easily be performed by the puppet characters on inanimate objects,[2] (verbs and frequencies can be found in Table A.1).

The construction chosen for this investigation was the past tense object-cleft construction *It was the [OBJECT] that the [SUBJECT] [VERBed]*[3] as instantiated by such sentences as *It was the cup that the frog took*. Note that, in common with constructions, such as the *for* dative and the *by* passive, this construction is partially abstract and partially concrete: Only the nouns and the verb vary, and the morphological *–ed* marker is often present on the verb. This construction was chosen for three reasons. Firstly, the construction is relatively complex and infrequent and as such, is unlikely to be previously known to young language learners. Secondly, the construction does not conform to canonical English word order. Thus, children are less likely to revert to a

[2] The verbs used were 14 of the 22 most frequent appropriate verbs in the corpus. The remaining eight (*wash, cut, break, bash, hit, throw, eat and push*) were selected for use in additional tests of repetition and priming which were discontinued after pilot testing.

[3] For clarity, this notation is used to represent past tense forms. In fact, 6 of the 14 verbs used have irregular past tense forms.

canonical *SVO* construction (such as a simple transitive), when attempting to copy or produce this construction, than they would be for constructions that conformed to *SVO* ordering (such as the subject-cleft construction: *It was the frog that took the cup*). Finally, the object-cleft construction has the pragmatic function of drawing attention to the patient of a transitive action. Thus, it was relatively easy to construct an experimental scenario where the use of this construction would be natural and pragmatically felicitous.

1.1.3. Design and procedure

In a warm-up session, children were introduced to the two experimenters and invited to name the animal puppets and objects to be used in the study, which the vast majority were easily able to do. As well as serving as a warm-up, this ensured that all children knew the names of the animals and objects before the start of the experiment. As a further warm-up, children were asked to repeat, following the second experimenter (E2), five intransitive utterances, each of which included one of the animal characters as the subject (e.g., *The duck is flying, The bear is sleeping*). This was to introduce the children to the procedure of the training section of the experiment, where they would be required to repeat E2 but not the first experimenter (E1).

Each child (except for those in the control group) then participated in 10 training trials. Each of these trials comprised one exposure to the past tense object-cleft construction schema *It was the [OBJECT] that the [SUBJECT] [VERBed]* (with the sentence repeated). For each trial, the procedure was as follows: the experimenter (E1) selected the appropriate animal puppet and two objects—the patient object and the distracter object. E1 then made the puppet perform the appropriate action on the patient object, whilst describing the action performed using a simple transitive sentence in the past tense (e.g., *The frog took the ball*). However, E1 always (apparently mistakenly) named not the actual patient object but the distracter object. The second experimenter (E2) then corrected the experimenter, always using a past tense object-cleft construction (e.g., *No, it was the cup that the frog took!*). E2 then repeated this utterance to give the two sentences constituting one training trial for the utterance type and invited the child to repeat the utterance. The procedure for each of the 10 trials is summarized in the following example.

E1: (*produces frog, cup, ball; frog takes cup*) The frog took the ball.
E2: No, it was the *cup* that the frog took. It was the *cup* that the frog took.
C: It was the *cup* that the frog took.

The verb, the subject (agent), the patient object and the distracter object were always selected at random by computer, with the stipulation that each verb could appear only once, each subject twice and each object twice (once as a patient object and once as a distracter object). The schedule on which the 10 training trials were presented was manipulated as a between-subjects variable. Trials were presented either all in one session with one immediately following the other (massed condition), or on a schedule of two trials per day, presented consecutively, for five consecutive days (distributed-pairs condition). Children in a control condition were given a single training trial. Ideally, these children would have received no training and completed only the test phase. However, pilot testing revealed that children who had received no training were not able to comprehend the procedure of the test phase (an issue to which we will return after introducing the test phase procedure).

For the children in the distributed-pairs condition, the experimenters took care to minimize any extra contact time as compared with children in the massed condition, giving only minimal greetings and no new instructions.

Generally, children learned quickly to repeat E2 with minimal or no prompting. For all except the first trial, E2 presented the sentence twice, whether or not children successfully repeated it.

For the first trial, E2 presented the sentence as many times as was necessary to elicit a successful repetition, up to a maximum of five presentations (except for the control group, for whom the sole utterance was presented twice). Children's attempts at repetition during training were recorded and scored according to the same criteria used for utterances produced during the test session, as outlined below.

The elicited-production test was presented immediately after the end of the training phase (i.e., on Day 5 for the distributed-pairs group, and on the one and only day for the massed and control groups). The rationale of the test was to investigate the extent to which children had acquired a productive, partially abstract object-cleft construction by assessing their ability to use a verb not presented during training in this construction.

Immediately after the final training trial, E1 introduced the test session saying "Now we're going to do some more, but this time E2 isn't going to tell me if I get it wrong, because I want you to tell me. But I want you to tell me exactly like E2 would, try and say it exactly like E2". E1 then followed the same procedure as for the verbs in the training session, using four verbs not used during training (*move, kick, drop, find*) with the child attempting to provide a correction for each of the experimenter's four erroneous transitive utterances. All children readily understood the game and attempted to provide a correction at the relevant point. As during training, the order of presentation of the four verbs and the selection of the subject and objects was randomized. The reason why children in the control group were given one training trial (as opposed to zero) should now be apparent. Without at least one training trial, children in the control group would not know that they should use the object-cleft construction to correct E1 during the test phase, even if this construction were part of their productive repertoire. Indeed, as mentioned previously, pilot testing revealed that, without exception, children who were given no training produced no response (other than a simple "no" or a shake of the head) at test.

1.1.4. Scoring

In the course of the test phase, four utterances were elicited from each child. However, every utterance produced by each child in this phase was recorded and included in the analysis (in practice, almost every child in both this study and Experiment 2 produced four utterances, with only five children across both studies producing five or six). Retraced part utterances were not counted. For example, the utterance *It was <the frog uh> [//] the cup that the frog bit* would be classified as a target object-cleft utterance. Using the scoring criteria shown in Table A.2, every utterance that used a target object-cleft construction, or one of four non-target constructions, was classified into one of the following mutually exclusive categories.

- Target object-cleft, for example, *It was the cup that the frog took.*
- Object-NP clause (or *reduced cleft*), for example, *It was the cup.*
- Subject-cleft, for example, *It was the frog that took the cup.*
- Transitive, for example, *The frog took the cup.*
- Subject–object error: Swap or duplication of subject (agent) and/or object (patient), for example, *It was the frog that the cup took.; It was the frog that the frog took.*

If, for a given trial, a child gave no response, or a response that did not fit into any of these categories, the trial was scored as *other*. This procedure was designed to avoid large number of irrelevant utterances entering into the analysis, whilst ensuring that a minimum of four data points were recorded for each child.

Utterances in which a child substituted particular items for different but semantically appropriate items from the same syntactic category were scored as if no such substitution had occurred (see Table A.2 for permitted substitutions). For example, an utterance such as *That's a cup what he dropped* would be scored as a correct form of the target utterance *It was the cup that the frog dropped* since the substitutions of *that* for *it* (demonstratives), *is* for *was* (forms of *BE*), *a* for *the* (articles), *what* for *that* (complementizers) and *he* for *the frog* (NPs) still result in a grammatical object-cleft utterance, and do not affect the underlying syntactic structure of the utterance. The only exception to this was if, in the test session, a child substituted a verb that had been presented in an object-cleft construction during training. In this case, the utterance was scored as unclassifiable (other) as it cannot be taken as evidence for the formation of an abstract object-cleft construction. Substitutions for verbs not previously presented during the experiment were allowed. Errors of verb agreement and tense were also ignored (in fact, such errors virtually never occurred, perhaps because E1 always supplied the appropriate verb form in his erroneous transitive sentence). On no occasion did a child substitute the distracter object for the patient object, or for any other object; neither did any child substitute the subject for another character.

One scoring decision that requires clarification is the decision to classify *what* as a correct complementizer in this construction (e.g., *It was the cup what the frog dropped*). This is a reflection of the fact that in many Northern dialects of British English, including the Manchester dialect spoken by the majority of the participants, *what* is an acceptable relative pronoun in this and related constructions. *What* and *that* were used with approximately equal frequency overall, with some children using exclusively *what* for the entire course of the experiment, even though, for much of the study, children were attempting to copy verbatim the experimenter's utterances in which *what* never occurred.

Children's attempts to repeat E2 during the training phase of the experiment were also recorded and scored according to the same criteria. A few children who had difficulty in producing an appropriate utterance adopted a strategy of "shadowing" E2, copying an utterance at the same time as it was being produced. Such cases were scored as *other*. In all cases, one experimenter wrote down the child's utterances, whilst the other interacted with the child.

In order to check for inter-rater reliability, test sessions for 20 children (10 each from Experiments 1 and 2) were transcribed and coded by a second coder, blind to the hypotheses under investigation. All non-target responses were collapsed together, so that each response was classified as either a target or non-target response. The overlap between the two coders was 98.8% ($K = 0.96$, $p < 0.001$), denoting agreement on all but one utterance.

1.2. Results

Statistical analyses were conducted on the proportion of each child's utterances that utilized the target object-cleft construction as a function of the total number of utterances produced by that child. As all analyses were conducted on proportional data, a natural logarithmic transformation was applied to the data. This transformation also corrects for heterogeneity of variance, which occurred for some groups on some measures as a result of floor/ceiling effects. All figures and tables show untransformed proportional data.

1.2.1. Analysis of correct responses

Fig. 1 shows the mean proportion of utterances which matched the target object-cleft construction schema for the different training schedule and age groups. Data from the control group are not included in this figure as no child produced a single target utterance.

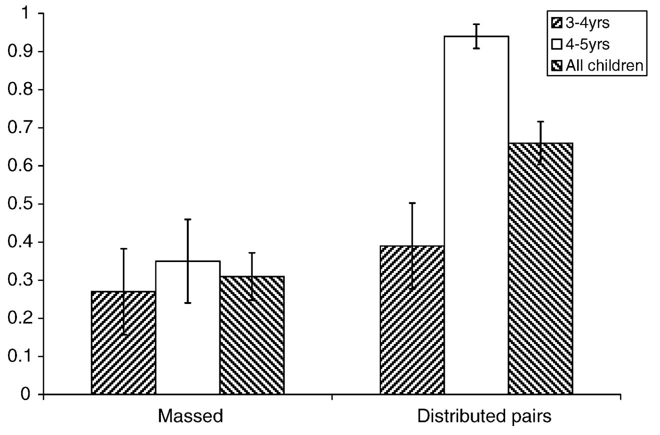

Fig. 1. Experiment 1. Mean proportion of target object-cleft utterances as a function of each child's total number of utterances by age group and condition (error bars represent standard error). Note: data from the control group are not shown, as no child produced a target response.

A 3 × 2 ANOVA was calculated to investigate the effects of training schedule and age on the proportion of each child's utterances that used the target object-cleft construction. This yielded significant main effects of training schedule ($M_{\text{massed}} = 0.31$, $M_{\text{distributed-pairs}} = 0.66$, $M_{\text{control}} = 0$, $F_{2,71} = 30.35$, partial $\eta^2 = 0.48$, $p < 0.001$) and of age ($M_{3\text{–}4\text{ years}} = 0.22$, $M_{4\text{–}5\text{ years}} = 0.43$, $F_{1,72} = 7.80$, partial $\eta^2 = 0.10$, $p = 0.007$) and a significant training schedule by age interaction ($F_{2,66} = 5.28$, partial $\eta^2 = 0.14$, $p = 0.007$). Post hoc tests (Fischer's LSD) revealed that the distributed-pairs group significantly outperformed the massed group ($p < 0.001$), with both of these groups outperforming the control group at $p = 0.001$ or better. To correct for multiple comparisons, a significance level of $p = 0.01$ was adopted for all post hoc tests.

To investigate the nature of the interaction, pairwise comparisons (Fischer's LSD) were used to compare the performance of the different training schedule groups at each age. For the younger children, the performance of the distributed-pairs ($M = 0.39$) and massed ($M = 0.27$) groups did not differ significantly ($p = 0.28$, n.s.). However, it is noteworthy that the distributed-pairs group significantly outperformed the control group ($M = 0$, $p = 0.001$), whereas the massed group did not ($p = 0.03$), at least at the adopted significance level of $p = 0.01$.

For the older children, all comparisons reached statistical significance at $p = 0.005$ or better. The distributed-pairs group ($M = 0.94$) outperformed the massed group ($M = 0.35$), with both groups outperforming the floor-level control group ($M = 0$).

Since the production of one single object-cleft utterance with a verb that was not presented in this construction during training constitutes evidence of a child having formed some kind of abstract construction, perhaps a more appropriate comparison is between the number of children in each training group that produced one or more such utterances. These data are shown in Table 1.

For both age groups, more children in the distributed-pairs training condition than the massed training condition produced one or more target object-cleft utterances at test. No child in the control group produced a target utterance. A series of chi-square calculations compared each of the three conditions for the number of children at each age that produced (or failed to produce) at least one object-cleft. The findings mirror those of our previous analysis. For the younger group, significantly more children in the massed and distributed-pairs conditions than the control condition produced at least one object-cleft at test ($\chi^2 = 6.31$,

Table 1
Experiment 1: number of children producing at least one target utterance

Age	Training schedule	Group (n)	Children producing at least one cleft
3–4 years	Control	12	0
	Massed	12	5
	Distributed-pairs	12	8
4–5 years	Control	12	0
	Massed	12	5
	Distributed-pairs	12	12
All children	Control	12	0
	Massed	24	10
	Distributed-pairs	24	20

$p = 0.03$ and $\chi^2 = 12.0$, $p < 0.001$, respectively). The massed and distributed-pairs conditions did not differ significantly ($\chi^2 = 1.51$, n.s.). For the older children, significantly more children in the distributed-pairs condition than the massed condition produced at least one cleft ($\chi^2 = 9.88$, $p < 0.01$), with more children in each of these two conditions than the control condition producing a target utterance ($\chi^2 = 6.31$, $p = 0.03$ and $\chi^2 = 24$, $p < 0.001$ for the massed and distributed-pairs conditions, respectively). It is particularly interesting to note that every 4- to 5-year-old child who followed a distributed-pairs training schedule produced at least one object-cleft utterance, as compared to less than half of those who followed a massed training schedule.

Considering all the children together, twice as many children in the distributed-pairs condition as the massed condition produced one or more object-cleft utterances at test, and this reached statistical significance at $p < 0.01$ ($\chi^2 = 8.89$). The difference between each of the training groups and the control groups was significant at $p < 0.001$ in both cases, with χ^2 values of 12.67 and 34.29 for the massed and distributed-pairs conditions, respectively.

1.2.2. Analysis of non-target responses

In many cases, where children did not produce a target object-cleft utterance, they instead used an alternative construction that appeared to be related in some way to the target construction. Children's use of non-target constructions in this study, then, has potential theoretical implications for accounts of language acquisition under which children acquire a hierarchically ordered network of interrelated constructions (Abbot-Smith & Behrens, in press; Croft, 2001; Elman et al., 1996; Goldberg, 1995; Langacker, 1988, 1991, 2000; Ruhland et al., 1995). Table 2 shows the mean proportion of utterances that were classified into each of the five non-target categories. These data, and their implications for *construction conspiracy* approaches will be considered in Section 3.

Table 2
Experiment 1: mean proportions of non-target utterances as a function of each child's total number of utterances, and corresponding standard deviations

Object-NP clause		Subject-cleft		Transitive		Subject–object error		Other error	
Mean	S.D.	Mean	S.D.	Mean	S.D.	Mean	S.D.	Mean	S.D.
0.15	0.29	0.06	0.19	0.13	0.25	0.01	0.03	0.17	0.29

1.2.3. Analysis of training phase

During the training phase of the experiment, all children were asked to repeat each of the cleft-construction sentences produced by the experimenter. It is possible, then, that the significant difference between the performance of the distributed-pairs and massed groups is simply a product of the differential ability of the two experimental groups to repeat these cleft sentences during training. This would not seem likely for two reasons. First, children in the control group necessarily achieved a 100% success rate in repeating E2's sole cleft utterance during training (else they did not proceed to the test phase) but did not produce any such utterances at test. Second children in the remaining two groups proceeded to the test phase only if they had successfully imitated at least four of E2's utterances during training.

Nevertheless, to investigate the possibility that children in the distributed-pairs training group produced more object-cleft utterances during training than did children in the massed-presentation group, a 2 × 2 training schedule by age ANOVA was calculated for the proportion of children's attempts to repeat E2 during training that successfully matched E2's utterance. (The control group was not included in the ANOVA as, due to our exclusion criteria, it was not possible for control group participants to score less than 100% on this measure.) This ANOVA yielded no significant main effect of training schedule ($M_{\text{massed}} = 0.72$ [corresponding to a mean of 7.62 correct repetitions per child], $M_{\text{distributed-pairs}} = 0.75$ [mean of 8.29 repetitions], $F_{1,44} = 0.81$, partial $\eta^2 = 0.02$, n.s.), nor any interaction (but did reveal a significant main effect of age, such that the older children successfully imitated E2 on a higher proportion of occasions during training than did the younger children: $M_{4\text{–}5\text{ years}} = 0.88$, $M_{3\text{–}4\text{ years}} = 0.62$, $F_{1,44} = 18.60$, partial $\eta^2 = 0.29$, $p < 0.001$). Thus it cannot be argued that the main effect of training schedule observed for the test session can be attributed to the differential ability of the different training schedule groups to successfully follow the training procedure.

2. Experiment 2: construction learning—massed versus distributed-pairs versus distributed

Experiment 1 demonstrated that a distributed-pairs presentation schedule facilitated learning of a partially abstract syntactic construction when compared with a massed presentation schedule. Experiment 2 was designed to compare each of these training schedules to a more widely distributed schedule (of one training trial per day), to see if an even more distributed schedule would help or hinder children's learning. Under Gentner's hypothesis, this new distributed condition should be difficult because children have no chance to compare across sentence types within a reasonable time frame. In this study we also manipulated the variability of the variable element – in this case the verb – during training. This manipulation (henceforth referred to as the *type frequency* manipulation) was designed to test the hypothesis that increased variability of the slot-filler item would facilitate the acquisition of the abstract slot (see Bybee, 1995; Gomez, 2002; Onnis et al., 2004).

2.1. Method

2.1.1. Participants

Participants were 72 children aged 4;0–5;0 ($M = 4;6$) consisting of an approximately equal number of males and females. This age group that was selected as an informal analysis of the results of Experiment 1 revealed that children of this age demonstrated most variability and did not show floor or ceiling effects. Again, all participants had English as their first language and attended primary schools in Manchester, England. In total, 148 children were tested, of whom 76 were subsequently excluded. Children were excluded if they were unable to complete the warm-up session (24 children), displayed uncooperative behaviour having begun the study (15 children), were absent from school for one or more training sessions (16 children), or having begun the study, failed to repeat at least four target utterances during training (21 children). Again, not counting children who were excluded after the warm-up, or due to absence, the drop-out rate for this study was 24%.

Given that the two modifications introduced for Experiment 2 (the addition of an additional experimental group and of a type frequency manipulation) do not affect the control group, it was deemed unnecessary to recruit a new control group for Experiment 2. Thus, in the statistical analyses reported for Experiment 2, we included as our control the data from the control group participants of Experiment 1 aged between 4;0–5;0 (the age-range of the other three groups for Experiment 2). Twelve children (three from the 4–5 age group and nine from the 3–4 age group of Experiment 1) meet this criterion.

2.1.2. Design

The 72 newly-recruited children were randomly assigned to one of the three experimental conditions corresponding to the three different training schedules: massed, distributed-pairs and distributed. The massed and distributed-pairs conditions were identical to the corresponding conditions in Experiment 1. The new distributed condition consisted of 10 training trials presented on a schedule of one per session for 10 “daily” sessions (in fact trials were presented for 5 days a week – Monday to Friday – for 2 weeks).

In Experiment 2, verb-type frequency was manipulated as a second independent variable with two levels: high (10 verb types, identical to Experiment 1) and low (2 verb types, selected at random from the 10). For children in the low condition, the two verbs were used alternately. This ensured that whilst children in the distributed-pairs condition heard the same two verbs used every day, each of the two instantiations of the construction presented on a particular day did not use the same verb.

2.1.3. Procedure

Apart from the modifications associated with the new training schedule condition, and the new verb-type frequency variable, the procedure for Experiment 2 was identical to that of Experiment 1 with respect to the training, test and scoring procedures, and the materials used.

2.2. Results

As with Experiment 1, statistical analyses were conducted on the proportion of each child’s utterances that utilized the target object-cleft construction as a function of the total number of utterances produced by that child (again, a natural logarithmic transformation was applied to the data). Since a 3×2 ANOVA (excluding the control group who heard only one verb used during training) revealed that the independent variable of verb-type frequency was not associated with any main effects or interactions, all subsequent analyses used one-way ANOVAs to investigate the training schedule manipulation, collapsing across the two verb-type frequency conditions.

2.2.1. Analysis of target responses

Fig. 2 shows the mean proportion of utterances that matched the target object-cleft construction schema for the different training schedule groups. Again, the data for the control group, that produced no target responses, are not shown.

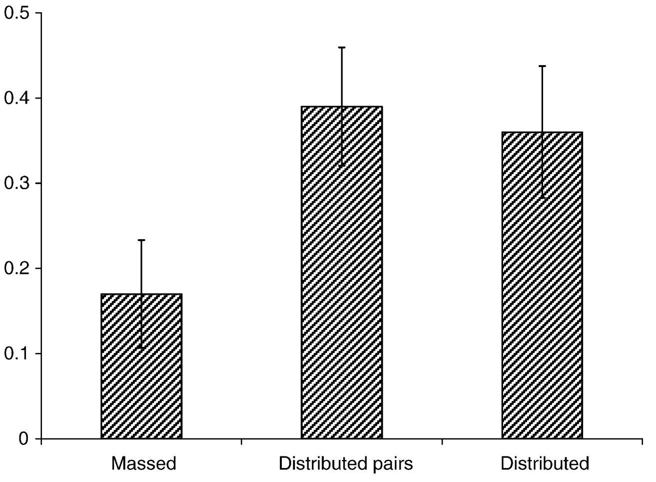

Fig. 2. Experiment 2. Mean proportion of target object-cleft utterances as a function of each child's total number of utterances by condition (error bars represent standard error). Note: data from the control group (a subset of the control group for Experiment 1) are not shown, as no child produced a target response.

A one-way ANOVA was calculated to investigate the effect of training schedule on the proportion of each child's utterances that used the target object-cleft construction. This yielded a significant main effect of training schedule ($M_{\text{massed}} = 0.17$, $M_{\text{distributed-pairs}} = 0.39$, $M_{\text{distributed}} = 0.36$, $M_{\text{control}} = 0$), $F_{3,80} = 6.09$, partial $\eta^2 = 0.19$, $p = 0.001$). Post hoc tests (Fischer's LSD) with a significance level of $p = 0.01$ revealed that both the distributed-pairs and distributed groups outperformed the control group at $p = 0.001$ or better, whereas the massed group did not ($p = 0.12$, n.s.). The distributed-pairs group significantly outperformed the massed group ($p = 0.01$), although the difference between the distributed and massed groups narrowly failed to reach significance ($p = 0.03$) at the adopted level of $p = 0.01$. Whilst this finding provides limited support for the distributed-pairs-advantage hypothesis, there was no tendency for the distributed-pairs group to outperform the distributed group ($p = 0.74$, n.s.).

Again, since the production of one single target object-cleft construction using a verb that was not presented in this construction during training constitutes evidence of the child having formed some kind of abstract construction, perhaps a more appropriate measure is the number of children in each training group who produced one or more such utterances. These data are shown in Table 3.

As Table 3 shows, more children in the distributed-pairs and distributed training conditions than the massed training condition produced one or more target object-cleft utterances

Table 3
Experiment 2: number of children producing at least one target utterance

Training schedule	Group (n)	Children producing at least one cleft
Control	12	0
Massed	24	7
Distributed-pairs	24	16
Distributed	24	14

Table 4
Experiment 2: mean proportions of non-target utterances as a function of each child's total number of utterances, and corresponding standard deviations

Object-NP clause		Subject-cleft		Transitive		Subject–object error		Other error	
Mean	S.D.	Mean	S.D.	Mean	S.D.	Mean	S.D.	Mean	S.D.
0.11	0.28	0.07	0.18	0.17	0.32	0.04	0.12	0.30	0.31

during the elicited-production test, whilst children in the control group produced no such utterances.

This analysis did reveal a significant difference between both the distributed and distributed-pairs groups and the massed group. More children in the distributed group ($\chi^2 = 4.19$, $p < 0.05$) and the distributed-pairs group ($\chi^2 = 6.76$, $p < 0.01$) than the massed group produced at least one target utterance at test. No significant difference was found between the distributed and distributed-pairs groups ($\chi^2 = 0.35$, n.s.), whilst the difference between each of the three training groups and the control group was significant (massed $\chi^2 = 4.34$, $p < 0.05$; distributed $\chi^2 = 11.45$, $p < 0.001$; distributed-pairs $\chi^2 = 14.4$, $p < 0.001$).

2.2.2. Analysis of non-target responses

Table 4 shows the mean proportion of utterances that were classified into each of the five non-target categories. The implications of these errors will be taken up in Section 3.

2.2.3. Analysis of training phase

An ANOVA calculated for the proportion of children's attempted repetitions of E2 during training that were successful imitations (again excluding the control group, who as an inclusion criterion displayed a 100% success rate with their sole training utterance) yielded a significant main effect of training schedule ($M_{\text{massed}} = 0.80$ [corresponding to a mean of 8.54 correct repetitions per child], $M_{\text{distributed-pairs}} = 0.63$ [7.42 correct repetitions], $M_{\text{distributed}} = 0.79$ [8.79 correct repetitions], $F_{2,69} = 3.32$, partial $\eta^2 = 0.09$, $p = 0.04$). However, since the distributed-pairs group, which exhibited the best performance at test actually produced the fewest correct repetitions during training, the main effect of training schedule observed for the test session cannot be attributed to differences between the groups with respect to the training procedure. It cannot be the case that producing a larger number of instantiations of the construction during training somehow inhibited the production of this construction at test (a habituation or interference effect) since the control group, who heard and produced only a single instance of the construction during training, failed to produce a single exemplar at test.

3. General discussion

The results of Experiments 1 and 2 provide clear evidence of a distributed learning effect for construction learning. Across both studies, significantly more children who followed either a distributed (1 trial per day for 10 days) or distributed-pairs (2 trials per day for 5 days) training schedule than a massed (10 trials in a single sitting) or control (1 trial) schedule produced at least one target object-cleft construction during the test phase. Additionally, for Experiment 2 and the older children of Experiment 1, children who followed a distributed-pairs schedule produced significantly more target utterances at test than children in the massed or control training groups. Children in the distributed group (Experiment 2) produced more target utterances than children in the control or the massed group, although this latter difference narrowly failed to reach significance at the adopted significance level of $p = 0.01$.

The results of Experiment 2 do not support the distributed-pairs-advantage hypothesis. No difference was found between the distributed and distributed-pairs groups with regard to either the number of target object-cleft utterances produced, or the number of children producing at least one such utterance. Additionally, both the distributed and distributed-pairs groups outperformed both the massed and control groups with regard to the number of children who were able to produce at least one target utterance at test. It could be argued that the fact that the distributed group did not produce significantly more utterances than the massed group provides support for the distributed-pairs advantage hypothesis, and counts against the interpretation that the findings provide support for the distributed learning effect. This conclusion is unwarranted for four reasons. First, since the production of at least one target utterance provides clear evidence that the child has formed an abstraction, the number of children producing one such utterance – a measure on which the massed and distributed groups differed significantly – is a more appropriate measure on which to compare the experimental groups than the total number of target utterances. Second, the distributed and distributed-pairs groups did not differ significantly on either of these two measures. Third, the difference between the massed and distributed groups with regard to the total number of target utterances produced only narrowly failed to reach significance at the adopted level of $p = 0.01$, and indeed was significant at the more lenient $p = 0.05$ level. Fourth, children in the distributed group, but not the massed group, produced significantly more target utterances than children in the control group.

It would seem, then, that the observed advantage for the distributed-pairs over the massed condition (as measured by the number of children able to produce a target utterance at test) was a simple consequence of the increased spacing of instantiations of the construction for the former training schedule group. It might be objected that each training trial consisted of a repetition of the same sentence, and so the distributed condition still enabled the child to compare instantiations of the construction within a small time window. But the fact is that these sentences were verbatim repetitions, and so there were not in fact two different instantiations to compare.

Perhaps the most important aspect of the present findings is their methodological implication. In the language acquisition literature it is common to encounter such claims as "the large majority of children under 3 years of age do not use ... verbs in the transitive construction ... [that] they have not heard in that construction" (Tomasello, 2000, p. 222) Such claims, though, are often based on the results of experiments conducted on a massed paradigm, with a large amount of training presented in one or two sessions. For example, Tomasello and Brooks (1998) presented 64 tokens of a novel verb in a single training session, and a further 24 presentations in a second session a few days later. The lesson of the present study is that just because children do not produce a certain linguistic item or structure after a certain amount of training, this is not to say that they could not have done so if the training had been more appropriately distributed. It should be clear from the present study (and those of Childers & Tomasello, 2002; Schwartz & Terrell, 1983) that the key determinant of the level of learning for training studies is not the number of presentations per se, but the number of different days (or, perhaps, sittings) on which at least one presentation is given.

Since the distributed learning effect has been shown to apply to construction learning and to many different types of learning including that outside the linguistic domain (e.g., Edwards, 1917; Menzel, Manz, Menzel, & Greggers, 2001; Shea, Lai, Black, & Park, 2000), another implication of this finding is that learning processes, such as schematization and analogy that have been shown to operate in non-linguistic domains (e.g., Markman & Gentner, 1993; Piaget, 1952) can plausibly be invoked to explain the acquisition of grammatical constructions. Of course, it is entirely possible that linguistic learning is quite different in many respects from other types of learning that benefit from distributed presentation, and we would not wish to claim that our findings provide evidence against the notion that linguistic learning is somehow privileged. The present findings, however, are broadly supportive of accounts under which children acquire language, in the form of a structured inventory of constructions, using skills of learning and categorization that are common to human (and animal) learning in general (e.g., Tomasello, 2003). With regard to Markman and Gentner's (1993) theory, the findings of the present study suggest that, as proposed by Tomasello (2003), the process of structure mapping would appear to be applicable in purely linguistic domains.

However, the data do not support Gentner's specific proposal that comparing contiguously presented exemplars is particularly beneficial for this mapping process. Counter-intuitively, perhaps, it would seem that, for grammatical constructions, children are more able to analogize across exemplars and extract a relational schema when those exemplars are more widely distributed in time than when they are temporally contiguous. It is not clear why this should be the case. One possibility is that the formation of non-linguistic schema and syntactic construction schema are different in some crucial respect, and that the experiments conducted by Gentner and colleagues are not fully analogous to the formation of linguistic constructions. If this is the case, then constructivist theorists must investigate further the ways in which the formation of grammatical constructions is similar to and different from other types of analogy formation.

Another possibility is that the advantage for distributed presentation found in the present study may simply reflect improved learning of the invariant lexical material (*It was the ... that the ...*), particularly given that Challis (1993) demonstrates that the distributed learning effect applies to whole sentences. The analogy-formation process required in the present study would be more similar to that required in the study of Markman and Gentner (1993) if different instantiations of the construction shared no common lexical material. Indeed, in addition to partially abstract constructions, which retain some invariant lexical material between instantiations (e.g., cleft constructions or the *by* passive), children must learn a number of entirely abstract constructions, different instantiations of which may share no lexical overlap. In fact, two such constructions, the *[SUBJECT] [VERB] [OBJECT]* transitive and *[SUBJECT] [VERB]* intransitive constructions, are almost certainly the two most common in adult speech. It would be interesting to see whether the distributed learning effect holds for wholly abstract constructions. Since even prelinguistic children will have had considerable exposure to the wholly abstract transitive and intransitive constructions, such experimental investigations would need to make use of novel (nonce) constructions (Akhtar, 1999; Goldberg, Casenhiser & Sethuraman, 2004) although, of course, such studies have their own associated problems, such as that of ecological validity.

When considered together with the findings of Schwartz and Terrel (1983) and Childers and Tomasello (2002), the findings of the present study are broadly supportive of the view that a single set of general learning and cognitive processes is responsible for the acquisition of both individual lexical items (the lexicon) and regular and irregular grammatical constructions (the grammar) (see Bates & Goodman, 1999; Langacker, 1987; Tomasello, 2003). Under the so-called dual-process theories (e.g., Clahsen, 1999; Pinker, 1991), irregular constructions (such as irregular past tense forms or constructional idioms) are acquired by general learning processes sensitive to effects of frequency, similarity, and so forth, which are also used to learn individual lexical items. Regular constructions (such as regular past tense forms or the transitive construction) are not acquired by means of general learning processes, but are formed using computational rules operating on variables that stand for innately specified syntactic categories and morphemes.

We do not wish to claim, of course, that our findings count against dual-process accounts of language acquisition. It remains entirely possible that word learning and construction learning are crucially different in one or more respects. Our claim is simply that the finding that an extremely robust and well-replicated effect for word learning applies also to construction learning increases the plausibility of so-called single-process accounts, under which lexical items and grammatical constructions are acquired using a single learning mechanism. Indeed, under accounts such as Langacker (1987), word learning and construction learning are conceptualized as a single process, operating on a different scale. In a highly inflected language such as Polish, where every noun consists of a stem plus a morphological marker, each noun can be viewed as a partially abstract (morphological) construction. Furthermore, in agglutinating languages (such as Turkish, Finnish or Hungarian) a single verb form often contains attached elements indicating participants, tense, number and so on, and so may be considered analogous to a syntactic argument-structure construction in English.

The nature and frequency of children's errors often provide important insights into the mechanisms underlying language acquisition. Children's uses of the various non-target construction schemas in this study have implications for a number of specific proposals. For example, utterances in which a child uses a simple object-NP clause or *reduced cleft* construction (*It was the cup*) may provide some support for a *construction-conspiracy* account of language acquisition (Abbot-Smith & Behrens, in press; Croft, 2001; Elman et al., 1996; Goldberg, 1995; Langacker, 1988, 1991, 2000; Ruhland et al., 1995). Under this account, children's acquisition of certain complex grammatical constructions is facilitated by prior knowledge of a number of source constructions which instantiate a part of the target construction. Utterances in the present study in which a child used a simple object-NP clause (or reduced cleft) construction (a mean of 13% of all utterances across both experiments) are relevant to this proposal, as this construction may serve as a source construction for the target object-cleft construction (*It was the cup that the frog took*). Thus, it would appear that some children who had not fully acquired the target construction instead used a shorter, simpler construction that instantiates a part of the target construction.

Subject-cleft constructions, such as *It was the frog that took the cup* (which constitute 6.5% of all utterances across both experiments) are perhaps the most interesting of the non-target constructions produced, as they reflect the fusion of the target object-cleft construction (*It was the cup that the frog took*) with the *SVO* transitive construction, often described as the *default* construction for the type of scene enacted in the present study (Fillmore, 1977; Hopper & Thompson, 1980). When producing a subject-cleft utterance, the child preserves the basic structure of the object-cleft construction but realigns the SUBJECT, OBJECT and VERB slots to conform to their ordering in the default, and, of course, much more frequent, *SVO* transitive construction (see Bever, 1970; Slobin & Bever, 1982, for experimental work along these lines). In the case of subject–object errors (2.5% of all utterances), the child again preserves the structure of the target construction but in this case realigns the SUBJECT and OBJECT slots, but not the VERB slot, to their canonical positions, resulting in an utterance with the AGENT and PATIENT roles reversed (*It was the frog that the cup took*). The presence of such errors then, provides some support for the claim that constructions are learned not in isolation, but as part of a hierarchical network of interrelated constructions which compete for activation (Abbot-Smith & Behrens, in press; Croft, 2001; Elman et al., 1996; Goldberg, 1995; Langacker, 1988, 1991, 2000; Ruhland et al., 1995).

Finally, the rate of production of transitive constructions simply represents the extent to which children default to the preferred or prototypical construction denoting a highly transitive event (Fillmore, 1977; Hopper & Thompson, 1980). The fact that children produced this construction with relatively high frequency (15% of all utterances across both experiments), despite fairly extensive training with an alternative construction, suggests that the argument that children's productions were, in many cases, influenced by a competing default *SVO* construction is plausible.

In conclusion, children acquire abstract grammatical constructions on the basis of the language they hear. The present study has demonstrated that the manner in which this input is temporally distributed significantly affects the learning process in a similar way for the learning of grammatical constructions, individual lexical items, and non-linguistic skills. This finding has implications for our understanding of the cognitive and learning processes that underlie language acquisition. Future research should attempt to specify in more detail both the nature of the input and the learning processes that enable young children to acquire a natural language.

Acknowledgements

We would like to thank the children, parents and teachers of the following schools who took part in the research reported here: Bishop Billsborrow RC Primary School, St Paul's CofE School, St Catherine's RC Primary School, William Hulme's Grammar School and Day Nursery, The Friars School, North Grecian Primary School, Collyhurst Nursery School, St Willabroads RC Primary School, Newall Green Infants School, Ashbury Community Primary School, St Gilbert's School, Haverley Hey Primary School and Christ the king RC Primary School. Our thanks are also due to the following for serving as E2 on various occasions: Eleanor O' Malley, Victoria Hulks, Anna Roby and Evan Kidd, and to Franklin Chang for a particularly helpful methodological insight. This research was funded by a Ph.D Scholarship from the Max Planck Institute for Evolutionary Anthropology, Leipzig, Germany and by an ESRC Postdoctoral Fellowship (PTA-026-27-0705), both to the first author.

Appendix A

See Tables A.1 and A.2.

Table A.1
Verbs used in the experiment

Verb	Experimental phase where used	Frequency (BNC spoken texts)
Take	Training	20878
Hold	Training	2140
Pull	Training	1423
Choose	Training	919
Touch	Training	583
Rub	Training	227
Grab	Training	206
Bite	Training	172
Punch	Training	62
Hide	Training	325
Find	Test	8119
Move	Test	4803
Drop	Test	1060
Kick	Test	445

Table A.2
Utterance categories and scoring criteria

Target object cleft	Demonstrative	BE	NP (PATIENT)		COMP	NP (AGENT)		VERB
	It That	was is 's	the a 0	PATIENT	that what which 0	the a 0	AGENT	VERB (any form, any semantically appropriate verb not presented in this construction)
			pronoun			pronoun		

Object NP clause (reduced cleft)	Demonstrative	BE	NP (PATIENT)	
	It That	was is 's	the a 0	PATIENT
			pronoun	

Subject cleft	*Demonstrative*	BE	NP (AGENT)		COMP	VERB	NP (PATIENT)	
	It That	is was 's	the a 0	AGENT	that what which 0	VERB (any form, any semantically appropriate verb)	the a 0	PATIENT
			pronoun				pronoun	

Transitive	NP (AGENT)		VERB	NP (PATIENT)	
	the a 0	AGENT	VERB (any form, any semantically appropriate verb)	the a 0	PATIENT
	Pronoun			pronoun	

Subject-object error	Demonstrative	BE	NP (PATIENT)		COMP	NP (AGENT)		Verb
	It That 0	was is 's 0	the a 0	AGENT AGENT PATIENT	that what which 0	the a 0	PATIENT AGENT PATIENT	VERB (any form, any semantically appropriate verb)
			pronoun			pronoun		

References

Abbot-Smith, K., & Behrens, H. How known constructions influence the acquisition of new constructions: The German periphrastic passive and future constructions. *Cognitive Science* in press.

Akhtar, N. (1999). Acquiring basic word order: Evidence for data-driven learning of syntactic structure. *Journal of Child Language*, *26*, 339–356.

Bates, E., & Goodman, J. C. (1999). On the emergence of grammar from the lexicon. In B. MacWhinney (Ed.), *The emergence of language*. Mahwah, NJ: Lawrence Erlbaum.

Bever, T. G. (1970). The cognitive basis for linguistic structures. In J. R. Hayes (Ed.), *Cognition and the development of language*. New York: Wiley.

Braine, M. D. S. (1976). Children's first word combinations. *Monographs of the Society for Research in Child Development*, 41(164).

Bybee, J. (1995). Regular morphology and the lexicon. *Language and Cognitive Processes*, *10*(5), 425–455.

Cameron-Faulkner, T., Lieven, E., & Tomasello, M. (2003). A construction based analysis of child directed speech. *Cognitive Science*, *27*, 843–873.

Challis, B. H. (1993). Spacing effects on cued-memory tests depend on level of processing. *Journal of Experimental Psychology: Learning, Memory and Cognition*, *19*, 296–389.

Childers, J., & Tomasello, M. (2001). The role of pronouns in young children's acquisition of the English transitive construction. *Developmental Psychology*, *37*, 739–748.

Childers, J. B., & Tomasello, M. (2002). Two-year olds learn novel nouns, verbs and conventional actions from massed or distributed exposures. *Developmental Psychology*, *38*(6), 867–978.

Clahsen, H. (1999). Lexical entries and rules of language: A multi-disciplinary study of German inflection. *Behavioral and Brain Sciences*, *22*(6), 991–1013.

Croft, W. (2001). Radical Construction Grammar: Syntactic theory in typological perspective. Oxford: Oxford University Press.

Dempster, F. (1996). Distributing and managing the conditions of encoding and practice. In E. Bjork & R. Bjork (Eds.), *Handbook of perception and cognition* (2nd ed.). New York: Academic Press.

Edwards, A. S. (1917). *The distribution of time in learning small amounts of material. In studies in psychology* (209–213). Worchester, MA: Wilson.

Elman, J., Bates, E., Johnson, M., Karmiloff-Smith, A., Parisi, D., & Plunkett, K. (1996). *Rethinking innateness*. Cambridge, MA: MIT Press.

Fillmore, C. (1977). *Topics in lexical semantics*. Bloomington, IA: Indiana University Press.

Gentner, D., Loewenstein, J., & Thompson, L. (2003). Learning and transfer: A general role for analogical encoding. *Journal of Educational Psychology*, *95*(2), 393–408.

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

Goldberg, A. E., Casenhiser, D., & Sethuraman, N. (2004). Learning argument structure generalizations. *Cognitive Linguistics*, *25*(3), 289–316.

Goldstone, R. L. (1996). Alignment-based nonmonotonicities in similarity. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *22*, 988–1001.

Gomez, R. L. (2002). Variability and detection of invariant structure. *Psychological Science*, *13*(5), 431–436.

Holyoak, K. J., & Koh, K. (1987). Surface and structural similarity in analogical transfer. *Memory and Cognition*, *15*, 332–340.

Holyoak, K. J., & Thagard, P. R. (1989). Analogical mapping by constraint satisfaction. *Cognitive Science*, *13*, 295–355.

Hopper, P., & Thompson, S. (1980). Transitivity in grammar and discourse. *Language*, *56*(2), 215–299.

Janiszewski, C., Noel, H., & Sawyer, A. (2003). A meta-analysis of the spacing effect in verbal learning: Implications for research on advertising repetition and consumer memory. *Journal of Consumer Research*, *30*, 138–149.

Kotovsky, L., & Gentner, D. (1996). Comparison and categorization in the development of relational similarity. *Child Development*, *67*, 2797–2822.

Langacker, R. W. (1987). *Foundations of cognitive grammar: Theoretical prerequisites.* Stanford, CA: Stanford University Press.

Langacker, R. W. (1988). A usage-based model. In B. Rudzka-Ostyn (Ed.), *Topics in cognitive linguistics* (pp. 127–161). Amsterdam: Benjamins.

Langacker, R. W. (1991). *Foundations of cognitive grammar. Vol. 2: Descriptive application*. Stanford: Stanford University Press.

Langacker, R. W. (2000). A dynamic usage-based model. In M. Barlow & S. Kemmer (Eds.), *Usage-based models of language* (pp. 1–63). Stanford: CSLI Publications.

Loewenstein, J., & Gentner, D. (2001). Spatial mapping in preschoolers: Close comparisons facilitate far mappings. *Journal of Cognition and Development*, *2*(2), 189–219.

Markman, A. B., & Gentner, D. (1993). Structural alignment during similarity comparisons. *Cognitive Psychology*, *25*, 431–467.

Menzel, R., Manz, G., Menzel, R. M., & Greggers, U. (2001). Massed and spaced learning in honeybees: The role of CS, US, the inter-trial interval and the test interval. *Learning and Memory*, *8*, 198–208.

Onnis, L., Monaghan, P., Christiansen, M. H., & Chater, N. (2004). Variability is the spice of learning, and a crucial ingredient for detecting and generalising in nonadjacent dependencies. In *Proceedings of the 26th Annual Conference of the Cognitive Science Society*.

Piaget, J. (1952). *The origins of intelligence in children*. Norton: New York.

Pine, J. M., Lieven, E. V. M., & Rowland, C. F. (1998). Comparing different models of the development of the English verb category. *Linguistics*, *4*(356), 807–830.

Pinker, S. (1991). Rules of language. *Science*, *253*(5019), 530–535.

Rowland, C. F., & Pine, J. M. (2000). Subject-auxiliary inversion errors and wh-question acquisition: 'What children do know?'. *Journal of Child Language*, *27*(1), 157–181.

Ruhland, R., Wijnen, F., & van Geert, P. (1995). An exploration into the application of dynamic systems modelling to language acquisition. In M. Verrips & F. Wijnen (Eds.), *Approaches to parameter setting. Amsterdam series in child language development* (p. 4).

Schwartz, R. G., & Terrell, B. Y. (1983). The role of input frequency in lexical acquisition. *Journal of Child Language*, *10*, 57–64.

Shea, C. H., Lai, Q., Black, C., & Park, J. (2000). Spacing practice sessions across days benefits the learning of motor skills. *Human Movement Science*, *19*, 737–760.

Slobin, D. I., & Bever, T. G. (1982). Children use canonical sentence schemas: A crosslinguistic study of word order and inflections. *Cognition*, *12*, 229–265.

Theakston, A. L., Lieven, E. V. M., Pine, J. M., & Rowland, C. R. (2001). The Role of performance limitations in the acquisition of verb-argument structure: An alternative account. *Journal of Child Language*, *28*(1), 127–152.

Theakston, A. L., Lieven, E. V. M., Pine, J. M., & Rowland, C. F. (2002). Going, going, gone: The acquisition of the verb "go". *Journal of Child Language*, *29*, 783–811.

Tomasello, M. (1992). *First verbs: A case study of early grammatical development*. New York: Cambridge University Press.

Tomasello, M. (2000). Do young children have adult syntactic competence? *Cognition*, *74*(3), 209–253.

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

Tomasello, M., & Brooks, P. J. (1998). Young children's earliest transitive and intransitive constructions. *Cognitive Linguistics*, *9*(4), 379–395.

Underwood, B. J. (1961). Ten years of massed practice on distributed practice. *Psychological Review*, *68*, 229–247.