UvA-DARE (Digital Academic Repository)

Lexical Acquisition through Acquisition order

Sassoon, G.W.

DOI

[10.1515/9783110308693.99](https://doi.org/10.1515/9783110308693.99)

Publication date

2013

Document Version

Submitted manuscript

Published in

Lexical bootstrapping: the role of lexis and semantics in child language development

[Link to publication](#)

Citation for published version (APA):

Sassoon, G. W. (2013). Lexical Acquisition through Acquisition order. In D. Bittner, & N. Ruhlig (Eds.), *Lexical bootstrapping: the role of lexis and semantics in child language development* (pp. 99-122). (Cognitive linguistics research; No. 50). De Gruyter Mouton. <https://doi.org/10.1515/9783110308693.99>

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

Lexical acquisition through acquisition order

Galit Weidman Sassoon

Abstract

It is often noted in the literature that typical instances of categories denoted by lexical items are acquired earlier than atypical instances. Yet, a line of empirical studies that were left relatively unnoticed provide evidence to an additional generalization. These studies show that early acquired members are often *interpreted as typical* of their category, thereby forming a basis for generalizations concerning the dimensions identifying category members. Based on this evidence, I propose a new intra-domain bootstrapping mechanism, which I call *the Learning Bias*. I propose that the bootstrapping of the meaning of simple and complex category type expressions like *bird*, *tall* and *tall bird*, is based on the order in which their early acquired members are classified. I review existing evidence supporting the hypothesis that the Learning Bias is at work in children of different ages, of both typical and atypical populations, and it remains operative in adults. The Learning Bias affects both linguistic and nonlinguistic categories, therefore being a domain-general learning mechanism. Finally, I propose that the reviewed data indirectly supports the lexical bootstrapping hypothesis. Along the paper, I draw attention to open questions and details, and provide concrete proposals for future research on the topics discussed.

1. Introduction

This paper focuses mainly on the acquisition of semantic knowledge in lexical items. In accordance, my focus is not merely on the number of lexical items a child produces at a given age, but on whether (and to what extent) the child's use of a given lexical item is conventional (see also Chen et al, this volume). Accordingly, while most of the papers in this volume study infants (Chen et al study older – three and four year old – children), this paper reviews findings pertaining to the use of lexical items in children whose age ranges between two and eleven years old.

Lexical items including nouns and adjectives such as *bird*, *apple*, *chair* and *tall*, as well as complex phrases such as *tall bird*, *birds which are pets* and *birds which are not pets*, denote *categories of objects*. The acquisition of their conventional use, then, includes the ability to correctly classify

under them newly encountered objects (Gelman et al, 1991). This ability is acquired gradually: When children have already learnt to correctly classify some entities, they still fail to do so for others. Section 2 provides background on classification and the order in which entities are learnt to be instances of lexical items. In section 3, I propose that the bootstrapping of meaning of simple and complex category-type expressions is often based on the *order* in which their early-acquired members are classified. This paper does not present new experimental results to this effect; instead, it reviews existing empirical evidence supporting the proposal that this intra-domain bootstrapping mechanism (which I call *the Learning bias*) is at work in infants and children. The reviewed evidence also suggests that the same mechanism remains operative in adult learning of new lexical-items' interpretations.

According to my proposal, syntactic knowledge is not a pre-requisite for acquisition of semantic knowledge related to complex category-denoting phrases (such as, e.g., *birds which are not pets*). It is suggested that, actually, if anything, quite the opposite is the case, in line with the lexical bootstrapping hypothesis (LBH). For example, the interpretation of complex expressions like conjunctions, modified nouns and modified verbs (*flying bird*, *sports which are games*, *slowly runs*), is systematically connected to their constituents' meanings (*bird*, *flying*,...), by *intersection* (e.g., the set of flying birds being the intersection of the set of birds and the set of flying things). The learning bias allows learning both simple and complex category meanings. Based on this semantic knowledge, the intersection rule (exemplified above) and similar rules for the interpretation of function words such as *and* and *which* can be extracted. Section 4 presents some evidence for this proposal, and further research is called for, pertaining specifically to learning-order effects in complex phrases.

Section 5 summarizes the main conclusions drawn in this paper and discusses broader implications of the data reviewed.

2. Learning orders

2.1. Background on classification: The graded structure of categories denoted by lexical items

Forty years of research in cognitive psychology show that our categories possess a graded structure (Mervis and Rosch 1981; Lakoff 1987; Markman 1989; Murphy 2002). For example, when speakers are asked to rate

items by the degree to which they exemplify a property like being a *bird*, robins and sparrows rank higher (i.e., they are considered *more typical* or *better examples* of birds) than ostriches or penguins. Similarly, bats rank higher (they are considered *more similar* or *related* to birds) than cows.

These ordering judgments show up in numerous unconscious on-line processing effects. To name one, verification time for sentences like *a robin is a bird* is faster than for sentences like *an ostrich is a bird* where subjects determine membership of a less typical bird (Rosch 1973; Rosch, Simpson and Miller 1976; Armstrong, Gleitman and Gleitman 1983).

Last, but not least, speakers regard certain properties as *typicality dimensions* ('features') of categories. For example, dimensions like *small size*, *flying* and *singing* are considered as typical of birds. The more typical birds are more typical in these dimensions. Their mean degree in these dimensions is higher than that of less typical members or nonmembers (Rosch 1973; Rosch and Mervis 1975; Murphy 2002).

Crucially, the ability to correctly classify newly encountered entities under (the category denoted by) a given lexical item is based on the ability to acquire the category's dimensions. Speakers are more likely to classify under a category the entities that average better in its dimensions (Hampton 1998; Murphy 2002).

2.2. The learning-order effects

Interestingly, robust empirical findings show that the mapping of entities to typicality degrees is tightly coupled with *the order in which items are learnt to be instances of the given category, whether directly or by inference*. Typical instances are acquired earlier. Mervis and Rosch (1981: 97-100) maintain that the learning-order effects are very robust. Rosch (1978) writes:

Rate of learning of new material and the naturally obtainable measure of learning (combined with maturation) reflected in developmental order are two of the most pervasive dependent variables in psychological research (Rosch 1978: 38).

Section 2.2.1 reviews the relevant findings.

2.2.1. The learning-order effects in natural categories

Studies such as Anglin (1977) show that, developmentally, children tend to learn the typical members of natural categories earlier. For example, for the

על הטקסט *Überschrift 2;Chap Autor* שגיאה! השתמש בכרטיסיה בית כדי להחיל 4
 שברצונך שיופיע כאן.

category *bird* this means that birdhood is normally determined first for bird types such as robins and pigeons, later on for chickens and geese, and last for ostriches and penguins. Similarly, non-birdhood is determined earlier for cows than for bats or butterflies. Thus, a normal acquisition order for the category *bird* is highly indicative of its typicality structure (figure 1).

Figure 1. A normal acquisition order for the category bird is highly indicative of the typicality structure.

Anglin (1977, experiment 5.1) has studied the order of classification in 20 children between two years and 10 month old and six years and six month old, and in 10 adults. The stimuli shown to the children consisted of pictures of real objects belonging to four natural categories (*animals, clothing, food* and *birds*). The pictures were selected based on adult ratings of familiarity and typicality of the objects in the pictures with respect to the given categories. The pictures belonged to the following four sets:

- (i) Familiar and typical instances (like cows, horses and cats for the category *animal*)
- (ii) Unfamiliar but typical instances (like wombats, aardvarks and anteaters)
- (iii) Familiar and atypical instances (like ants, butterflies and starfish)
- (iv) Unfamiliar and atypical instances (like crustacean, hydra and centipede).

Let us call *underextension* any case in which a speaker does not identify a category member as such (for example where a child fails to recognize an ant as an *animal*). The instances of each category were positively classified by all the adults as category members, except for 11 underextension responses pertaining to the membership of unfamiliar atypical instances. Yet, the children failed to classify instances as belonging to their target categories in many more cases, thereby exhibiting immature semantic knowledge. The children in the given experiment produced 260 underextension responses, most of which occurred between age 3 to 5. Rates of underextension gradually decreased with age (Anglin 1977: 143). We see that the competence allowing a conventional use of lexical items is acquired gradu-

ally and reaches an adult state not earlier than age 6, even in relatively simple (perceptual) categories.

Furthermore, interestingly, while children produce but few underextension responses on *typical* instances, regardless of whether they are familiar or not, they make many for both familiar and unfamiliar *atypical* category members, as shown in table 1.

Table 1. Total number of underextensions made by children in Anglin (1977: 144).

	Typical	Atypical
Familiar	17	133
Unfamiliar	19	91

There were more underextension responses in the set of atypical familiar instances than in the set of atypical unfamiliar instances. In that respect, children responses differ from those of adults. Yet the familiarity effect among the typical examples did not reach significance. A careful analysis of the children's interviews revealed one factor contributing to the (significant) familiarity-typicality interaction. Occasionally, the children excluded from a category atypical members for which a more specific name was available (like birds who were known to be of the bird type *Kiwi*). This reluctance did not occur in typical members.

Finally, children consistently classified as category members unfamiliar instances that they had never seen before, providing that they were typical (that they exemplified well the category dimensions). Thus, for example, children correctly classified newly encountered typical animals like wombats and anteaters, when they were not able to classify more accessible, yet atypical animals like ants as being *animals*.

On a par with Anglin's (1977) results, Mervis and Rosch (1975) show that children learn the good examples of basic colour categories before learning the poor examples. In addition, Rosch (1973, experiment 4) obtained similar results with different measures. The subjects were 20 children of age 9-11 year old and 24 students. The stimuli consisted of many more natural categories, such as *toy, bird, fruit, sport, crime, vehicle, sickness*, etc. Briefly, Rosch (1973) has measured error rates and speed in categorization judgments. Error rates and speed of subjects verifying categorization statements such as *a hen is a bird* were both inversely related to typicality, in children and adults.

We see that the order in which entities are learnt to be category members correlates with typicality. Section 2.2.2 reviews further evidence for this effect, in artificially construed stimuli.

2.2.2. *Learning-order effects in adults and artificially construed categories*

The use of artificially construed (or otherwise unfamiliar) categories allows testing predictions concerning learning order in adults, as well as controlling for factors such as familiarity, word frequency, and object frequency (in particular, frequency of occurrence of items as category members).

Experiments show that the coupling between typicality and learning order persists in adults. For example, it was found in adult learning of *form* categories in cultures that do not possess them (Rosch 1973: 129, experiment 2). The subjects in this experiment were 94 Dani adults. The learning effects were also found in western adults, who learnt invented categories such as dot patterns and stick figures (Rosch and Mervis 1975, experiments 5-6; Rosch, Simpson and Miller 1976; Mervis and Pani 1980). In all these experiments subjects learnt new categories, and typicality correlated with acquisition order, measured in terms of number of training trials and error frequency. Generally, as typicality increases, the number of errors in classification of items in their target categories reduces. Category learning reaches a criterion (a low enough error rate) earlier in trained typical instances (Mervis and Rosch 1981; Murphy 2002).

In experiment 2, Rosch, Simpson and Miller (1976) controlled for the frequency of occurrence of items in the training sessions. Crucially, participants observed the items which were better in the category dimensions (in categories of letter strings or stick figures) or in overall resemblance to a prototype example (in dot-pattern categories), in fewer training sessions. Typicality ratings, verification time of category membership, order of production and error rate were all correlated with the "category structure" (average in the dimensions or overall resemblance to the prototype), and **not** with frequency of occurrence of an item as a category member.

2.2.3. *Theoretical implications*

Why do we find these learning effects? Recall that, according to the prototype theory, classification in natural categories is a process in which it has to be decided whether the mean typicality degree of an item in the category dimensions is high enough (reaches the classification criterion; Hampton

1998; Murphy 2002). Prototype theories typically assume that the category dimensions are those properties which are frequent within the category and infrequent outside of it (Rosch 1973; Murphy 2002). A child or an adult who possesses enough knowledge about the dimensions may be able to tell that certain items reach the threshold for classification, in which case she can automatically infer that they fall under that category. It would often be easier for a child to tell of a given typical example of the category that it reaches the threshold in that category, than it would be for her to tell of a given atypical exemplar that it reaches threshold. For instance, encountering a hairy, four-legged, tail-wagging, barking creature (i.e., a creature with a high average in the dimensions of *dog*), a child may be able to tell that that creature reaches threshold in *dog*, and therefore, that it is a dog. The same child, on encountering a bald, unusually small, and silent dog (with a low average in the dimensions of *dog*), may not be able to tell that it reaches threshold, and therefore not classify it as a dog. So the child will infer about more typical members of a category that they are members before she can infer about less typical members that they are members. This creates a correlation between typicality and the order in which items are learned to fall under the given category.

Next we will see additional factors that contribute to this correlation.

3. Classification order as a bootstrapping mechanism

3.1. My proposal: The Learning Bias

I propose that in some circumstances the inference direction is inverted: Classification order serves as a hint (a bootstrapping mechanism) for category structure (for learning the set of category dimensions, the classification criterion, and thereby, the set of instances). In a nutshell, I propose that in the lack of knowledge about the dimensions, a child who is taught the category of an object can infer that it is more typical than any other available object whose status in this category is still unknown. When the latter is classified, it is represented as less typical.

Learning a category (the dimensions and classification criterion) is a complicated task, but inferences concerning the typicality of items based on their being early-acquired instances provide straightforward means to this end. Initially, the early classified members are assumed to be best in the category dimensions. Consequently, the category dimensions are assumed to be precisely those properties in which these members rank higher

על הטקסט *Überschrift 2;Chap Autor* שגיאה! השתמש בכרטיסיה בית כדי להחיל
שברצונך שיופיע כאן. 8

than other familiar objects that were classified later. Then, classification of new entities is automatically inferred, based on their average in these presumed dimensions. For example, normally, robins are classified earlier under *bird* and hence are considered more typical birds than chicken, as was previously demonstrated in Figure 1. Consequently, dimensions in which robins average better than chicken (like *small size* and *flying*) are linked to *birds*: Categorization is based on average in these dimensions.

3.2. Empirical evidence for the Learning Bias

Evidence for this learning-bias is formed by experiments showing that *unless the dimensions are directly taught*, acquisition is delayed if early exposure is to atypical items rather than to typical items (Homa & Vosburgh 1976; Goldman and Homa 1977; Mirman 1978; Mervis and Pani 1980). Moreover, some of the studies show that acquisition is delayed even if early exposure is to the whole category in a random order, but not to the typical items first! (Mirman 1978; Mervis and Pani 1980; Hupp and Mervis 1981 and 1982).

These results were obtained for different types of categories, whether linguistic category types (Mervis and Pani 1980) or category types who are not usually denoted by linguistic items like dot patterns (Mirman 1978). This suggests that the learning-order-based strategy for category acquisition is an important domain-general learning strategy. Finally, Hupp and Mervis (1981) have replicated the results in severely handicapped children, showing the advantages of a therapy based on learning orders.

In section 3.3, I describe one of these studies and its theoretical implications in some more detail (Mervis and Pani 1980). In section 3.4 and onwards, I devote my efforts to stating explicitly the ingredients of a process of acquisition based on a learning-bias, thereby pointing at issues for future research.

3.3. Mervis and Pani (1980) – a detailed description

3.3.1. *The study*

The stimuli in Mervis and Pani's (1980) study consisted of 24 items (real objects) belonging to six artificially construed categories of toys, characterized by their form, material and the noises they produced. The categories were designed to have a graded structure like natural categories, such that items' means on a set of dimensions could predict categorization, and some

elements were typical (averaged well in the dimensions), some were atypical (averaged poorly), and some were neither. The subjects in experiment 1 were 20 five year olds (mean age 5;9) and 20 students. The subjects in experiment 2 were 30 five year olds (mean age 5;8) and 30 students.

The initial stage of experiment 1 consisted of two types of conditions. In the good example condition (GE) subjects were taught the category membership of the six *most typical* examples of the categories first, and then they were presented with the remaining items (which were not named for them) one by one in a random order. In the poor example condition (PE) subjects were taught the category membership of the six *most atypical* examples of the categories first, and then they were presented with the remaining items. Experiment 2 included an additional condition (ALL), where subjects were taught the category membership of all the examples one by one in a completely *random order*.

Following this stage, the subjects' production and comprehension abilities were tested. To test production abilities, the subjects were asked to name objects. To test comprehension abilities, the subjects were asked to touch one of the objects in response to a name stated by the experimenter. Finally, to test for *generalization of the category name* (the ability to apply it to newly encountered objects which are "the same kind of things"), items whose names the subjects were not directly taught were tested for production and comprehension. The entire procedure (including the naming stage and the tests for comprehension and production) was repeated until either subjects could demonstrate perfect comprehension and production or for maximum four trials in experiment 1, and five trials in experiment 2.

The main results are the following. Ease of correct generalization of the category name was significantly greater for the good example (GE) condition than for the random-order (ALL) condition, and it was significantly greater for the ALL condition than for the poor example (PE) condition.

These generalizations were assessed in the second experiment by three different measures of ease of correct generalization. The main results of the comprehension tests are presented in table 2. The same results obtained also in the production tests. The measures included number of trials required in order to correctly generalize a name (main effect for condition: $F(2,54) = 18.02$, $p < 0.01$), percentage of correct comprehension-generalizations out of the total number of attempts ($F(2,54) = 19.72$, $p < 0.01$), and number of category names which were correctly generalized on the first attempt ($F(2,54) = 24.25$, $p < 0.01$). Age by condition interactions in the first two measures indicated that the differences between conditions were particularly pronounced for the children (as the task was easier for

על הטקסט Überschrift 2;Chap Autor שגיאה! השתמש בכרטיסיה בית כדי להחיל 10
שברצונך שיופיע כאן.

the adults). In the 'number of trials' measure the ALL condition performed non-significantly better than the GE condition.

Table 2. Number of trials, mean percent correct generalization, and correct first attempt, in comprehension tests (Mervis and Pani 1980: 512-513, exp. 2).

Number of trials	GE	ALL	PE
Children	0.63	1.32	2.30
Adults	0.34	0.20	0.52
Correct generalization	GE	ALL	PE
Children	79.0%	60.9%	32.6%
Adults	97.6%	89.7%	82.2%
First attempt	GE	ALL	PE
Children	3.8	2.2	1.1
Adults	5.8	4.9	4.0

Finally, for subjects who received equal exposure to all category members (the ALL condition), the good examples were learnt to be category members before the poor examples.

In conclusion, initial exposure to good examples results in a more rapid and more accurate (at least at first) category learning, compared to initial exposure to poor examples, or even to the whole category in a random order.

3.3.2. Theoretical implications: The Learning Bias

When children or adults cannot calculate entities' mean degrees on the noun dimensions, acquisition is based on the early acquired entities. More specifically, the choice of a dimension-set is based on the earliest-acquired items.

In the good example (GE) condition, the earliest acquired items are in fact representative (have high means on the actual dimensions), so they make a good basis for generalization, i.e., similarity to them correctly predicts category membership.

In the random-order (ALL) condition, the order of presentation of category members is arbitrary. In order to make a correct generalization, then,

the objects have to be divided to categories based on alternative (frequency-based) strategies, such that each category will be characterized by a set of dimensions which are frequent within it and infrequent outside (maximizing within-category similarity and minimizing between-category similarity; Rosch 1978). This frequency-based procedure may be more difficult or time consuming than the learning-order based strategy, where generalization is simply based on the early acquired items.

Acquisition was more significantly delayed in the poor example (PE) condition. The proposal that acquisition is based on a learning-bias (that generalization is based on the early acquired items) predicts that with early exposure to poor examples, the poor examples provide a wrong cue as to the category dimensions. This happens because poor examples have low means on the actual dimensions, and so they trigger the abstraction of an incorrect dimension set (one in which they indeed have a high mean). At a later stage, the incorrect inferences concerning the dimension-set have to be corrected, and only then an alternative strategy (as in the ALL condition) can be employed. Therefore, category learning in this condition is correctly predicted to be the slowest and least accurate.

The anecdotal evidence provided by Mervis and Pani's (1980) subjects further support these conclusions. For example, one child who had the poor examples named for her, upon generalization, handed the good examples to the experimenter contemplating that it would have been much easier if she had named these objects for her in the first place.

The Prototype Theory, advanced by researchers like Mervis and Pani (1980), has triggered the discovery of these facts. Yet, this theory's probabilistic criterion (the view that dimensions are inferred from their observed frequency within and outside the category, cf. section 2.2.3 above) does not go well with the fact that the typical examples have facilitated acquisition more than exposure to the whole category has. Nor can a 'prior-knowledge' criterion for dimension selection (Murphy 2002) explain these findings, as the subjects in this study possessed no prior knowledge about the dimensions. Only a criterion which states that properties of early-acquired entities are selected for the dimension set, as the Learning Bias proposal predicts, directly explains the data (Sassoon 2006; 2007).

3.4. The Learning Bias - A detailed description

Let us recapitulate. I propose that the generation of typicality orderings for natural categories are governed not only by mean in the category dimensions, but also by learning orders, as stated in (1a-b). I propose that the two

principles in (1a-b) are at work simultaneously. Wherever possible, categorization is based on both of them.

(1) a. *The Prototype-Theory principle:*

The typicality degree of any given entity in a category equals its mean degree in the category's dimensions (Rosch 1978).

b. *The Learning principle:*

The typicality ordering (or alternatively, the ordering between the entities' typicality degrees) reflects the order in which entities are learnt to be category members, whether directly or by inference (Sassoon 2006; 2007).

c. *The Frequency of Occurrence principle:*

The typicality ordering reflects the frequency of occurrence of entities as category members (cf. Nosofsky 1988).

My proposal can best be grasped within an optimality framework (cf. Beaver 2004), within which we can state that speakers attempt to employ several competing strategies, ranked by their relative importance. When the dimensions are known, the Prototype-Theory principle (1a) prevails. When the dimensions are unknown, the Learning principle (1b) prevails. The early acquired items are assumed to be the most typical (best on the dimensions, whatever they are). In accordance, their properties are regarded as dimensions and their values (degrees) on these properties are regarded as the category's ideal values. When early exposure is to an atypical item, the inferred dimensions and ideal values are wrong and need to be corrected later on. This delays acquisition. Finally, when neither information about means in the dimensions nor information about learning order is available (or when information needs to be corrected as will be explained below), other category-learning strategies are employed, based on factors such as frequency of occurrence of entities as category members, as stated in (1c).¹

Section 3.4 demonstrates these proposals with detailed examples. Where possible, it provides descriptions of supporting evidence, and where not, it brings out the issues in question with the goal of advancing future research on the topic.

3.4.1. *Learning from early acquired items, a detailed example*

What exactly happens when, say, Sam's teacher shows her a robin and tells her that this is a bird? Sam can tell that there is an entity in front of her². Yet, this entity is only partially accessible. Sam can see the entity's shape

and colours; she can feel its texture or odour, see the manner of its movement, etc. But Sam may not know many other properties or property-values pertaining to the entity (its weight, what it is made of, how it behaves or functions, what its evolutionary origins are, etc.) Moreover, her teacher cannot tell her *all* the facts about the entities she sees, because the entities are only partially accessible to her teacher, too. But in spite of her partial knowledge, Sam's teacher positively classifies the entity as a bird. For that reason, Sam can presume that classification depends on the dimensions that are accessible to her and her teacher. This means that any of the entity's accessible properties may be part of the dimension-set of *bird*, the entity's values may form the ideal values for birds, and the mean of these values may form the threshold (criterion) for category-membership. Thus, the membership of things that deviate from these ideal values even a bit is still questioned. If Sam is told about several examples of robins simultaneously that these are birds, she can probably eliminate from the dimension-set any of the properties along which these robins differ from one another. They are now known not to be necessary conditions for membership. In addition, they may also be taken not to play a role at all in determining typicality, since items with different values in these dimensions are added to the category earliest, which means that they are supposed to all be best examples (equally good examples). There is little likelihood that all these entities will end up having equal typicality degrees in *bird*, unless all those dimensions in which some of them have lower degrees than others are ignored.

If at a later stage Sam is told that a certain pigeon that she sees is a bird too, she can tell that the threshold values in many potential dimensions that she has in mind are not the tightest possible. Perhaps she can also tell that properties in which the pigeon scores better than the robin are not bird dimensions or are dimensions with low weights. How can she tell that, for instance, cooing like a pigeon is not a bird-dimension? For all she knows, most birds may coo. After all, even if the learning principle tells her that robins are more typical birds than pigeons are, it remains possible that in that particular dimension a pigeon scores better than a robin. Still, the weight for this dimension should be low enough, so as not to render pigeons more typical birds than robins. We see that it is not always logical inferences that the child makes. She may just take guesses that are in line with what she knows, and learning orders may strongly bias these guesses. Dimensions of pigeons may be eliminated simply because robins are the earliest-acquired members and so their dimensions are favoured. These issues call for a future experimental investigation, which will flesh out in more detail the type of inferences triggered by the Learning Bias.

After a while, Sam can obtain a partial set of *bird* dimensions (and negative dimensions, i.e. properties that she positively classifies as not being *bird* dimensions), ideal dimension-values (and negative values, i.e. values that she positively classifies as not ideal), a set of potential weights for the dimensions (and sets of negative weights, i.e. numbers that cannot form these weights), and negative threshold-values (values that she already positively classifies as not forming the threshold for membership in *bird*). These partial sets can adequately predict many facts about membership of entities in *bird*. Sam may remain uncertain only concerning the membership of entities with low rates in the dimensions she obtains.

3.4.2. Error correction: Contradictory inferences delimit the learning bias effects

What happens if, for example, Sam knows nothing about the characteristics of birds, and her initial exposure to birds is through ostriches? According to the Learning Bias (and the review in 3.3), she will classify items by their similarity to ostriches. That is, she will think that the ostrich is a representative bird, and that its known dimensions and dimension-values (*running*, *ostrich size*, etc.) are the dimensions and ideal values of the category *bird*.

But Sam is not doomed to "remain in the cave" for ever. Mean functions are such that, all other things being equal, an increase in one of the values increases the mean. Thus, the proposal that nominal categories are linked with mean functions (the prototype-theory principle, (1a)) predicts that (all other things being equal) the more typical an entity is of a certain dimension (say, *flying*), the more typical this entity is of the category (*bird*). Thus, in initial exposure to ostriches, Sam will expect that (all other things being equal), entities which are, say, more typical runners (given that *running* is a feature of ostriches), will be more typical *birds*. These inferences will be cancelled later on, when non-birds (or poor examples of birds) will be discovered to be better on this dimension than (good examples of) birds. Sam is taught about many flying and non-running animals that they are birds. She may, initially, wrongly consider them atypical birds. But she is also taught about many running and non-flying animals that they are not birds. These non-birds are more similar to the examples that she wrongly considers typical birds than to the examples that she wrongly considers atypical, in these, and maybe also other, dimensions (and consequently, in their weighted means). This eventually *forces* her to abandon her initial assumption that ostriches are the prototypical birds.

The role of frequency in such cases is discussed below. Deference may enhance this decision too, if indeed Sam's language community considers ostriches nonrepresentative.

In sum, we see that the learning bias is a potent mechanism, but it is not too strong. In light of contradictory inferences, wrong inferences based on learning orders are discarded. But this process slows down acquisition. In fact, for some children in Mervis and Pani's (1980) study, early exposure to atypical members completely blocked acquisition, at least within the time given to them during the experiment. They refused to abandon inferences that were based on the early-acquired members.

Further evidence that corrections of earlier assumptions take place following early exposure to atypical examples is formed by the fact that sometimes adults for whom poor examples were explicitly named ultimately excluded these objects from the generalized category (Mervis and Pani 1980).

3.4.3. The role of frequency of occurrence as a category member

Frequency of occurrence (e.g., principle (1c)) affects acquisition when it comes to error corrections. Corrections occur when one is faced with two contradicting hypotheses about the facts. For example, if one has been exposed to the bird-hood of ostriches earlier than to the birdhood of pigeons or chicken, and at the same time has discovered that the latter score higher in the set which one presumes to be the set of typicality dimensions (e.g. flying, perching, etc.³), what would one infer? Would one eliminate the problematic dimensions from the dimension-set and infer that indeed ostriches are more typical than pigeons or chicken, or would one leave the dimension-set as is, and infer that ostriches are not more typical?

In such cases, considerations other than the learning order may be called for. For instance, if one encounters birds which are similar to pigeons or chicken exceedingly more frequently than birds which are similar to ostriches, the latter (the assumption that the ostrich is not more typical than pigeons or chicken) might seem more attractive. However, the situation is different if one observes that the 'problematic' dimensions are significantly less frequent in birds than in other categories (mammals, insects, etc.) In such a situation the former (the assumption that in fact the ostrich is more typical) may seem more attractive. In this way, frequency does play a role, albeit a less central one, in the determination of typicality judgments.

Evidence for the smaller importance of frequency considerations compared to learning-orders comes from Hupp and Mervis's (1982) study of

severely handicapped children. Training based on one or more typical items resulted in significantly more accurate generalization than training with a range of items. In fact, the latter did not result in generalization above chance levels. Still, training based on multiple typical items tended to lead to more accurate generalization than training based on a single typical item.

My hypothesis about the role of frequency is reminiscent of Markman's (1989: 215) view of the role of frequency in the acquisition of novel predicates. For example, Markman suggests that children assume a principle of *mutual exclusivity*, according to which each object has but one label. When encountered with a novel label but no novel object which it may label, mutual exclusivity may cause the child to infer that the label is related to a part of the object, or to one of its dimensions (colour, size or so on). However, Markman does assume that mutual exclusivity may be abandoned and a second label accepted, when hearing a second label (*animal*) applied repeatedly to an object with a known label (*bird*).

Thus, when the dimension-set isn't directly taught, the learning-order bias is the default strategy. When no learning-order cues are available that produce consistent concept structures (e.g. in PE and ALL conditions), the learning-order bias is abandoned. For further discussion of the role of probabilistic criteria in categorization see Sassoon (2007, chapters 2 and 4.3).

3.4.4. The final stage: Classification of typical entities by inference

Once the dimensions and selected values are set, they are used in order to infer facts about membership of new items. At this stage, the *actual* order in which items are learnt to be members becomes irrelevant. This does not mean that the learning principle ceases to apply. What happens is that a newly-encountered item that is good enough in the dimensions is *automatically regarded as a member*, and it is regarded as an *already known* member. Let us see evidence for these effects.

Studies of damaged neural network simulations and of aphasic patients (Kiran and Thompson 2003 and references therein), show that **following training with sets of typicality dimensions**, exposure to *atypical* items results in spontaneous recovery of categorization of untrained *typical* items, but *not vice versa*; exposure to typical items *does not* result in recovery of categorization of untrained atypical items. We see that the membership of typical instances is indirectly acquired earlier – Typical items are good in the known dimensions so they are automatically classified,

while other items are not (cf. the pattern of children's underextension responses reported in Anglin 1977). Crucially, when the category dimensions are learnt or taught, the Prototype-Theory principle (1a) prevails.

Yet, the Learning principle does not cease to apply. In healthy adults, often typical items which are seen for the first time are falsely thought to already be known (Reed 1972; 1988). For example, participants presumed to have identified criminals in a line-up, who in truth they never saw before, only because they obtained characteristic dimensions of the given category of criminals. Why? Since the newly encountered items were typical (good in the dimensions), and since, according to the learning principle, typical items are classified relatively early, once encountered – typical items are presumed to already be known (classified).

Finally, in addition to their early acquisition and their importance in triggering inference, speakers also remember best the typical instances. They are most likely to be listed from memory, and their dimensions affect future remembrance of new entities and their dimensions (Heit 1997). For example, when speakers are initially exposed to *joggers* that *wear expensive running shoes*, they frequently falsely recollect joggers that do not wear expensive shoes as *non-joggers* or as *joggers that do wear expensive shoes*. In this case, new facts were corrected so as to match ones based on early-acquired entities.

To conclude, when subjects are asked which stimuli they have seen previously, percentage of false recognition responses and degree of confidence in seeing the stimuli are both correlated with degree of typicality (Mervis and Rosch 1981). This is directly predicted by the proposal that speakers attempt to generate a typicality ordering satisfying both (1a) and (1b) (reflecting mean in the dimensions as well as learning order). Speakers 'pretend', so to speak, that newly encountered entities that are good in the dimensions are early-acquired.

4. Classification order in complex expressions

According to my proposal syntactic knowledge is not a pre-requisite for semantic acquisition of complex category denoting phrases, such as, e.g., *non-birds*, *pet birds* and *birds which are not pets*. The data actually suggests the opposite, in line with the lexical bootstrapping hypothesis (LBH), and more specifically, the idea of emergence of the grammar from the lexicon (Bartsch, this volume).

4.1. My proposal: The leaning bias applies to complex concepts

Consider, for example, the interpretation of complex expressions like negated nouns (*non-birds*, *animals which are not birds*), modified nouns (*flying bird*, *sports which are games*), disjunctions (*sports or games*) and modified verbs (*slowly runs*). I propose that, initially, the interpretation of such expressions is extracted using the same means as in the bootstrapping of interpretation of basic lexical items (the Learning Bias). When enough knowledge about simple and complex expressions is accumulated, the systematic connections between the meaning of the complex expressions and the meanings of their constituents (*not*, *or*, *bird*, *flying*, ...) can be extracted.

For example, the fact that the set of members of a negated category like *not-a-bird* is the *complement* of the set of members of the (non-negated) category *bird*, constitutes the core of the semantics of the functional word *not* in phrases like *animals which are not birds*. This fact can be extracted from knowledge about the meanings of pairs of negated and non-negated categories (Chierchia and McConnell-Ginnet 2002). According to this scenario, the availability of semantic knowledge is independent of syntactic rules; as such, it can, perhaps, even assist the extraction of the later.

Similarly, the fact that the set of members of a conjunctive category like *pets which are birds* is the *intersection* of the set of members of the constituent categories (*pets* and *birds*) constitutes the core of the semantics of the functional words *which are*. The same holds true of the functional word *and* in phrases like *a pet and bird*, and of the adjectival modification construction in phrases like *pet birds*. Again, we see that the availability of semantic knowledge is independent of, but most probably can assist the extraction of, syntactic rules.

It remains for future research to directly employ Anglin's (1977) and Mervis and Pani's (1980) paradigms with stimuli consisting of complex phrases, so as to support or refute the hypotheses submitted in this section. For now, let us review existing evidence that indirectly supports these hypotheses.

4.2. Evidence for a learning-based acquisition of complex concepts

Evidence for the use of the learning-order bootstrapping mechanism in complex expressions is formed by the fact that modified nouns (or the categories they denote) are linked with new dimensions which are related neither to the modifier, nor to the noun (Hampton 1997; Murphy 2002).

For example, *pet-birds* are characterized by properties like *lives in cages* and *can talk*, which are neither typical of pets nor of birds alone. Similarly, *small-spoons* are typically *made of metal* and *large-spoons* are typically *wooden*. *Boiled-eggs* are *hard* whereas *boiled-potatoes* are *soft*. None of these dimensions characterizes any of the separate constituents. The same phenomenon (emergent dimensions) characterizes negated categories and disjunctions, too.

The emergent dimensions are usually viewed as refuting the idea of compositionality for dimension-sets, because they derive from experience with category members, rather than being logically entailed by Boolean composition rules for dimension-sets (Hampton 1997). On the present proposal, the emergent dimensions are understood as evidence for the bootstrapping of independent representations (dimension-sets) for complex expressions. The construction of a dimension-set for a complex expression is not fully based on the dimension-sets of the constituents (it is not fully compositional), because it is based on a productive bootstrapping strategy, the Learning Bias, *which is available for any type of expression*.

Sassoon (2006; 2007) shows that the learning principle helps to explain other empirical effects related to typicality and categorization in complex predicates, such as *the conjunction fallacy* (Tversky and Kahneman 1983). More generally, studies of complex expressions support the hypothesis that the characterization of complex categories (in terms of category structure and processing) is similar to that of lexical categories; examples include Barsalou (1983), who studied complex phrases like *food not to eat in a diet* and *things to take from home in a case of a fire*, and Smith et al (1988), who studied complex phrases like *nonred fruit*.

Finally, the present proposal predicts that delays in acquisition of lexical categories will result in delays in bootstrapping the meaning of function words denoting logical connectives such as *and*, *or* and *not*, as well as of syntactic constructions such as the adjective-modification construction. This prediction is derived in virtue of the fact that in the given proposal, the acquisition of the semantics of function words and syntactic structures depends on the acquisition of the semantics of open-class lexical items. The latter form a prerequisite for the bootstrapping of meaning of the former. See Bittner 2010 (this volume), for further discussion supporting the view that knowledge of closed class semantics is a prerequisite of grammatical development.

5. Concluding remarks

In this paper, the graded structure of categories is characterized as reflecting the order in which entities are learnt to be category members, either directly or by inference. This characterization is supported by developmental trends and by experimental studies of artificially construed categories. The learning principle directly predicts a variety of learning-order effects.

There are compelling reasons (appraised in sections 2 and 3) to believe that the learning order mechanism is equally available for infants, children and adults. Adults may utilize this bootstrapping mechanism also as listeners within a discourse, when they need to disambiguate the utterances' context dependent interpretation, or, in other words, when they acquire the "speaker-meaning". As such, my proposal can be viewed as supporting the homotypic continuity approach (Kagan 1971; Bates et al 1997).

Furthermore, while the Learning Bias is in essence an intra-domain bootstrapping mechanism (order of acquisition of category members influences category learning), in section 4 I proposed that it also triggers inter-domain bootstrapping, linking lexico-semantic and syntactic development. Evidence about emergent dimensions, conjunction fallacies, and order of acquisition of open- and closed-class words suggests that the Learning Bias affects the acquisition of complex phrases, facilitating acquisition of conventional uses of some closed-class words, and thereby of syntactic constructions.

Finally, the learning-bias proposal presupposes that there are tight connections between language and cognition, in that learning the core part of lexical meaning is reduced to conceptual learning (learning about categorization of the world). Sassoon (2007) uses a formal model representing the gradual growth of information about word meaning. She shows that in such a model it is hardly possible to distinguish between world-knowledge and linguistic semantic knowledge, in accordance with the fact that the Learning Bias is at work in both linguistic and nonlinguistic categories (cf. section 3.2).

Notes

1. Homa and Vosburgh (1976) and Goldman and Homa (1977) have failed to find an advantage for GE over ALL conditions. Subtle experimental design may affect the choice of strategy, biasing towards the use of either learning orders or frequency-based strategies. At any rate, several other studies did find

- an advantage for GE over ALL conditions (Mirman 1978; Mervis and Pani 1980; Hupp and Mervis 1981; 1982), establishing the importance of the leaning-bias.
2. At this initial stage, an ability to learn by direct teaching is required, including, for example, the ability to correctly identify the entity who is being named among all possible entities in the scene. This competence is studied by Markman (1989). My paper focuses on the next stage, namely, the one involving indirect learning (generalization to newly encountered entities). The Learning Bias is at work once the first member is classified. It allows generalizing.
 3. These dimensions can be inferred based on earlier exposure to, e.g., robins.

References

- Anglin M. Jeremy
1977 *Word, object and conceptual development* New York: W.W. Norton.
- Armstrong, Lee, Lila Gleitman, and Henry Gleitman
1983 What some concepts might not be. *Cognition* 13: 263-308.
- Barsalou, Lawrence
1983 Ad hoc categories. *Memory and Cognition* 11: 211-227.
- Bates, E. and Goodman J.C.
1997 On the inseparability of the grammar and the lexicon: Evidence from acquisition, aphasia and real time processing, *language and Cognitive processes* 12(5/6): 507-584.
- Bartsch, Susanna
2010 The Lexical Bootstrapping Hypothesis. In: Dagmar Bittner and Nadja Kuehn (Eds.) *Lexical bootstrapping, on the central Role of Lexis and Semantics in Child Language development*, Mouton De Gruyter.
- Beaver, David
2004 "The Optimization of Discourse Anaphora", *Linguistics and Philosophy* 27(1): 3-56.
- Bittner, Dagmar
2010 Closed class words as a prerequisite of grammatical development. In: Dagmar Bittner and Nadja Kuehn (Eds.) *Lexical bootstrapping, On the central Role of Lexis and Semantics in Child Language development*, Mouton De Gruyter.
- Chen, Ping, Lauren Tonietto, Maria A. Parente, Karine Duvignau, Bruno Gaume,
2010 The Lexical Bootstrapping Hypothesis and conventionality: The contribution of a cross-linguistic study in verb acquisition by Chinese- and Portuguese-learning children. In: Dagmar Bittner and Nadja Kuehn (Eds.) *Lexical bootstrapping, on the central Role of*

על הטקסט Überschrift 2;Chap Autor שגיאה! השתמש בכרטיסיה בית כדי להחיל
 שברצונך שיופיע כאן.

- Lexis and Semantics in Child Language development*, Mouton De Greuyter.
- Chierchia, Gennaro and Sally McConnell-Ginet
 2000 *Meaning and grammar. An introduction to Semantics*. 2nd edition. MIT Press, Cambridge, MA.
- Gelman, Susan A. & James P. Byrnes,
 1991 *Perspectives on language and thought: interrelations in development*, Cambridge University Press
- Goldman, D. and Homa, Dunbar.
 1977 Integrative and metric properties of abstracted information as a function of category discriminability, instance variability, and experience *Journal of Experimental Psychology: Human Learning and Memory* 3: 375-385.
- Hampton, James,
 1997 Conceptual Combination. In Koen Lamberts and David Shanks (eds.), *Knowledge, Concepts and Categories*. The MIT Press. Cambridge, MA: 135-162.
- 1998 Hampton, James (1998), Similarity based categorization and fuzziness of natural categories. *Cognition* 65: 137-165.
- Heit, Evan
 1997 Knowledge and concept learning. In: *Knowledge, Concepts and Categories*, Koen Lamberts and David Shanks (eds.), 135-162. Cambridge, MA: The MIT Press.
- Homa, Dunbar and Vosburgh, R.
 1976 Category breadth and the abstraction of prototypical information. *Journal of Experimental Psychology: Human Learning and Memory* 2: 322-330.
- Hupp, Susan C., and Mervis, Carolyn B.
 1981 Development of generalized concepts by severely handicapped students. *Journal of the Association for the Severely Handicapped (JASH)*, volume 6/1: 14-21.
- 1982 *Child Development*, Vol. 53, No. 3: 760-767.
- Kagan, J.
 1971 *Change and continuity in infancy*. New York: John Wiley and Sons.
- Kiran, S. and Thompson, C. K.
 2003 The role of semantic complexity in treatment of naming deficits: Training categories in fluent aphasia by controlling exemplar typicality, *Journal of Speech Language and Hearing Research* 46: 608-622.
- Lakoff, George
 1987 *Women, Fire and Dangerous Things: What Categories Reveal about the Mind*. Chicago University Press.
- Markman, M. Ellen
 1989 *Categorization and Naming in Children, Problems of Induction*. Chapter 9: 187-215. Cambridge, Massachusetts: The MIT Press.

על הטקסט שברצונך Überschrift 1; Chap No השתמש בכרטיסיה בית כדי להחיל
 שיופיע כאן. 23

- Mervis Carolyn B. and J.R Pani
 1980 Acquisition of basic object categories, *Cognitive Psychology* 12: 496-522.
- Mervis Carolyn B., and E. Rosch
 1981 Categorization of natural objects, in: *Annual review of psychology* 32, 89-115.
 1975 Development of the structures of colour categories. *Developmental psychology* 11: 54-60.
- Mirman, M.C.
 1978 The effect of number and distortion level of category exemplars on concept abstraction. BS thesis. Uni. Of Illinois, Urbana Champaign.
- Murphy, L. Gregory
 2002 The big book of concepts. Massachusetts: The MIT Press.
- Nosofsky, Robert, M.
 1988 Similarity, frequency, and category representation, *Journal of Experimental Psychology, Learning, Memory and Cognition* 14: 54-65.
- Reed, Stephen K.
 1972 Pattern recognition and Categorization. *Cognitive Psychology* 3: 382-407.
 1988 *Cognition, Theory and Applications*, Brooks/Cole, Pacific Grove.
- Rosch, Eleanor
 1973 On the internal structure of perceptual and semantic categories. In: *Cognitive Development and the Acquisition of Language*, T. E. Moore (ed.), New York: Academic Press.
 1975 Cognitive Representations of Semantic Categories, *Journal of Experimental Psychology: General* 104, 192-232.
 1978 Principles of categorization. In Eleanor Rosch & Barbara B.Lloyd (Eds.), *Cognition and Categorization*, N.J.Erlbaum, Hillsdale: 17-48.
- Rosch, Eleanor, and Carolyn Mervis
 1975 Family Resemblance: Study in the Internal Structure of Categories. *Cognitive Psychology* 7: 573-605.
- Rosch, Eleanor, Carol Simpson and R. Scott Miller
 1976 Structural Bases of Typicality Effects. *Journal of Experimental Psychology: Human perception and performance* 2: 491-502.
- Sassoon, W. Galit
 2006 The logic underlying the typicality effects, *Proceedings of Sinn und Bedeutung* 11, Pompeu Fabra University, Barcelona.
 2007 Vagueness, Gradability and Typicality, A comprehensive Semantic analysis. Unpublished PhD dissertation, Tel Aviv University.
- Smith, Edward, Daniel Osherson, Lance Rips, Margaret Keane
 1988 Combining prototypes, a selective modification model. *Cognitive Science* 12: 485-527.
- Tversky, Amos & D. Kahneman,

על הטקסט *Überschrift 2;Chap Autor* שגיאה! השתמש בכרטיסיה בית כדי להחיל 24
שברצונך שיופיע כאן.

- 1983 "Extensional Versus Intuitive Reasoning: The Conjunction Fallacy in Probability Judgments", *Psychological Review* 90, 293-315.
- Wittgenstein, Ludwig,
- 1968 Reprint: *Philosophical Investigations*, Translated By G.E.M Anscombe, 3rd edition, Oxford: Blackwell. Original edition, 1953.