A Top-Down Approach to Lexical Acquisition and Segmentation

Helen Gaylard **Peter Hancox**

School of Computer Science

University of Birmingham

Abstract

A major objection to top-down accounts of lexical recognition has been that they are incompatible with an account of acquisition, it being argued that bottom-up segmentation must precede lexical acquisition. We counter this objection by presenting a top-down account of lexical acquisition. This is made possible by the adoption of a flexible criterion as to what may constitute a lexical item during acquisition, this being justified by the extensive evidence of children's under-segmentation. Advantages of the top-down account offered over the bottom-up alternatives are that it presents a unified account of the acquisition of a lexicon and segmentation abilities, and is wholly driven by the requirements of comprehension. The approach described has been incorporated into an integrated model of acquisition processes, the incremental learning of which captures the gradual nature of child language development.

Introduction

We discuss top-down and bottom-up approaches to lexical recognition and the bottom-up approach to lexical acquisition. A top-down account of lexical acquisition can be seen to be required both to complete the top-down approach and to overcome the inadequacies of previous bottom-up accounts of lexical acquisition. We present such an account and describe its implementation in a computational model of child language acquisition. Learning in this model is used to illustrate that, as well as providing a unified framework for the acquisition of a lexicon and segmentation abilities, the top-down account of lexical acquisition suggests explanations for some of the observed features of child language development.

Top-Down *versus* Bottom-Up Approaches to Lexical Recognition

In bottom-up accounts of lexical recognition (e.g., Grosjean & Gee 1987; Cutler & Mehler 1993), segmentation is guided by prosodic cues and precedes lexical lookup. The major alternative to this is the top-down approach (e.g., Cole & Jakimik 1980; Tyler & Marslen-Wilson 1982; McClelland & Elman 1986) in which segmentation is predicted from knowledge of possible lexical items. The latter approach has also been termed “postlexical”, but we avoid this term since there is no necessary connection between the top-down use of lexical information and a naïve, strictly left-to-right recognition strategy. Here, we briefly examine the relative strengths and weaknesses of the alternative approaches.

According to the bottom-up approach to lexical recognition, lexical lookup is triggered, bottom-up, by cues such as stressed syllables (in the case of English). This idea is supported by experimental evidence on the role of stress in segmentation (e.g., Cutler & Norris 1988, Cutler & Butterfield 1992). Furthermore, it has been argued, on the basis of a comparative study of lexical access strategies (Briscoe 1989), that the constraints provided by stressed syllables are necessary to keep to a reasonable number the lexical candidates considered.

One argument against the bottom-up approach is that it sacrifices the goal of achieving accurate lexical recognition to that of constraining lexical access:

> "While function and content words have metrical characteristics, the distribution of such words is controlled by syntax. Any prelexical strategy for characterizing words which has as its strength the fact that it is autonomous will have as its weakness the fact that it fails to use the appropriate higher-level information."
>
> (Bard 1990, p.204)

Related to this is the focus upon the problem of segmentation and the failure to provide a unified account of segmentation and lexical access. The arguments in favour of the necessity of constraining lexical access are, anyway, weak, since humans entertain a large number of incorrect lexical hypotheses (Shillcock 1990). Furthermore, while Briscoe (1989) rules out lexical lookup triggered by syllables (as opposed to *stressed syllables*), these have been proposed as the lexical access unit for French (Cutler & Mehler 1993).

A further criticism of the bottom-up approach relates to the language-specificity of the cues to segmentation proposed. This implies that segmentation strategies, based upon language-specific distributional information, need to be acquired prior to the acquisition of, and thus in the absence of, a lexicon. These issues are discussed in a separate section below.

According to the top-down approach to lexical recognition, lexical access precedes, and provides the basis for, segmentation. We use the term "top-down" to include interactionist approaches, such as the TRACE model (McClelland & Elman 1986), in which competition amongst lexical items utilises lower-level, phonemic information as well as higher-level knowledge of lexical items. While in the bottom-up approach segmentation is viewed as a separate process preceding that of lexical lookup, the strength of the top-down approach is that the mechanism underlying lexical access is simultaneously responsible for segmentation (Bard 1990). This is the reason why the top-down approach is able to utilise the kinds of information implicated by the bottom-up approach and thus account for the *appearance* of a metrical segmentation strategy. It has been argued that a model like TRACE will naturally exploit the relative *intelligibility* (Bard 1990) and *informativeness* (Altmann 1990) of stressed syllables.

The major weakness of top-down models of lexical recognition has been the lack of an associated account of lexical acquisition in children. It has been argued that, while top-down approaches to recognition rely upon the lexicon, the acquisition of the lexicon itself presupposes segmentation (Mehler *et al* 1990; Cairns *et al* 1994). Acquisition of the lexicon from isolated words in the input is not regarded as plausible, since function words, for instance, are not used in this way (Jusczyk 1993). Below we discuss how the top-down approach may be extended to incorporate an account of acquisition.

The Bottom-Up Approach to Lexical Acquisition

The bottom-up approach to acquisition can be summarized as the proposal that segmentation abilities, which precede lexical acquisition, are *bootstrapped* on the basis of prosodic/suprasegmental and phonotactic/segmental information in the language:

> "We have suggested that it may be the case that the characteristic pattern of a language is sufficiently salient to assist the newborn child in segmenting the continuous speech stream into discrete units."
>
> (Cutler & Mehler 1993, p.105)

Prosodic information is viewed as useful in acquisition due to the correlations which exist between prosodic units and syntactic or lexical units. Phonotactic information is viewed as useful at a lower level where knowledge of legal and illegal phoneme clusters can be used to distinguish phonemes within the same unit (syllable or word) from those which belong to different units. Below we evaluate accounts of the roles of each of these kinds of information in acquisition.

There is empirical evidence to support the claim that infants are sensitive to correlations between prosody and syntax in "motherese", with sensitivity to clausal units developing at around 6 months, and, to phrasal units, later, at around 9 months (Hirsh-Pasek *et al* 1987; Nelson *et al* 1989; Jusczyk *et al* 1992). Accounts of lexical acquisition based upon prosody appear problematic, however, in that they require the assumption that sensitivity to word and syllable boundaries follows sensitivity to phrasal boundaries, in the same way that the latter follows sensitivity to clausal units. The hypothesis that the recognition of lexical and syllabic boundaries has a prosodic, rather than, for instance, a phonotactic, basis is thus one that must be treated with caution. A further difficulty with prosodically-based accounts of lexical acquisition is that they require, not only that syllable boundaries be recognised, but also that syllables, as units with a special status in lexical acquisition (Mehler *et al* 1990), be recognised as such.

Infants' preferences for legal over illegal phoneme clusters provide evidence in support of the hypothesis that they are sensitive to phonotactic as well as prosodic information (Friederici & Wessels 1993). A number of computational models have been developed which use statistical analyses to simultaneously acquire phonotactic knowledge about a language and use this in segmenting the input into syllabic and lexical units (Wolff 1988; Cartwright & Brent 1994; Cairns *et al* 1994). In the absence of a lexicon, segmentation works on the assumption that frequent sequences of phonemes are likely to be word-internal, whereas infrequent phoneme sequences are likely to indicate word or syllable boundaries. Cartwright and Brent (1994) find that performance in segmentation is optimised when both these kinds of information are used in the analysis of child-directed speech. The advantage for child-directed speech is attributed to the large number of repetitions it contains, e.g.,

> "Do you see the kitty? See the kitty? Do you like the kitty?"
>
> (Cartwright & Brent 1994, p.2)

An interesting result of the work of Cairns *et al* (1994) is the suggestion that the *appearance* of a role for prosody in lexical acquisition may emerge from a model which uses only lower-level kinds of information. Input to the model described is represented accurately by a complex matrix of sub-phonemic features. The model tends to place boundaries before strong rather than weak syllables, as predicted by metrically-based accounts.

Phonotactic approaches appear to provide the most promising basis for a bottom-up account of lexical acquisition since they suggest how language-specific segmentation strategies may be bootstrapped on the basis of the input alone. There is, however, a general weakness which all bottom-up approaches to acquisition share with the bottom-up account of lexical recognition. This is that they focus upon the acquisition of segmentation while paying insufficient attention to issues in lexical acquisition, thus failing to provide a unified account of these processes. It remains to be demonstrated that a purely bottom-up approach to acquisition can yield the accuracy in segmentation required, and it further remains to be shown how meanings are to be attached to the syllabic and lexical units resulting from this process.

The Top-Down Approach to Lexical Acquisition

The top-down approach to lexical recognition can be seen to have a number of advantages over the bottom-up approach. It presents a unified account of lexical access and segmentation and is able to make use of lower- as well as higher-level kinds of information. In order for the top-down approach to be shown to be adequate, however, requires that a top-down account of lexical acquisition be given. We outline such an account below.

It has been assumed that segmentation must precede lexical acquisition and, thus, that any account of acquisition will be bottom-up:

> "it is difficult to reconcile the interactionist approach with the development of segmentation since a lexicon is presupposed."
>
> (Cairns *et al* 1994, p.4)

This assumption evinces an instantaneous view of acquisition in which segmentation of the input into the units in the adult lexicon precedes lexical acquisition. Taking into account the gradual nature of child language development, it seems more likely that lexical acquisition and segmentation are incremental. This suggests that the first lexical items acquired need not correspond to adult lexical items. If input utterances may constitute the first lexical items acquired, then segmentation need not precede lexical acquisition. The adoption of a flexible criterion as to what may constitute a lexical item during acquisition forms the basis of our top-down account.

There are a number of considerations in favour of the assumption that utterances may constitute the first lexical items acquired. The rote, unproductive appearance of the earliest utterances produced by children suggests that segmentation does not precede the acquisition of the earliest lexicon and grammar. A gradual view of segmentation also appears to be required in order to account for a period of correct usage of functional morphemes preceding the onset of functional morpheme omission and overregularization (Gaylard 1995). Furthermore, the assumption that a lexical entry in acquisition may subsume a number of words is consistent with one of the observed features of child language development, which is the extensive evidence of under-segmentation:

> "The first units of language acquired by children do not necessarily correspond to the minimal units (morphemes) of language described by conventional linguistics. They frequently consist of more than one (adult) word or morpheme."
>
> (Peters 1983, p.89)

The finding that sensitivity to clauses precedes sensitivity to phrases is also consistent with the suggestion that children's earliest lexical units may consist of unanalysed utterances.

The incremental, top-down account of lexical acquisition may be summarised as follows. The earliest lexical items acquired consist of unanalysed utterances. These utterances may be of various grammatical types, in keeping with the finding that child-directed utterances include sentences, phrases (especially noun phrases), and isolated nouns (Newport 1977). Lexical recognition uses existing lexical items where possible, with lexical acquisition being triggered by failure of recognition. For example, if "the cat" is input, and "cat" is recognised, then an attempt will be made to acquire a lexical entry for "the". Given the simplifying assumption that the meaning of child-directed utterances may be inferred from context, lexical acquisition may succeed either where the lexical item to be acquired corresponds to the whole of the input utterance or where its meaning can be inferred on the basis of existing lexical items. Thus, acquisition of "the cat" may use an existing lexical entry for "the cat sat on the mat". Similarly, acquisition of "the" may, in turn, use the lexical entry "the cat". In each case, a side-effect of the acquisition of the former is the segmentation of the latter.

The top-down approach to lexical acquisition having been outlined, its implementation in a computational model is described below.

Implementation of Lexical Acquisition

Inputs to acquisition consist of an utterance paired with its semantic representation. Lexical recognition uses existing lexical entries where possible, with lexical acquisition only being triggered given failure of normal lexical recognition. Segmentation of input utterances takes place as a "side-effect" of lexical recognition or acquisition.

Lexical acquisition, in those cases where it is possible, works by using the semantic input to acquisition along with existing lexical knowledge to infer the lexical items which combine to produce the input utterance. Thus, in order to describe lexical acquisition, it is necessary to first outline those processes which it attempts to reverse.

Representation in the model developed uses the Lexical-Functional Grammar (LFG) formalism (Kaplan & Bresnan 1982). In LFG, there are two basic levels of representation, Constituent ("C-") Structure, corresponding to an annotated phrase-structure tree, and Functional ("F-") Structure, consisting of a shallow semantic representation. The F-Structure for an utterance provides a convenient semantic input to acquisition, embodying the simplifying assumption that the meaning of a child-directed utterance can be inferred from context.

The F-Structure for an utterance is built, through the process of *unification,* from the F-Structures of its lexical entries. In LFG, the lexical entry contains, in addition to the word (represented here as a list of phones) and its syntactic category, a number of defining feature:value pairs; for instance:

Cat1 → [dd, sw] ("the") [def:pos]

Cat2 → [k, a, t] ("cat") [pred:cat]

The F-Structure for the phrase "the cat" represents the unification of these features:

[def:pos] U [pred:cat] = [def:pos, pred:cat]

Unification embodies the constraint of uniqueness, which means that any feature may only take a single value:

[num:pl] U [num:sg] = FAILS

In lexical acquisition, the task is, given the utterance and its F-Structure, to infer the word/F-Structure pairs representing its constituents. The nature of unification forms the basis for the inferences made. The inferences licensed by unification mean that the process of acquiring the meanings of words is necessarily incremental. We illustrate this point with the following example.

Existing Lexical Entries

Cat3 → [dd, sw, d, o, g] ("the dog") [def:pos, pred:dog]

Cat2 → [k, a, t] ("cat") [pred:cat]

Input

[dd, sw, k, a, t] ("the cat") [pred:cat, def:pos]

In the case of complete failure to recognise the input utterance, the input utterance/F-Structure pair would be stored as a new lexical entry. However, in this case, "cat" is recognised and so lexical acquisition is triggered for "the". This uses information contained in the input along with the existing lexical entry for "the dog", which subsumes the lexical item to be acquired. Unification licenses the inference that any feature:value pairs not shared by "the cat" and "the dog" are not defining feature:value pairs of the common constituent "the". Shared feature:value pairs, however, may or may not derive from the lexical item in question. The possibilities are illustrated below:

(i) Common feature derives from common constituent "the"

[pred:dog] U [def:pos] = [pred:dog, def:pos]

(ii) Common feature derives from the unshared constituent

[pred:dog, def:pos] U [] = [pred:dog, def:pos]

(iii) Common feature derives from multiple sources

[pred:dog, def:pos] U [def:pos] = [pred:dog, def:pos]

In the model implemented, no lexical acquisition takes place save that which is required for the task of recognising utterances input. This means that only “the” is acquired, with there being no attempt to acquire a lexical entry for “dog” until this also becomes necessary. We can narrow down the possible lexical entries for “the” to the following two:

Cat1 → [dd, sw] (“the”) [def:pos]

Cat1 → [dd, sw] (“the”) []

The first of these, which represents the assumption that common features derive from the common constituent “the”, is that which is acquired by the model. Incremental lexical acquisition is required since the processing of later inputs may reveal this assumption to be incorrect. However, it is precisely because this assumption, if incorrect, may later be recognised as such that we prefer it over the alternative hypothesis which later inputs may not correct.

The Incremental Nature of Lexical Acquisition

Lexical acquisition and segmentation are incremental in the account offered since the lexical items initially acquired are analysed in the process of acquiring later lexical entries. It is also incremental in another respect which is that feature:value pairs found to be non-defining may be removed from existing lexical items at any stage in the process. This aspect of acquisition is necessitated by the nature of the process of attempting to reverse unification which, as outlined above, means that inferences about the meanings of constituents are uncertain. If we initially assume all possible defining feature:value pairs for a lexical item, then those which are incorrect will eventually lead to failure of unification and can thus be removed.

Incremental meaning acquisition is implemented by viewing lexical lookup and acquisition, not as separate processes, but as variations of the same process, differing along the dimension of knowledge acquired/required. When a lexical entry is accessed, its F-Structure is compared with the input F-Structure. If the lexical F-Structure fails to unify with the input F-Structure, then a new lexical entry is created containing just those feature:value pairs common to the two F-Structures. In this way, any feature:value pairs which would lead to failure of unification in the processing of the current input utterance are removed. An example is outlined below.

Existing Lexical Entries

Cat3 → [dd, sw, k, a, t] (“the cat”) [def:pos, pred:cat, num:sg]

Cat4 → [z] (“s”) [num:pl]

Input

[dd, sw, k, a, t, z] (“the cats”) [pred:cat, def:pos, num:pl]

In this case, the F-Structure of the existing lexical entry for “the cat” fails to unify with the input F-Structure since the number values conflict. A new lexical entry for “the cat” is created containing just those features shared by the existing lexical entry and the new input:

Cat3 → [dd, sw, k, a, t] (“the cat”) [def:pos, pred:cat]

Recovery from Mis-segmentation

Recognition of a known lexical item may be constrained by its left and right contexts, where these also consist of known words. Lexical recognition is, however, impaired in the context of acquisition since unknown lexical items cannot constrain the interpretations of adjacent lexical items. Furthermore, since known lexical items may be only a small subset of actual lexical items, there may be a failure to recognise ambiguity (i.e., between a known lexical item and unknown alternatives). An example of a possible mis-segmentation resulting from the context of acquisition would be the segmentation of “chased the dog”, given knowledge of the lexical item “chase” but not “chased”:

[ch, aa, s | t, dd, sw, d, o, g]

[ch, aa, s]

Such mis-segmentations do occur in the model, but are not regarded as problematic. The analysis of “chased the dog” as containing the lexical item “chase” is appropriate. While the target adult segmentation suggests that “ed the dog” is not an appropriate unit, further lexical acquisition will result in its segmentation without the need for separate mechanisms for recovery from mis-segmentation.

An Integrated Model of Acquisition Processes

The unified account of lexical acquisition and segmentation described above has been integrated with an account of syntax acquisition in a computational model of child language development. We describe the overall model here in order to illustrate the wider implications of the top-down approach to lexical acquisition for an account of child language development. We illustrate that the incremental account of lexical acquisition and segmentation implies, in parallel with it, an incremental account of the acquisition of a phrase-structure grammar. Thus, by modelling these processes simultaneously, we arrive at an appropriately gradual view of child language acquisition.

The basis for the model is a deterministic left-corner parser. There are two basic actions in parsing in the model: proposing a constituent, bottom-up, on the basis of its leftmost constituent, or attaching a constituent to the existing C-Structure (i.e., parse tree). When parsing using existing rules fails, grammar acquisition is triggered, in the same way that failure of lexical recognition triggers lexical acquisition. Grammar acquisition involves the same basic actions as parsing, but guided by the semantic input to acquisition in the absence of the appropriate rules. F-Structure, as semantic input, provides information about the semantic relations underlying syntactic relations, but in an unordered, language-neutral form. The additional ordering information required at the level of C-Structure is provided by the input utterance.

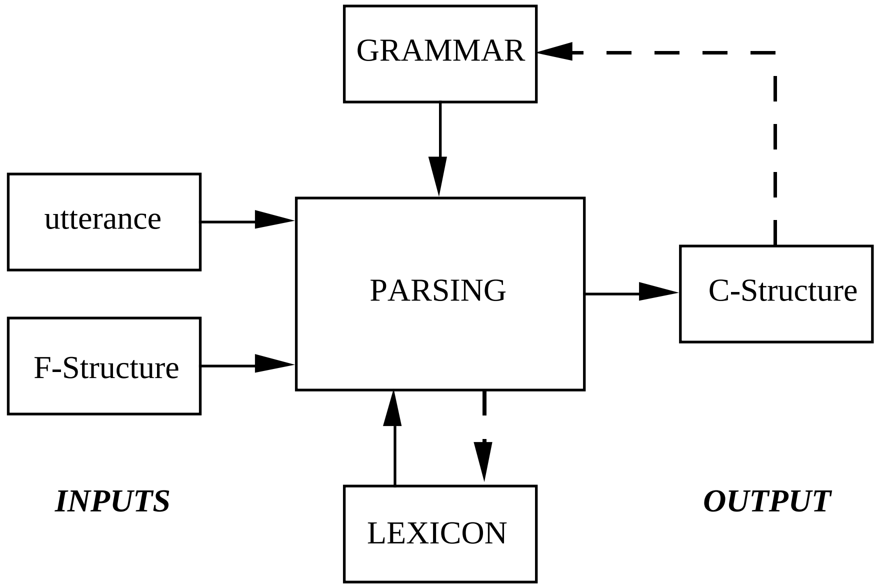

$\rightarrow$ [d, o, g] [pred:dog]

Grammar Rule Acquired

Node2 $\rightarrow$ Node1

$\uparrow = \downarrow$

Input 2

[sw, m, ow, s] ("a mouse") [pred:mouse, def:neg]

Lexical Item Acquired

Node3 $\rightarrow$ [sw, m, ow, s] [pred:mouse, def:neg]

Grammar Rule Acquired

Node4 $\rightarrow$ Node3

$\uparrow = \downarrow$

Input 3

[dd, sw, k, a, t, ch, aa, s, t, f, ii, d, oo] [pred:chase(subj,obj),

("the cat chased fido") tense:past,

subj:[pred:cat, def:pos],

obj:[pred:fido]]

Lexical Item Acquired

Node5 $\rightarrow$ [dd, sw, k, a, t, ch, aa, s, t, f, ii, d, oo] [pred:chase(subj,obj),

tense:past,

subj:[pred:cat, def:pos],

obj:[pred:fido]]

Grammar Rule Acquired

Node6 $\rightarrow$ Node5

$\uparrow = \downarrow$

Inputs 1-3, the first utterances parsed, illustrate straightforward lexical and grammar acquisition given the failure to apply any existing knowledge. Lexical acquisition for such simple cases has already described above. We outline the basics of grammar acquisition here. As the examples above illustrate, no syntactic categorization is assumed by the model so that, for example, each lexical item acquired is initially assigned a unique category label. The induction of syntactic categories is an issue discussed below.

|
Node5
$\uparrow = \downarrow$
|
the cat chased fido

$\uparrow = \downarrow$” annotating Node5 states that its F-Structure is to be unified with that of the node above, Node6. This, in effect, means that all of Node5’s features, deriving from the lexicon, are passed up the tree to Node6. The C-Structure is built in acquisition as follows. Lexical acquisition results in a constituent, Node5, to be processed. The options are to attach it to the existing C-Structure, propose a new constituent above it, or return it as the top-level C-Structure. The first of these options is obviously ruled out in this case. The latter option is also ruled out, since a distinction is built into the model between lexical and non-lexical nodes. Given that a constituent is to be proposed above Node5, the equation annotating Node5 remains to be determined. Since Node5’s features directly unify with that of the input F-Structure (as opposed to with an F-Structure nested within the top-level F-Structure, as in the following examples), it can be assumed that Node6 is the top-level C-Structure node and that Node5 is related to it by the equation “ $\uparrow = \downarrow$”.

Input 4

[sw, m, ow, s, ch, aa, s, t, f, ii, d, oo]	[pred:chase(subj,obj),
(“a mouse chased fido”)	tense:past,
	subj:[pred:mouse, def:neg],
	obj:[pred:fido]]

Lexical Item Acquired

Node9	$\rightarrow$	[ch, aa, s, t, f, ii, d, oo]	[pred:chase(subj,obj),
			tense:past,
			obj:[pred:fido]]

Grammar Rules Acquired

Node8	$\rightarrow$	Node7	Node9
		$(\uparrow \text{subj}) = \downarrow$	$\uparrow = \downarrow$
Node7	$\rightarrow$	Node3	
		$\uparrow = \downarrow$	

The C-Structure which is the output from processing Input 4 is given by Figure 3. In this case, “a mouse” is recognised and Node3 is proposed using the newly acquired lexicon. The existing grammar could at this point be used to propose the constituent Node4 above Node3, but the top-down constraints and lookahead built into parsing mean that, instead, parsing failure is recognised and the switch to acquisition mode triggered. In acquisition, the F-Structure for Node3 is recognised to match the subject function nested in the input F-Structure: a new non-lexical node corresponding to the subject function, Node7, is thus proposed above it and Node3 is annotated with the equation “ $\uparrow = \downarrow$”. Node7 in turn proposes a new node, Node8, corresponding to the top-level F-Structure for the utterance, and is itself annotated with the equation “$(\uparrow \text{subj}) = \downarrow$” (which states that its F-Structure unifies with the subject function of the F-Structure of the node above). Lexical acquisition of “chased fido” is triggered, and this is successful due to the existence of a lexical entry for “the cat chased fido”. The F-Structure of the resulting constituent, Node9, unifies directly with the input F-Structure, as does Node8, and so Node9 is attached to Node8 and annotated with the equation " ↑ = ↓".

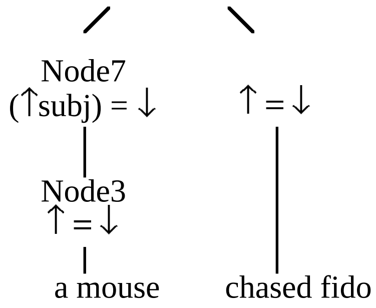

→ [sw] [def: neg]

Grammar Rules Acquired (prior to generalization, outlined below)

Node12	→	Node11	Node9
		(↑subj) = ↓	↑ = ↓
Node11	→	Node10	Node1
		↑ = ↓	↑ = ↓

Figure 4 gives the C-Structure output from processing Input 5. Here, recognition of “dog” triggers acquisition of a lexical entry for “a”. This is slightly more complex than in previous examples, not because a functional morpheme is involved, but because the required defining features are to be found nested within the subject function of the input F-Structure. Acquisition of “a”, which uses the existing lexical entry for “a mouse”, proceeds as follows. Of the top-level input F-Structure and those F-Structures nested within it, “a mouse” *uniquely* shares features (in this case, a single feature) with the subject function, and these features go to make up the lexical entry acquired. Completion of the C-Structure involves the same processes described in the previous examples. The phrase-structure tree given in Figure 4 serves to illustrate how the acquisition of an increasingly complex grammar relates to lexical acquisition and segmentation in the model.

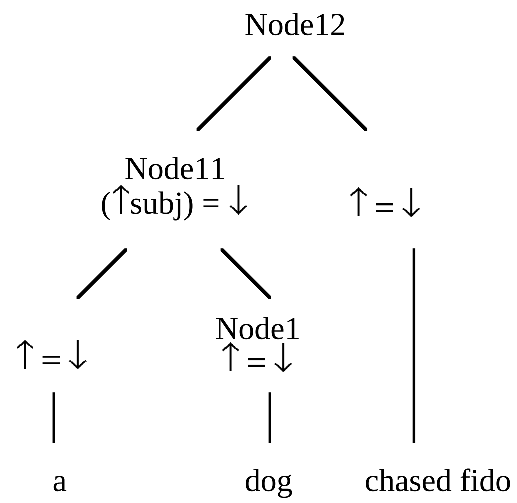

Lexicon

Node1	→	[d, o, g]	[pred:dog]
Node3	→	[sw, m, ow, s]	[pred:mouse, def:neg]
Node5	→	[dd, sw, k, a, t, ch, aa, s, t, f, ii, d, oo]	[pred:chase(subj,obj), tense:past, subj:[pred:cat, def:pos], obj:[pred:fido]]
Node9	→	[ch, aa, s, t, f, ii, d, oo]	[pred:chase(subj,obj), tense:past, obj:[pred:fido]]
Node10	→	[sw]	[def: neg]

Grammar

Node2	→	Node1 $\uparrow = \downarrow$	
Node6	→	Node5 $\uparrow = \downarrow$	
Node12	→	Node13 $(\uparrow subj) = \downarrow$	Node9 $\uparrow = \downarrow$
Node13	→	Node3 $\uparrow = \downarrow$	
Node13	→	Node10 $\uparrow = \downarrow$	Node1 $\uparrow = \downarrow$

The Integrated Model as an Account of Child Language Development

We have suggested that incremental learning is a desirable feature in a computational model of child language acquisition, given the gradual nature of children's acquisition. Both lexicon and grammar are acquired incrementally in the model described, so that, at a general, qualitative level, it can be viewed as embodying a reasonable theory of child language development. The model can also be evaluated with respect to relatively specific language acquisition phenomena, since the integration of learning processes in the model suggests explanations for some of these. Here, we briefly consider the model's account of functional morpheme acquisition, focusing upon the explanation offered of the onset of phenomena like functional morpheme omission and overgeneralization.

Existing accounts of functional morpheme omission have failed to account for the observation that children may include functional morphemes in their utterances before they omit them. Gerken (1987) cites this observation in her criticism of the "telegraphic perception hypothesis", the idea that children fail to produce functional morphemes because they fail to perceive them. The idea of a "production bottleneck" as responsible for functional morpheme omission (e.g., Pinker 1984) has avoided some of the criticisms of the telegraphic perception hypothesis, but, crucially, it also fails to explain why functional morphemes should be produced before the onset of their omission.

Overgeneralization is, similarly, a phenomenon which may follow a period of correct use of regular and irregular forms. Its onset has been linked with the acquisition of a generalized (as contrasted with word-specific) rule for the formation of the regular past tense, in the case of overregularization of the past tense "ed" (Marcus 1992). However, accounts which seek to explain overgeneralization in terms of the acquisition of this rule fail, crucially, to explain why the latter is acquired when it is (Plunkett 1994).

Since the top-down account of lexical acquisition incorporates no innate distinction between content and function words, functional morpheme acquisition is identical to other cases of lexical acquisition in the model. The integrated model of acquisition processes incorporates an account of the acquisition of rules like that for the formation of the regular past tense. While there is no explicit model of language production, the lexicon and grammar acquired provide constraints on the utterances that could be produced at any stage in acquisition.

The model's account of the onset of functional morpheme omission is a competence-based account which links it with the development of functional morpheme recognition, or acquisition. In the model, those units which can be either produced or omitted correspond to lexical entries acquired. The nature of these units changes during the course of acquisition. It can be predicted that, so long as "Daddy goes" exists as an unanalysed lexical unit, the verb inflection will not be omitted from it. Functional morpheme omission only becomes a possibility once "goes" has been segmented into the constituent verb stem and inflection. Thus the model is consistent with the observation that functional morphemes are produced during the early stages of acquisition. The model's explanation of functional morpheme omission will not be discussed in detail here. Briefly, it is linked with the model's predictions concerning order of acquisition of functional morphemes, which consider at what point sufficient knowledge is acquired to enable selection of the appropriate lexical item in language production (Gaylard 1995).

The model provides a simple account of how the acquisition of the past tense rule and the possibility of overregularization arise following a period of correct use of regular and irregular past tenses. In the early stages of acquisition, inflected verbs like "walked" and "went" exist as unanalysed units. There is no rule for the formation of the past tense and thus no possibility of overregularization. Lexical acquisition and segmentation in the model eventually result in the acquisition of lexical items corresponding to verb stems such as "walk" and "go" and inflections like "ing" and "ed". Syntax acquisition in turn results in rules for the formation of the progressive and the past tense. Since the rule for the formation of the progressive applies to both "walk" and "go", these come to be considered instances of the same syntactic category. The consequence is that the rule for the acquisition of the past tense, acquired in relation to verb stems like "walk", applies equally to a verb like "go" for which an irregular past tense form exists, and thus the possibility of overregularization arises.

The assumption that children begin learning language by producing utterances by rote is central to the top-down account of lexical acquisition. Explanations of the onset of functional morpheme omission and overgeneralization arise naturally from a model of acquisition in which a progression is predicted from a relatively unproductive repertoire of utterances, which does not allow for the possibility of errors, to a more productive system.

Conclusion

We have described a top-down, incremental approach to lexical acquisition. This serves to address the criticism of top-down accounts of lexical recognition that they are incompatible with an account of acquisition. At the same time, the top-down approach to lexical acquisition avoids a major deficiency highlighted in relation to bottom-up accounts of lexical acquisition. While these focus upon the problem of segmentation in acquisition, paying insufficient attention to the development of the lexicon, the top-down account is a unified account of the development of the lexicon and segmentation abilities. Such a unified account is, in fact, intrinsic to the top-down approach, in which an account of the acquisition of the lexicon automatically incorporates an account of segmentation during acquisition. A related issue is that, while the bottom-up approach requires the assumption of an innate propensity to acquire initially meaningless units through segmentation of the input, all acquisition processes in the top-down account are driven directly by the requirements of comprehension.

The incorporation of the top-down approach to lexical acquisition in a computational model of child language acquisition serves to illustrate how it may contribute to an account of the observed features of child language development. Incremental learning means that development in the model is gradual, predicting a progression from the rote reproduction of utterances to the acquisition of increasingly complex grammatical constructions. Syntactic categories are also gradually acquired in the model so that the “lexical” nature of early child language (Ninio 1988) is captured. Furthermore, the simultaneous acquisition of segmentation abilities with the grammar suggests a simple explanation of the onset of phenomena like functional morpheme omission and overgeneralization following periods of correct use of functional morphemes.

References

Altmann, G. & Steedman, M. (1988). Interaction with context during human sentence processing. *Cognition, 30(3)*, 191-238.

Bard, E.G. (1990). Competition, Lateral Inhibition, and Frequency: Comments on the Chapters of Frauenfelder and Peeters, Marslen-Wilson, and Others. In G.T.M. Altmann (Ed.), *Cognitive Models of Speech Processing: Psycholinguistic and Computational Perspectives,* pp.185-210. London: MIT Press.

Briscoe, E.J. (1989). Lexical Access in Connected Speech Recognition. In *Proceedings of the 27th ACL Congress,* Vancouver, June 1989, pp.84-90. [s.l.]: ACL.

Cairns, P. *et al* (1994). Bootstrapping word boundaries: a bottom-up corpus-based approach to speech segmentation. Submitted to *Cognitive Psychology.*

Cartwright, T. & Brent, M.R. (1994). Segmenting Speech Without a Lexicon: The Role of Phonotactics and Speech Source. In *Proceedings of the 1st Meeting of the Association for Computational Phonology (Post-Conference SIGPHON Workshop of the 32nd Annual Meeting of the ACL),* New Mexico State University, July 1994, http:/xxx.lanl.gov/abs/cmp-lg/94 12005.

Cole, R.A. & Jakimik, J. (1980). A Model of Speech Perception. In R.A. Cole (Ed.), *Perception and Production of Fluent Speech,* pp.133-63. Hillsdale, NJ: Erlbaum.

Crain, S. & Steedman, M. (1985). On not being led up the garden path. In D.R. Dowty *et al* (Eds.), *Natural Language Parsing*, pp.320-58. Cambridge: Cambridge University Press.

Cutler, A. & Butterfield, S. (1992). Rhythmic Cues to Speech Segmentation. *Journal of Memory and Language, 31(2),* 218-36.

Cutler, A. & Norris, D. (1988). The Role of Strong Syllables in Segmentation for Lexical Access. *Journal of Experimental Psychology: Human Perception and Performance, 14(1),* 113-21.

Cutler, A. & Mehler, J. (1993). The Periodicity Bias. *Journal of Phonetics, 21(1/2)*, 103-8.

Friederici, A. D. & Wessels, J.M.I. (1993). Phonotactic Knowledge of Word Boundaries and its Use in Infant Speech Perception. *Perception & Psychophysics, 54(3),* 287-95.

Gaylard, H.L. (1995). *Phrase Structure in a Computational Model of Child Language Acquisition.* Submitted for the degree of Doctor of Philosophy, University of Birmingham, March 1995.

Gerken, L. (1987). Telegraphic Speaking Does Not Imply Telegraphic Listening. *Papers and Reports on Child Language Development, 26,* 48-55.

Grosjean, F. & Gee, J.P. (1987). Prosodic Structure and Spoken Word Recognition. *Cognition, 25(1/2),* 135-55.

Hirsh-Pasek, K. *et al* (1987). Clauses are Perceptual Units for Young Infants. *Cognition, 26(3),* 269-86.

Jusczyk, P.W. (1993). From general to language-specific capacities: the WRAPSA model of how speech perception develops. *Journal of Phonetics, 21(1/2),* 3-28.

Jusczyk, P.W. *et al* (1992). Perception of Acoustic Correlates of Major Phrasal Units by Young Infants. *Cognitive Psychology, 24(2),* 252-93.

Kaplan, R.M. & Bresnan, J. (1982). Lexical-Functional Grammar: A Formal System for Grammatical Representation. In J. Bresnan (Ed.), *The Mental Representation of Grammatical Relations*, pp.173-281. London: MIT Press.

McClelland, J.L. & Elman, J.L. (1986). The TRACE Model of Speech Perception. *Cognitive Psychology, 18(1),* 1-86.

Marcus, G.F. *et al* (1992). Overregularization in Language Acquisition. *Monographs of the Society for Research in Child Development, 57(4),* 1-165.

Mehler, J. *et al* (1990). Constraining Models of Lexical Access: The Onset of Word Recognition. In G.T.M. Altmann (Ed.), *Cognitive Models of Speech Processing: Psycholinguistic and Computational Perspectives,* pp.236-62. London: MIT Press.

Nelson, D.G.K. *et al* (1989). How the prosodic cues in motherese might assist language learning. *Journal of Child Language, 16(1),* 55-68.

Newport, E.L. (1977). Motherese: The Speech of Mothers to Young Children. In N.J. Castellan *et al* (Eds.), *Cognitive Theory*, Vol. 2, pp.177-217. Hillsdale, NJ: Erlbaum.

Ninio, A. (1988). On Formal Grammatical Categories in Early Child Language. In Y. Levy, M. Izchak & M.D.S. Braine (Eds.), *Categories and Processes in Language Acquisition*, pp.99-119, London: Erlbaum.

Peters, A.M. (1983). *The Units of Language Acquisition.* Cambridge: Cambridge University Press.

Pinker, S. (1984). *Language Learnability and Language Development*. London: Harvard University Press.

Plunkett, K. (1994). Learning the Arabic Broken Plural: The Case for Minority Default Mappings in Connectionist Nets. In P. Broeder & J. Murre (Eds.), *Workshop on Cognitive Models of Language Acquisition*, Tilburg University, April 1994, pp.47-51. Tilburg: Katholieke Universiteit Brabant.

Shillcock, R. (1990). Lexical Hypotheses in Continuous Speech. In G.T.M. Altmann (Ed.), *Cognitive Models of Speech Processing: Psycholinguistic and Computational Perspectives,* pp.24-49. London: MIT Press.

Tyler, L.K & Marslen-Wilson, W.D. (1982). Speech Comprehension Processes. In J. Mehler *et al* (Eds.), *Perspectives on Mental Representation: Experimental and Theoretical Studies of Cognitive Processes and Capacities,* pp.169-84. London: Erlbaum.

Wolff, J.G. (1988). Learning Syntax and Meanings Through Optimization and Distributional Analysis. In Y. Levy, M. Izchak & M.D.S. Braine (Eds.), *Categories and Processes in Language Acquisition*, pp.179-215. London: Erlbaum.