1 *Missing Objects in Child Language*

1.1 General Goals

Three core premises guide this book. First is the idea, universally accepted in contemporary formal grammar, that the missing components of a sentence must be recovered. Second, that lexical knowledge, syntactic structure and contextual understanding all interact in this recoverability. The third premise is that the growth of children's grammars can be best characterized in terms of the developmental interplay of these three domains, or "submodules," of grammar. In this book, we aim to show that this interplay is especially apparent in the domain of object omission. This chapter presents the problem of missing objects in child language, explains the nature of the learning problem these missing objects represent and articulates the learnability approach, which serves as the basis for our methodology.

The following scenario happened long ago. A small boy, just shy of eighteen months, waddled into the kitchen and said to his mother, "*Want*." The mother turned away from the dishes to ask him, "*What*?" The child repeated, this time louder, "*Want!!*" This happened twice until the little boy started to cry. This conversational exchange between parent and child obviously failed. It did not fail because the sentence was spoken in bad English but, rather, because the child's conversational partner was unable to identify the referent of some missing elements in it.

When the child said, "*want*," both subject and object were left out (__ *want* __). This is an ordinary occurrence, at least if you happen to be around people under the age of two. However, for syntacticians and psycholinguists who study early language development, this is an interesting phenomenon. About the missing subject, the syntacticians would say that it was likely recoverable as referring to the speaker. In the case of the direct object, an obligatory companion of the

transitive verb *want*, the content was neither easily recoverable from grammar nor from context. Specialists in language development would note that this initial stage in language acquisition is characterized by pervasive optionality of arguments (i.e., objects and subjects) and functional categories. These investigators are aware that early telegraphic speech evolves quickly into fairly well-specified sentences. The question is: How does this happen? We feel that to answer this question we must start by addressing a more basic issue: What is the difference between the grammar of this boy and that of his mother?

Returning to the above scenario, we will start by giving you a hint of how things like this happen. A speaker can delimit, from the meaning of the verb *want*, the possible interpretations of the object to the set of desirable things. The context of the utterance (the kitchen, with Mom as the interlocutor) helps to further restrict the set of desirable things to only those relevant in the specific circumstances. The child is likely to want something from the set of things available in this kitchen, at this moment, and hopefully, from the kinds of things Mom is believed to control. The problem remains, even with contextual information, that the kitchen contains many objects that this two-year-old might find interesting, even leaving aside the extra-interesting stuff that Mom will never agree to. The information in the baby's short sentence is just not sufficient and the result is a failure to communicate.

Actually, the whole field of language acquisition, or developmental linguistics, as called by some, is captivated by small sentences like these. Developmental linguistics is concerned with how children start their life with a language that comes packaged in tiny sentences that are difficult to anchor to a context, and learn to build longer sentences that include all the parts necessary to negotiate meaning and achieve grammaticality. Missing objects and subjects, such as described above, are reported again and again in the speech of young children. During most of the field's history, investigators focused on missing subjects and paid little attention to missing objects. For us, and many others, missing objects are at least as interesting as null subjects, and possibly much more mysterious. Accounts of why children omit objects are even more complicated and contentious than those proposed for missing subjects. In order to tell the story of missing objects in child grammar, we need to delve right down to the very nature of the fundamental elements that form the mental grammar and the mental lexicon, and to unravel how sentence grammar interacts with the extra-grammatical context. This trip will not resemble a tidy tour organized by a travel agency, with every stop fully planned and carefully scheduled. Instead, we hope for a rambling trek taking detours along obscure

roads, occasionally backtracking to earlier points on the road and including as many stops in interesting places as possible.

We start the story of missing objects in child language with a startling admission: contrary to what has been assumed by many linguists for years, at least about languages such as English, adults do it too (Cummins and Roberge 2005). Consider the examples in (1), where the "__" indicates the missing object.

(1) a. Here, read __.
b. I'm hungry, I want to eat __.
c. Let's do dishes. I'll wash__, you dry__.

A keen observer will note that some of the verbs above (*read* and *eat,* at the very least) are normally characterized as "optionally transitive" verbs. However, even verbs that are often described as "obligatorily" transitive can appear with a missing object (2).

(2) a. Did you lock __? (Wife to husband, in the car, about to leave for work)
b. There are those who annihilate __ with violence, who devour __.
(British National Corpus; Cummins and Roberge 2004: 124)

Unlike in the child example presented earlier, English-speaking adults omit objects in ways that: a) allow interlocutors to negotiate the meaning of the missing object; and b) do not trigger intuitions of ill-formedness in speakers. Thus, such cases go largely unnoticed. When we start to pay attention to them, these implicit objects can provide crucial insights into how meaning is constructed and how speakers make sense of missing object arguments through a variety of mechanisms, including through deixis ("this thing," in 1a) and generics, from the prototypical meaning of the verb ("food," in 1b), or the discourse context ("dishes," in 1c). This book is an exploration of how children sort out the fine details of when object arguments can be omitted and when they cannot.

Explaining language acquisition becomes an even more challenging enterprise when we move beyond a single language and consider the range of possibilities observed in human languages. It turns out that across different languages there is substantial variation in how missing objects are licensed and how they are interpreted. Analyzing cross-linguistic variation represents the first step in mapping out the complexity of the learning task that children face. Here we are obliged to state not one but two obvious things. Firstly, since a child must come prepared to acquire any language, an explanation that can fit only some languages will never be fully satisfying. Secondly, the

child does not know at the outset how the language she has to learn works. Therefore, our account must map out the relationship between the following three dimensions:

- What underlies children's ability to learn language (i.e., the capacity);
- What speakers specifically need to learn about implicit objects and transitivity in the target language (i.e., the representation); and
- How the language experience is configured in a specific language (i.e., the input).

In other words, triangulating from these dimensions of input, representation and capacity, language acquisition theories must be able to account for the starting point, the process of how experience leads to change in representations and the eventual outcome of language development.[1]

Our point of departure is then a learnability-based approach. Learnability approaches consider the structure of the domain to be learned in order to deduce what type of steps must be involved in the task of learning the relevant properties. It is a bit like tasting a dish and trying to imagine from its flavors what might have happened in the kitchen. In the sciences, this is known as "reverse engineering." We assume that the meaning of sentences in general (and consequently, of sentences with missing objects) is compositional in nature. A compositional view of meaning states that the meaning of a sentence results from the combination of the meanings of the parts (the lexical content, provided by words and morphemes) and the way these are combined (the syntax). Furthermore, we hold that syntactic structure defines how meaning can be tuned to context, so that syntax not only organizes words but also constrains how access to context happens. To a certain extent, sentence meaning and context interact as part of the general human ability to make inferences about context and communication (pragmatics). Other aspects of how syntax and meaning interact are specific to each language, and thus, part of the learning task. The goal of this book is to understand how children learn the specific aspects of syntax and of the syntax-discourse interface.

We take a particular view of language development where ontogenetic stages resemble the shape of potential mature states. This is known as the continuity assumption, or the strong continuity hypothesis (Pinker 1984). A bit like the little tail that appears in immature stages in tadpoles and

[1] On how stimuli and input interact with the linguistic system in the language acquisition process, see Carroll (2001).

human embryos, to subsequently disappear during development, but remains to the mature stages of dogs and lizards. Some readers might recognize in this a version of Ernst Haeckel's recapitulation theory ("ontogeny recapitulates phylogeny"; Haeckel 1866). In its raw form, the continuity hypothesis in language acquisition suffers from the same limitations as Haeckel's recapitulationist model: ontogeny does not replicate actual adult stages, but embryonic stages from other points in the evolutionary pathways. Leaving the biological dimension aside, in linguistics the continuity assumption has proven a useful tool to analyze target-deviant behavior in children. It happens with clear regularity that nonadult structures in the production of children learning a given language share features with adult structures in a different language.

Language development never starts with nothing, but rather starts with a variant of something. We posit that this something is a default, a minimal structure used with the widest range of meanings (Lebeaux 1988, 2000; Roeper 2007, 2008). We explore the hypothesis that "learning" consists of selecting one member of a set of predetermined hypotheses. In Fodor's (1998) language of unambiguous triggers, these predetermined hypotheses are simply pieces of syntactic tree structure. We propose that children start with the simplest representation, which is structurally the most basic and, therefore, the most flexible possible structure. With the input provided by experience, functional structure develops, giving specificity to the original default representation. From the various possible mechanisms of recoverability available for missing elements some are reinforced while others undergo attrition or blocking. The semantic representation of sentences becomes narrower, its relation to context more constrained, more precisely defined and more language-specific.

To provide a precise characterization of how this happens in the case of null objects, we propose the following steps. The first step is to consider the general analyses that have been proposed to account for missing objects in linguistics. The second is to characterize object omission phenomena cross-linguistically so as to provide a complete picture of the learning task involved. Third, we propose to examine in more depth, and with more semantic details, what children do in a few languages where null objects in acquisition have been examined carefully. As a preamble to taking this route, in this chapter we introduce some fundamental grammatical notions, a bit of the history of the field of language acquisition and make explicit our basic assumptions about how all of this should be approached.

1.2 What Are Verbs and How Are They Learned?

1.2.1 Verbs and Nouns

> Missus Gloria Pearson went to Livingstone College in North Carolina. She told me. I guess that's why she know so much about how to fix my English speaking [. . .] Missus Pearson say she will teach us all the rules. She say English is governed by rules of grammar, and the rules, she say, go special with nouns and verbs.
>
> (Verdelle 1995: 186)

Few linguists these days would agree that the purpose of grammar is to fix anybody's way of speaking. Most, however, will concede that languages are governed by rules and that grammar is but a description of speakers' rule-governed behavior. There is also wide consensus that within the rules of grammar nouns and verbs indeed occupy a special place. Nouns generally refer to entities, which can be abstract or concrete, and verbs are supposed to refer to actions, occurrences or states. Noun phrases (NPs) denote individuals or sets of individuals, whereas verb phrases (VPs) denote situations. The VP, in conjunction with the functional structure it is anchored in, establishes the association between the verb and the various participants in a situation, which are known as arguments of the verb. Such meaning-based definitions are intrinsically limited: some nouns denote actions (such as the noun *action* itself), some verbs don't. English *be*, for instance, is considered a verb but it has nothing to do with actions. Indeed, it denotes very little if anything at all: the copula verb *be* can be seen as just a linker between other contentful elements, or a placeholder for grammatical tense. In various languages, like Arabic or Irish or even some varieties of English (Becker 2004), the copula need not be present at all.

If a verb does not have to refer to actions to be a verb, what does it have to do? Since notional definitions of verbs and nouns are weak, the standard approach is to define verbs and nouns in terms of the syntactic or morphological behavior of the lexical class. Verbs are relational categories that select arguments of a given type. The main classification criterion of a lexical verb is according to how many argument noun phrases (NPs) it typically associates with. This is called the argument structure of a verb.

(3)
a. Ditransitive: [3 arguments: Subject, Direct Object, Indirect Object]
   The witch gave an apple to the girl.
b. Transitive: [2 arguments: Subject, Object]
   Snow White ate the apple.
c. Intransitive: [1 argument: Subject]
   She swallowed, then fainted.

Intransitives are further classified into two types, according to the role of the subject. For *swallow*, the subject is the agent. For *faint,* it is not. Swallowing is something Snow White did; fainting, instead, happened to her. The subjects of the first class of verbs (known as unergative verbs) share many syntactic properties with the subjects of regular transitive verbs. Across languages, subjects of the second type of verbs (known as unaccusative verbs) share a range of syntactic properties with direct objects.

Verbs thus play a crucial role in determining not just the number of arguments in a sentence but how those arguments are situated within sentence structure. Interestingly, children seem to know how verbs organize sentences from the beginning of their syntactic development, at the very onset of the ability for combinatorial speech. Researchers who work on this early stage of language development debate whether children at this age have mental representations for grammatical categories such as nouns and verbs. It is worthwhile to take a moment to consider the evidence pertinent to the existence (or lack thereof) of abstract categories at the earlier stages of syntactic development. Starting from the no-categories camp, researchers such as Olguin and Tomasello (1993) taught novel verbs and nouns to children aged two. This classic experimental approach allows researchers to infer what children know under carefully controlled circumstances. The method dates back to Berko Gleason (1958), who presented invented words to children in order to verify whether they understood the rules of English morphology.

(4) Here is a wug.
Now, there are two . . .

Since in such cases, the given words are completely new, there is no question of prior experience. For Olguin and Tomasello, the question was whether children exhibited comparable syntactic flexibility for the two core grammatical categories. These authors observed that what children learned was different for nouns and verbs. Two-year-olds treated the newly learned nouns creatively, using them in varying syntactic positions and playing different semantic roles. In contrast, when taught novel verbs, the children most often reproduced the same combinatorial pattern they had heard for each specific verb, roughly 90 percent of the time. When they did combine a new verb with a known noun, they often failed to follow the canonical word order. The implication, according to these authors, is that young children are productive with their early language in some ways, but not in others. The overall conclusion extracted in this and related work is that children are not primarily creating an abstract

category for verbs for purposes of syntax, but are instead extracting lexically specific schemas (Tomasello, Akhtar, Dodson and Rekau 1997; Tomasello 2000).

There are some (rather obvious) arguments in favor of the opposite conclusion. Since verbs are distributionally more restricted than nouns, the very same data can be interpreted to say that even two-year-olds are aware of this fundamental asymmetry between nouns and verbs. It is not that babies are more cautious with verbs because they lack an abstract representation of verbs, but because they are already sensitive to the nature of verbs. It is almost as if they know that so much more is at stake when using a new verb. As Jean Aitchison tells us, "Verbs are inextricably linked with syntactic structure" (Aitchison 1994: 111). Nouns go where the grammatical context sends them. Verbs are the context. The same event can be described differently: *the hunter chases the rabbit*, or *the rabbit flees from the hunter*. Bad luck for the rabbit, in either case. For the speaker, however, the choice of verb frames the shape of the sentence. Beyond determining the number of arguments and their roles in the sentences, the verb also determines their semantic types. Snow White may be able to eat the apple, but the apple will not eat Snow White. Metaphorical extensions aside, only animate entities can perform the action of eating, and only certain things are eaten. In other words, verbs determine or "select" the properties of their subjects and objects. As a consequence, the acquisitions of verbs and their syntactic distributions are narrowly interlocked (Gleitman, Nappa, Cassidy, Papafragou and Trueswell 2005).

1.2.2 Verbs and Objects

> There is a slim but tantalizing link between a verb's [argument] structure and its meaning.
>
> (Aitchison 1994: 11)

We are now closer to the matter that interests us, which is verbs and their relationship to their objects. Verbs may be at the center of sentence structure, but one element is closer to the verb than any other: the object. The object is an inherent part of the event described by the verb and the verb restricts what can appear in the object position. As we will see, this happens even when the objects are left unpronounced. The distributional relation between verbs and objects is known as transitivity.

In Chapter 2, we will discuss a particular formal approach to transitivity, but for now, we start by introducing some basic facts. The first is that there is a range of relations between verbs and objects. On the extreme of transitivity,

we have verbs whose meanings are fundamentally grammatical and for which the object is essential for the event to be fully characterized:

(5) a. Do a dance/math/household chores/some gardening/dinner.
b. Finish a cake/a book/my chores.
c. Get ice cream/a headache/a car/a boyfriend/some money.

These verbs denote highly abstract properties of event structure, but express little lexical or encyclopedic meaning. Without the objects, we don't really have a clue what action or situation is actually involved. Other verbs may be lexically contentful, but still considered obligatory transitives. We still don't have a clear sense of what determines the difference between these verbs and verbs whose direct objects are variably expressed. A well-known pair is *eat* vs. *devour*. These denote fundamentally the same event. The first is neutral while the latter says something additional about manner. English speakers have fairly robust intuitions that there is a contrast between the two. Speakers know that you are as likely as not to pronounce the object of *eat.* Both options feel equally grammatical. In contrast, eliminating the object of *devour* leads to intuitions of incompleteness, or ungrammaticality (represented conventionally by an asterisk). Consider the examples in (6).

(6) a. The dwarves were eating their supper when Snow White arrived.
b. The dwarves were eating Ø when Snow White arrived.
c. The dwarves were devouring their supper when Snow White arrived.
d. *The dwarves were devouring Ø when Snow White arrived.

Yet, even the objects of *devour*-type verbs can go missing, but only given the right sort of context, as we saw in (2b). We use the Ø symbol throughout this book to represent the position where the object should have appeared. This symbol thus indicates a syntactic position that is not overtly occupied and for which it is not useful or necessary to provide a specific type of empty element for the purpose of our presentation. In the next few paragraphs, we hope to explain why it is a good idea to make reference to an empty position and, later on, we provide more specific characterizations of potential null elements for this position.

For the optionally transitive verbs such as *eat*, the encyclopedic meaning of the verb does not generally change depending on whether or not the object is included: the emphasis is on the activity when the object is missing and the sentence is more informative when the object is present (*eating some stuff* vs. *eating a cookie*). However, the presence or absence of the object does

help define the type of event. In verbs with incremental objects (consumption and creation verbs fall into this category), the object defines the natural end point for activity. The dwarves' eating ends when the supper is all gone. Building a house ends when the house is finished. Drawing a circle ends when the circle is complete. Erasing a circle works the same, but in the opposite direction. When the object is implicit we say that the verb is interpreted as describing an unbounded activity. The object, when present, provides the boundary or culmination for the event. This transitivity alternation represents two different aspectual classes of dynamic events: activities (7a) vs. accomplishments (7b).

(7) a. To write *in one summer/for months. (van Hout 2008)
b. To write a script in one summer/*for months.

At one extreme of the transitivity continuum we have verbs that never occur without their objects, because the characterization of the situation heavily depends on it. This is the case of light verbs such as *have, get*, which have little meaning without their complements (cf., *get a book/get some money/get the flu*). At the opposite extreme we find the so-called unergatives, or agent-intransitive verbs: *dance, run, sing*, whose object is so predictable that it is seldom realized. The standard, lexical view of transitivity holds that these verbs have a single argument, the subject. We adopt an alternative view according to which unergatives are held to have implicit objects. According to Hale and Keyser (2002), these can be analyzed at the morphological level as having a hyponymic noun in the position of the object noun. The object is not a fully referential nominal phrase. Instead, the object is reduced to the lexical content of the root; basically, the nonreferential noun that corresponds to the action denoted by the verb. In other words, *to dance* means to *dance a dance*. This is not as extreme as it may seem at first. In fact, this is not that different from the earlier step we took to analyze *eat*-type verbs. Furthermore, there is evidence for this analysis. Unergative verbs commonly allow a modified cognate object, a regular DP or measure phrase to appear as their complement (Massam 1990):

(8) a. Sleep a restless sleep.
b. Slept the whole night.
c. Live a good life.
d. John ran a good race.
e. Then he swam a mile.
f. She danced a tango.