3.3 The Acquisition Problem: The Poverty of the Stimulus

The child acquires a finite, generative system – a grammar – which generates structures which correspond more or less to utterances of various kinds. Children acquire these grammars despite a poverty of stimulus on three levels.

First, the child hears speech from adults, peers, and older children. This stimulus does not consist uniformly of complete, well-formed utterances; it also includes sentence fragments, slips of the tongue, incomplete thoughts, ill-formed utterances from people who do not know the child's language well, and even sentences artificially simplified, supposedly for the benefit of children. Even if only 5 percent of the expressions the child hears are of this latter type, there will be a significant problem in generalizing to the set of grammatical sentences of the language, because the pseudo-sentences do not come labeled as defective.

Second, a child encounters only a finite range of expressions, but she comes to be able to produce and understand an infinite range of novel sentences, going far beyond the sentences heard in childhood. We know this at the intuitive level when we recognize that we constantly encounter novel sentences that we have not heard or used before. Consider the sentences on this page; it is unlikely that the reader has encountered any of them before in the precise form in which they occur here.

More formally, to understand that there is an infinite number of English sentences, one has only to realize that, in principle, any given sentence may be of indefinite length. Three iterative devices permit this and they may occur in various combinations:

Relativization: This is the dog that chased the cat that killed the rat that caught the mouse that nibbled the cheese that . . .
Complementation: I think that Jim asked me to tell Kim that Tim thought that I said that . . .
Co-ordination: Jay went to the movies and to the concert, and Ray and Kay went out for dinner, and Fay stayed at home, and . . .

If a sentence may be of indefinite length, then it follows that people have a capacity to use and express an indefinite number of sentences. Since a person's experience is finite and the mature capacity ranges over infinity, the stimulus alone cannot fully determine the mature capacity.

Third, people come to know things subconsciously about their language, things for which no direct evidence is available in the data to which they are exposed as children learning to speak. People eventually understand

and utter complex, ambiguous sentences, identify paraphrases, and distinguish sentences that may occur in their language from ones that may not. We have just spent a few pages illustrating this kind of thing, and the subconscious, mature capacity that ordinary people have involves language properties made explicit by linguists. However, the crucial properties lie outside the primary linguistic data available to young children. Children are not systematically informed that hypothetical sentences do not in fact occur (e.g. **Who do you wanna go?*, **Kim's happier than Tim's*), that a given sentence is ambiguous, or that certain sets of sentences are paraphrases of each other. Also, many legitimate, acceptable sentence types may never occur in a child's linguistic experience. Such data are not available to preschool children and are not part of their verbal experience. The distinction between the range of data known to the linguist and the much more limited data available to the beginning speaker is of vital importance for the biological view of language development.

This third deficiency is quite crucial. The first two, the imperfection and finiteness of the stimuli, are not decisive kinds of data deficiencies. They do not deny that relevant experience for language learning is available; they simply assert that the experience is "degenerate," hard to sort out. The fundamental deficiency is the third, which says not that relevant experience is degenerate but that in certain areas it does not exist at all. This deficiency shapes hypotheses about the linguistic genotype.

Some structural principle prevents forms like **Who do you wanna go?* from occurring in the speech of English speakers, as we have seen. Children are not exposed to pseudo-sentences like this and informed systematically that they are not said. Speakers come to know subconsciously that they cannot be said, and this knowledge emerges somehow, even though it is not part of the input to the child's development. Furthermore, it is hard to imagine how the inventory of sentences and sentence fragments that constitute the child's linguistic environment could provide even indirect evidence that such sentences do not occur. It is not enough to say that people do not utter such forms because they never hear them. This argument is insufficient, because people say many things that they have not heard, as we have noted. Language is not learned simply by imitating or repeating what has been heard.

This third deficiency of the stimulus is of particular importance in defining our approach to language acquisition. A good deal of evidence exists that the contrast between the child's experience and the range of data available to the linguist is quite substantial. Over the last 40 years, much of the linguistic literature has focused on areas where the best description cannot be derived directly from the data to which the child has access, or is underdetermined by those data, as in the examples with the clitics *to* and *'s* and the pronouns. If the child's linguistic experience does not provide the basis for establishing a particular aspect of linguistic knowledge, another source must exist for that knowledge. That aspect must be known a priori, in advance, in the sense that it is available independently of linguistic experience. We tentatively assume that it is available through genetic prescriptions, while not ignoring other possibilities – for example, that it arises as a consequence of other, nonlinguistic experience.

All this is not to say that imitation plays no role: just that it does not provide a sufficient explanation, given the third data deficiency. This is worth emphasizing, because antagonists sometimes caricature this approach to language acquisition as "denying the existence of learning," when in fact its adherents merely deny that learning is the whole story – a very different matter. The quotation is from a remarkable article in *Science* magazine, in which the authors assert that "Noam Chomsky, the founder of generative linguistics, has argued for 40 years that language is unlearnable," and that they, on the other hand, have "rediscovered" learning (Bates and Elman 1996)!

Caricatures of this type show up in the writing of people who claim that *all* information is derived from the environment, and that there is no domain-specific genetic component to language acquisition. These people deny the poverty-of-stimulus problems, claiming that children derive all relevant information from their linguistic environment. Bates and Elman provide a recent, particularly clear and striking instance of this line, claiming that artificial neural networks can learn linguistic regularities from imperfect but "huge computerized corpora of written and spoken language."[5]

Others have appealed to a structured input which allows children to circumvent the poverty-of-stimulus problems (Snow 1977). Parents and other people often adopt a simple, sometimes artificial style of speech when addressing children; but it is scarcely plausible that this "Motherese" provides sufficient structure for language acquisition to take place on a purely inductive basis. Children do not simply generalize patterns without the aid of genetically determined principles.

There are at least four reasons why this kind of pattern generalization is not the answer to how children acquire speech. First, although children no doubt register only part of their linguistic environment, there is no way of knowing exactly what any individual child registers. Therefore, there is no factual basis for the claim that children register only what is filtered for them through parents' deliberately simplified speech. Children have access to more than this, including defective utterances, sentence fragments, and all the imperfections that the world throws at us. Second, even supposing that they register only perfectly formed expressions (and hence that the first data deficiency does not hold), this isn't enough to show that the child has a sufficient inductive base for language acquisition. The third data deficiency still holds, and the child would need to know that *wanna* fails to occur in certain contexts. If children learn by induction, we must ask why quite ordinary inductive generalizations like this – that *want to* may be pronounced *wanna* – break down. The artificial, simplified speech of the Motherese style does not show where inductive generalizations must stop. Third, if the child registered only the simplified, well-formed sentences of Motherese, the problem of language learning would be *more* difficult, because the child's information would be more limited. Fourth, careful studies of parents' speech to children (like Newport et al. 1977) show that an unusually high proportion consists of questions and imperatives, and that simple declarative sentences are much rarer than in ordinary speech. This suggests that there is very little correlation between the way the child's language emerges and what parents do in their speech directed to children. The existence of Motherese in no way eliminates the need for a genetic basis to language acquisition. The child is primarily responsible for the acquisition process, not parents or playmates.

Nobody denies that the child must extract information from her environment; it is no revelation that there is "learning" in that technical sense. My point is that there is more to language acquisition than this.

Children react to evidence in accordance with specific principles. It is not at all clear what role induction plays. Induction does not enable a child to determine what a well-formed sentence is; nor does it explain how children "learn" the meanings of even the simplest words. Children do not have sufficient evidence to induce the meaning of *house*, *book*, or *city*, or of more complex expressions, even if we grant everything to advocates of Motherese or those who argue that it's all data processing of huge corpora (see n. 5). The biggest problem with ordinary sentences like *Kim is too clever to catch* or *Everybody saw him* is not that they exist, but to characterize their meaning and the meanings of the individual words. There is another problem with the Motherese hypothesis, which is more trivial than this but nonetheless real: children typically acquire the language not of their parents, but of their older siblings and peers.

The problem demanding explanation is compounded by other factors. Despite variation in background and intelligence, people's mature linguistic capacity emerges in fairly uniform fashion, in just a few years, without much apparent effort, conscious thought, or difficulty; and it develops with only a narrow range of the logically possible "errors." Children do not test random hypotheses, gradually discarding those leading to "incorrect" results and provoking parental correction. In each language community the non-adult sentences formed by very young children seem to be few in number and quite uniform from one child to another, which falls well short of random hypotheses. Normal children attain a fairly rich system of linguistic knowledge by five or six years of age, and a mature system by puberty. This is impressive when compared with the laborious efforts of squads of adult linguists who have made it through graduate school and try with only moderate success to characterize explicitly what people know when they know Estonian, Hopi, or Bengali. In this regard, language is no different from, say, vision, except that vision is taken for granted, and ordinary people give more conscious thought to language.

These, then, are the salient facts about language acquisition or, more properly, language growth. The child masters a rich system of knowledge without significant instruction and despite a triple deficiency of experiential data. The process involves only a narrow range of "errors" or false hypotheses and takes place rapidly, even explosively, between two and three years of age. The main question is how children acquire so much more than they experience.

Given these facts, especially the third data deficiency, certain properties must be available to the organism independently of linguistic experience. These properties permit language growth to circumvent environmental deficiencies and thus to take place quickly and not solely by trial and error. The environmental stimulus is impoverished, unstructured, and fairly random; the child hears a haphazard selection of sentences and pseudo-sentences and receives no significant instruction. So the environmental stimulus is just a trigger. Much of the ability eventually attained is determined by genotypical principles, which are activated by environmental stimuli.

The further fact that children can master *any* human language to which they happen to be exposed in infancy imposes strong limitations on the principles which the scientist can attribute to the genotype. An answer to the problem of acquisition does not lie merely in the properties of the specific language a particular child acquires. This would amount to a claim that the specific properties of, say, Hindi are innately prescribed, which permits no explanation of how a language with a significantly different structure – Polish, say – is acquired. The native principles must therefore be somewhat abstract and not language-specific. The principles must meet strict empirical demands: for each child, the principles must provide a basis for attaining the eventual mature grammar given exposure only to the haphazard and unstructured experience available in the child's community. To require the native principles to give an account of how a child may master *any* natural grammar under the conditions noted is to ask a lot.

The tight empirical demands make language particularly useful as a probe into the intrinsic properties of the human mind/brain.

A grammar represents what a speaker comes to know, subconsciously for the most part, about his or her native language. It represents the fully developed linguistic capacity, and is therefore part of an individual's phenotype. It is one expression of the potential defined by the genotype. Speakers know that certain sentences (in fact, an infinite number) may occur in their speech and that others may not; they know what the occurring sentences mean and the various ways in which they can be pronounced and rephrased. Most of this largely subconscious knowledge is represented in a person's grammar. The grammar may be used for various purposes, from everyday functions like expressing ideas, communicating, or listening to other people, to more contrived functions like writing elegant prose or lyric poetry, or compiling and solving crossword puzzles, or writing a book about language acquisition and change.

Universal Grammar (UG) represents the genetic equipment that makes language growth possible under the conditions assumed here (therefore part of the genotype) and delimits the linguistic knowledge that may eventually be attained – that is, the form and functioning of the grammar. The genotypical principles and parameters can be viewed as a theory of grammar. On this view, the theory of grammar is, in Chomsky's words,

> a common human attribute, genetically determined, one component of the human mind. Through interaction with the environment, this faculty of mind becomes articulated and refined, emerging in the mature person as a system of knowledge of language. To discover the character of this mental faculty, we will try to isolate those properties of attained linguistic competence that hold by necessity rather than as a result of accidental experience, where by "necessity" I of course mean biological rather than logical necessity. . . . The commitment to formulate a restrictive theory of UG is nothing other than a commitment to discover the biological endowment that makes language acquisition possible and to determine its particular manifestations. . . . we can explain some property of attained linguistic competence by showing that this property necessarily results from the interplay of the genetically-determined language faculty, specified by UG, and the person's (accidental) experience. (Chomsky 1977: 164)

It need hardly be pointed out that there is nothing necessary or God-given about this research goal; nor do I want to give the impression that all linguists adopt it. Crucially for our purposes, the vast majority of people who have worked on language change have not adopted this research goal. In fact, people have studied language with quite different goals in mind, ranging from the highly specific (to describe Dutch in such a way that it can be learned easily by speakers of Indonesian), to the more general, such as showing how a language may differ from one historical stage to another (comparing, for example, Chaucerian and present-day English). However, this is the goal I adopt, and this is the sense in which I shall construe a grammar, seeing it as a biological object; and it is important to keep this idea straight and not conflate it with other, more traditional notions of what a grammar is.

3.4 The Analytical Triplet

A grammar, for us, is a psychological entity, part of the psychological state of somebody who knows a language. For any aspect of linguistic knowledge, three intimately related items are included in the account. First, there is a formal, explicit characterization of what a mature speaker knows; this is the *grammar*, which is part of that speaker's phenotype. It is an internal system, what Chomsky 1986 called the I-language, as distinct from the external linguistic production, the E-language. Since the grammar is represented in the mind/brain, it must be a finite system, which can relate sound and meaning for an infinite number of sentences.

Second, also specified are the relevant principles and parameters common to the species and part of the initial state of the organism; these principles and parameters make up part of the *theory of grammar*, or *Universal Grammar*, and belong to the genotype.

The third item is the *trigger* experience, which varies from person to person and consists of an unorganized, fairly haphazard set of utterances, of the kind that any child hears (the notion of a trigger stems from ethologists' work on the emergence of behavioral patterns in young animals). The universal theory of grammar and the variable trigger together form the basis for attaining a grammar; grammars are attained on the basis of a certain trigger and the genotype.

In (9) I give the explanatory schema, using general biological terminology in (9a) and the corresponding linguistic terms in (9b). The triggering experience causes the genotype to develop into a phenotype; exposure to a range of utterances in, say, English allows the UG capacity to develop into a particular mature grammar. One may think of the theory of grammar as making available a set of choices; the choices are taken in the light of the trigger experience, or the primary linguistic data (PLD), and a grammar emerges when the relevant options are selected. A child develops a grammar by setting the parameters of UG in the light of her particular experience.

9 (a) linguistic triggering experience (genotype → phenotype)
  (b) primary linguistic data (Universal Grammar → grammar)

Each of the items in the triplet – trigger, UG, and grammar – must meet various demands. The trigger, or PLD, must consist only of the kinds of things that children routinely experience and includes only simple structures.[6] The theory of grammar, or UG, is the one constant and must hold universally, so that any person's grammar can be attained on the basis of naturally available trigger experiences. The mature grammar must define an infinite number of expressions as well-formed, and for each of these it must specify at least the sound and the meaning. A description always involves these three items, which are closely related; changing a claim about one of the items usually involves changing claims about the other two. This tight, ambitious system of description must meet many empirical demands. It is hard to imagine that we might have to choose between two or more descriptions which meet all empirical demands, hard to imagine problems of indeterminacy of the kind that plague the natural historian or the taxonomist who does not take a psychological view of grammars.

The grammar is one sub-component of the mind, which interacts with other cognitive capacities or modules. Like the grammar, each of the other modules is likely to develop in time and to have distinct initial and mature states. So the visual system recognizes triangles, circles, and squares through the structure of the circuits that filter and recompose the retinal image (Hubel and Wiesel 1962). Certain nerve cells respond only to a straight line sloping downward from left to right, other nerve cells to lines sloped in different directions. The range of angles that an individual neuron can register is set by the genetic program, but experience is needed to fix the precise orientation. In the mid-sixties David Hubel, Torsten Wiesel, and their colleagues devised an ingenious technique to identify how individual neurons in an animal's visual system react to specific patterns in the visual field (including horizontal and vertical lines, moving spots, and sharp angles). They found that particular nerve cells were set within a few hours of birth to react only to certain visual stimuli, and, furthermore, that if a nerve cell is not stimulated within a few hours, it becomes totally inert in later life. In several experiments on kittens, it was shown that if a kitten spent its first few days in a deprived visual environment (a tall cylinder painted only with vertical stripes), only the neurons stimulated by that environment remained active; all other optical neurons became inactive, because the relevant synapses degenerated, and the kitten never learned to see horizontal lines or moving spots in the normal way. Thus learning is a selective process: parameters are provided by the genetic equipment, and relevant experience