Meno's paradox and the acquisition of grammar

B. Elan Dresher

> From our point of view as English speakers, a language such as Chinese might seem totally different from our own. In fact, these two languages as well as all other human languages are nearly identical. The differences that seem all important to us are relatively minor. (Jonathan Kaye 1989: 54).

1. Plato's problem and Meno's paradox

In Plato's dialogue *The Meno*, Meno doubts that one can investigate what one does not know. Which of the things you do not know, he asks, will you propose as the object of your search? Even if you stumble across it, how will you know it is the thing you did not know? Socrates replies that there is a way out of this paradox: we can investigate what we do not know because, at some level, we already know everything; what we call learning is but recollection. He goes on to demonstrate the truth of this astonishing claim by showing that an ignorant slave boy actually knows the Pythagorean theorem, even though the boy does not know that he knows it, and in fact does not seem to know it until Socrates leads him through a series of questions about it. The implication is that Meno's paradox would indeed make learning impossible, unless we assume that we have knowledge from some source other than experience in this life.

More recently, Noam Chomsky has observed that the problem of how we come to know things remains with us. He has named this *Plato's problem* (Chomsky 1988: 4). In the words of Bertrand Russell (Russell 1948: 5), the problem is this: "How comes it that human beings, whose contacts with the world are brief and personal and limited, are able to know as much as they do know?" As Chomsky has shown over the years, this problem arises sharply in the case of language. For it can be shown that everybody is like Meno's slave boy, in that they know many things about their native language that they don't know they know; moreover, these are things they were never taught; nor does it appear that they could have had any other experience, in this life, that could suffice to account for their knowledge.

2.   Some examples of Plato's problem in language

To give a brief example (Chomsky 1975: 30-35): every native speaker of English knows how to form yes/no questions. Given a declarative sentence like (1a), the corresponding question is (1b); similarly, the corresponding question to (2a) is (2b):

(1)   a.   *The boy is tall.*
       b.   *Is the boy tall?*

(2)   a.   *Mary has been swimming.*
       b.   *Has Mary been swimming?*

It is clear that every speaker of English can do this for any declarative sentence, so this ability is not just a matter of memorising a long list of questions. Rather, every speaker of English must have at some point acquired a general rule for creating such questions. Some aspects of this rule must be based on experience, since not every language forms questions in the same way. We might imagine a child, on the basis of being exposed to simple questions of the sort mentioned earlier, (unconsciously) formulating a rule such as the following: to form a yes/no question, move the first auxiliary verb of the corresponding declarative to the front of the sentence. (This formulation presupposes that the learner has figured out what auxiliary verbs are.)

   The rule we have given, though, is not exactly correct. When we consider a more complex example, we see that we do not always move the first auxiliary verb to the front. Consider (3a). Applying our rule, we would move the first auxiliary verb *was* to the front, deriving the incorrect question (3b). This is of course totally ill-formed. What every speaker of English knows is that the correct question in this case is (3c). That is, one has to skip the first auxiliary verb *was*, which is in a subordinate clause *who was in the room*, and pick out the auxiliary verb of the main clause, which is *is*, and move that to the front. So the rule is that we must move the first auxiliary verb of the main clause to the front, not the first auxiliary verb in the sentence.

(3)   a.   *The man who was in the room is tall.*
       b.   **Was the man who in the room is tall?*
       c.   *Is the man who was in the room tall?*

If children learn this rule by observation, through hypothesis formation or trial and error, we might expect to find a learning sequence such as the following: on the basis of simple sentences, children arrive at the simple rule, "Move the first auxiliary verb to the front." Later on, when they start producing more complex sentences with subordinate clauses, we expect them to make mistakes, where they move the wrong verb to the front. Observations of children, and experiments designed to test how they in fact deal with such cases, reveal, to the contrary, that children never make such mistakes. They make all sorts of other mistakes, but they never try to move a verb from a subordinate clause rather than a main clause (Crain and Nakayama 1987). It follows, then, that children never entertain what appears to be the simplest hypothesis. Rather, they appear to know from the outset that rules of grammar that move elements around are sensitive to clause structure.

Here, then, is an example of Plato's problem, and Chomsky proposes a version of Plato's solution: the knowledge that rules of grammar are sensitive to structure is not gained in this lifetime, but sometime before, and is a part of our genetic inheritance. Supporting this idea is the further observation that the principle of sensitivity to structure appears to be universal, true of every human language.

Such cases are not limited to syntax, but arise in all aspects of language acquisition. Consider how one learns the meaning of words. In *Word and Object* (1960: 29–54), W.V.O. Quine made up a story about field linguists studying a completely unfamiliar language. They observe that a native says *gavagai* when a rabbit runs past, and guess that *gavagai* means 'rabbit'. But Quine observes that there are many other possibilities. *Gavagai* could be a whole sentence, such as 'Lo, there goes a rabbit!'. And even if the linguists are able to learn enough of the language to determine that *gavagai* is a word, not a sentence, Quine points out that it will still not be certain that *gavagai* means 'rabbit'. For it might mean 'temporal rabbit stages'; or 'all and sundry undetached rabbit parts'; or it might refer to 'that single though discontinuous portion of the spatio-temporal world that consists of rabbits'; or 'rabbithood'; or 'a momentary leporiform image'; and so on. Quine observes that all these possible meanings, and infinitely many more, are very hard to distinguish from each other and from plain 'rabbit', in all kinds of tests and situations — i.e., most things which are true of 'rabbit' are true of 'undetached rabbit parts'. So if we ask a native speaker, "Is this [pointing to something] an ex-

ample of *gavagai*?" the answer will be "yes" or "no" (or as Quine suggests, *evet* or *yok*) whether *gavagai* means 'rabbit' or some other of the above candidate interpretations.

Now, it is interesting to notice that linguists or travellers do not in actual practice encounter the *gavagai* problem when learning an unknown language; nor do children, since they arrive at the same meanings of words as their parents, siblings, and other members of the same speech community. How can this be? It must be the case that everybody draws similar conclusions about what a word means, and what a likely meaning is. That is, the odds are that *gavagai* means 'rabbit' and not 'a momentary leporiform image'.

As with syntax, children do make many mistakes in the course of learning what words mean. For example, they may overgeneralise or undergeneralise the meanings of words. They might, for example, suppose that the word *dog* applies also to horses and other such animals. But again, the misgeneralisations children actually make are very limited in comparison with the ones it is possible to make. For example, when a dog enters the room and someone says "Dog," no child supposes that *dog* means 'animal viewed from the front', or 'animal that has just entered a room'. Again, it appears that learners are constrained to entertain only a small subset of the conceivable hypotheses.

The *gavagai* problem does arise, however, when we try to decode the meanings of the vocalisations of other species. For example, Cheney and Seyfarth (1990), after much observation of vervet monkeys, are unable to decide if their leopard call means 'Behold! a leopard!', or 'Run into the trees!', or a range of other possibilities. It is unlikely that the language of vervet monkeys is so much more complex than our own; vervet monkeys do not do well at learning English yes/no questions, for example. But lacking the appropriate built-in constraints, we are at a loss to arrive at the correct answer, which must be obvious to any vervet monkey.

Acquisition of phonology, the sound system of a language, is also not just a question of memorising sounds, but involves learning patterns. Consider how words are stressed in English, for example. English stress, though complex and subject to numerous exceptions and special cases, follows certain patterns that we have come to know at some unconscious level. The reality of these patterns is made manifest when we borrow into the language a word that does not observe them. For example, the Russian word *bábushka* is more usually pro-

nounced by English speakers as *babúshka*. Similarly, the capital of Finland, *Hélsinki*, is frequently pronounced *Helsínki* in English. The shift in stress brings these words into conformity with English stress patterns.

If stress patterns of languages all involved simple generalisations, we could think that it would suffice to learn them from experience. But some extremely complex generalisations are quite common cross-linguistically, whereas myriad others of lesser or comparable complexity that one could imagine do not occur. For example, here is the rule for assigning main stress in Passamaquoddy, an Algonquian language spoken in Maine and New Brunswick (Stowell 1979, based on data from Philip LeSourd):

(4)   a. If the penultimate syllable of a word has a full vowel, stress it.

          b. Otherwise, if the antepenultimate syllable has a full vowel, stress it.

          c. Otherwise (if neither the penult nor antepenult has a full vowel), stress whichever one of these two syllables is separated by an even number of syllables from the last preceding syllable that has a full vowel, or — if there is no full vowel — from the beginning of the word.

This type of stress system is not at all rare; many languages have rules that are similar, with minor variations. As before, we seem led to the conclusion that the learner is constrained to look for particular types of patterns.

The idea that language learners must be benefiting from innate direction of some kind is supported further when we consider under what conditions language learners (here we are talking about small children) must master these complex and subtle linguistic generalisations. Unlike linguists, who can bring to bear on these questions all sorts of evidence, children must learn the rules of their language from exposure to whatever examples come their way, without any explicit discussion of what the rules are, or any systematic ordering of examples. Imagine trying to learn how to play chess by observing other people play, where nobody tells you explicitly what the rules are for moving a knight or for taking a pawn en passant; suppose moreover that some small percentage of moves you observed are actually illegal, but unnoticed, hence unremarked, and that many games are bro-

ken off or left incomplete for a variety of reasons, usually unstated; suppose also that grown-ups let you make many illegal moves without correcting you, because you are little and don't know better; and now suppose that the game you are learning is not chess but one that is many times more difficult. How much more difficult? Recall that a few years ago a computer beat the highest ranked chess player in history; but no computer can speak a language in a way that can even be compared to the most inarticulate human speaker, nor are there any prospects of creating such a program in the foreseeable future.

3.   Parameters in Universal Grammar

The above remarks, then, are intended to persuade you that Plato's problem arises with great force in the acquisition of language: people know things about their language that seem to go beyond their experience. It follows, then, that their knowledge of these things does not come from experience, but from their own minds. Following Chomsky, let us call these innate cognitive principles, whatever they are, *Universal Grammar*. What these principles are is an open question that is still far from solved; but the hypothesis that there is a Universal Grammar is, so far, the only hypothesis that offers any hope of solving Plato's problem in the domain of language.

Assuming that grammar is universal, though, presents us with a new problem, and that is to explain why languages are not all the same. The most striking way in which languages differ, of course, is in how sounds are paired with meaning to form words. A tree is denoted by the word *tree* in English, but by *arbre* in French and by *baum* in German. This is the phenomenon called the *arbitrariness of the sign* by de Saussure; that is, there is no inherent connection between the meaning of a word and its sound. It follows that the words of each language have to be learned by experience. Fortunately, Plato's problem does not arise here: English children have lots of opportunities to observe that the word for tree is *tree*.

Apart from vocabulary, there are other sources of cross-language variation that must be accounted for. For example, in some languages a verb precedes its object, while in others the verb follows its object. In some languages, like Italian, a subject is optional, whereas in others, like English, it is obligatory, even if has no semantic role, such as in the sentence *It's raining*. In Passamaquoddy, the rules of stress as-

signment distinguish between full vowels and reduced vowels; in other languages, similar distinctions are drawn between long vowels and short vowels, or between open syllables and closed syllables. These cross-language differences have the flavour of being variations on a theme. To account for such variation, Chomsky (1981b) has proposed that the principles of Universal Grammar are not rigidly fixed, but allow for parametric variation. These parameters are like open variables in a formula, whose value is to be set on the basis of experience. The possible values of a parameter are limited and given in advance, like choices on a menu that allows very limited substitutions. On this view, then, language acquisition is reduced to setting parameters to their appropriate values.

4.   Parameter setting

Parameter setting is a much more manageable learning problem than open-ended induction, or hypothesis formation and testing. Most research into Universal Grammar has quite appropriately concentrated on trying to figure out what the principles and parameters are. There is not much we can say about the learning problem until we know what it is that has to be learned. Nevertheless, a small subfield devoted to parameter setting has arisen in the last fifteen years which has begun to consider different models of how learners might go about setting parameters. The problem has turned out to be deceptively difficult, much more than one might have thought.

On the face of it, parameter setting would appear to be a simple matter. Suppose there is a principle of Universal Grammar that says that in each language the subject of a sentence may either precede or follow the verb — the learner must decide which it is. In English, the subject precedes the verb, as in the sentence *John kicked the ball*. So one might think that a learner just has to hear a few simple sentences like this to know how to set this parameter. But it's not that simple. The parameter fixes what we could call the canonical, or basic, word order of a sentence; however, in many languages the basic order can be disturbed by other rules that move elements around. English has constructions in which the subject follows the verb, as in *Hello, said Peter to his friend*. Further, in order to set the parameter correctly, we have to know what the subject and object in any given sentence are, but it is not obvious that a learner always knows this. In an impera-

tive sentence like *Watch yourself!* there is no overt subject; from the meaning, a language learner might mistakenly conclude that the subject is *yourself*, which follows the verb.

Cases of this kind occur in other languages. In French, for example, the object follows the verb, as in *Paul voit la table*. But when the object is a pronoun, it appears before the verb: *Paul la voit*. The explanation for this, as we understand it, lies in the fact that French pronouns are clitics, and clitics have special positions that are not the canonical ones. However, a learner might not know that, and be misled by this example.

In Dutch and German, it can be shown that the verb follows its object in the basic word order. However, this basic order can be observed mainly in subordinate clauses. In main clauses, the basic order is disturbed by a principle that requires the verb to occur in second position in the sentence.

Jonathan Kaye and I worked on a computational model for learning how to assign stress to words. We found that the relation between a parameter and what it does is rather indirect, due to the fact that there are many parameters, and they interact in complex ways. For example, in English main stress is tied to the right edge of the word. But that doesn't mean that stress is always on the last syllable, as in *chandelíer*. It could be on the penultimate syllable, as in *Manitóba*, or even on the first syllable, as in *Cánada*. How can *Cánada* be an example of stress on the right? In English, stress is assigned not to individual syllables, but to groupings of syllables called *feet*. An English foot is a *trochee*, consisting of a stressed syllable followed by an optional unstressed syllable. Other languages have other kinds of feet, depending on how the foot parameter is set. But then why are words like *Cánada* and *álgebra* stressed on the third to last syllable and not on the second to last? That is due to another parameter to the effect that the last syllable may be ignored in assigning stress. Why, then, are some words stressed on the second to last syllable, like *aróma* and *agénda*? That is due to their different syllable structures. In English, syllables with long vowels or closed by a consonant are considered to be *heavy*, like full vowels in Passamaquoddy, and heavy syllables tend to attract stress. A heavy syllable can be a foot by itself. Thus, words that have a heavy penultimate syllable have stress on that syllable, as in *aróma, Helsínki, babúshka*; words with a light penult have stress on the antepenult, as in *Cánada, álgebra*.

The point of all this is that the stress patterns of any language are the result of a number of interacting parameters. This interaction makes the relationship between a parameter and its effects non-transparent. Some surprising consequences follow from this fact.

The first one is that learners who have some incorrectly set parameters might know that something is wrong, but might not know which parameter is the source of the problem. Suppose, for example, that a learner of English mistakenly thinks that objects must precede the verb. The learner observes the sentence *John kicked the ball*. According to the learner's developing grammar, that should have been *John the ball kicked*. So the learner realises that something is wrong with the grammar: one or more parameters are set to the wrong values. But which ones? It could be the parameter that says whether the object should precede or follow the verb. But it could be something else entirely. Maybe English is like Dutch in requiring the verb to move to second position. Maybe that's the parameter that has to be adjusted, and not the word order parameter. This is known as the *credit problem*: a learner cannot reliably assign credit or blame to individual parameters when something is wrong.

There is a second way in which parameters can pose problems to a learner, somewhat reminiscent of Meno's paradox. When we talk about specific parameters, we presumably know exactly what they do, and we assume that learners ought to know what they do, also. But it is not obvious that this is the case, and indeed, it looks quite certain that this cannot always be the case. For some parameters are stated in terms of abstract entities and theory-internal concepts which the learner may not initially be able to identify. For example, the theory of stress is couched in terms of concepts such as heavy syllables, heads, feet, and so on. In syntax, various parameters have been posited that refer specifically to anaphors, or to functional projections of various types. These entities do not come labelled as such in the input, but must themselves be constructed by the learner. So, to echo Meno, how can learners determine if main stress falls on the first or last foot if they don't know what a foot is, or how to identify one? How can learners set parameters that control where anaphors can appear when they don't know which parts of the data represent anaphors? And if they happen to stumble across an anaphor, how will they know that that's what it is? This can be called the *epistemological problem*.

5.   Some parameter setting learning models

To summarise, to get out of Meno's paradox and to solve Plato's problem, we posited a theory of Universal Grammar with a set of open parameters. By doing so, we limit the role of experience to parameter setting. But now we have found that the same problems arise even in this limited domain. How do we solve them? Different learning algorithms have taken radically different tacks in trying to deal with the credit problem and the epistemological problem.

5.1. *A cue-based learner (Dresher and Kaye 1990)*

Our proposal is to put even more into the mind: not only the principles and parameters of Universal Grammar are innate, but learners must be born with some kind of a road map that guides them in setting the parameters. Some ingredients of this road map are as follows.

First, Universal Grammar associates every parameter with a *cue*, something in the data that signals the learner how that parameter is to be set. The cue might be a pattern that the learner must look for, or simply the presence of some element in a particular context.

Second, parameter setting proceeds in a (partial) order set by Universal Grammar: this ordering specifies a learning path (Lightfoot 1989). The setting of a parameter later on the learning path depends on the results of earlier ones.

Hence, cues can become increasingly abstract and grammar internal the further along the learning path they are. For example, in learning stress patterns, we suppose that children are able to recognise relative stress levels, and can tell if words with different types of syllables have the same stress patterns or different patterns. Such a developmental stage may have representations as in *Figure 1*. In these simple representations of English stress contours, each syllable is represented by *S*, and the height of the column of *x*'s indicates the relative perceived stress level of each syllable.

```
a. América              b. Mànitóba              c. agénda
    x                            x                      x              Line 2
    x                       x    x                      x              Line 1
x   x   x   x            x   x   x   x            x   x   x            Line 0
S   S   S   S            S   S   S   S            S   S   S            Syllables
A  me  ri  ca           Ma  ni to:ba             a gen da
```
Figure 1. Representations of English stress contours before setting metrical parameters

As learners acquire more of the system, their representations become more sophisticated, and they are able to build on what they have already learned to set more parameters, eventually acquiring the representations in *Figure 2*.

```
a. América              b. Mànitóba              c. agénda
    x                            x                      x              Line 2
   (x)                    (x      x)                   (x)             Line 1
x (x   x) <x>            (x   x) (x) <x>          x (x) <x>            Line 0
L   L   L   L             L   L   H   L            L   H   L           Syllables
A  me  ri  ca           Ma  ni to:ba             a gen da
```
Figure 2. Acquired representations

In these representations, the undifferentiated *S* has been replaced by *L*, which indicates a light syllable, and *H*, which stands for a heavy syllable. Scanning from right to left, syllables have been grouped into binary left-headed (trochaic) feet, and final syllables are marked as extrametrical (<x>).

As an example, consider (a), the word *America*. The learner has found that the word consists of four light syllables. On line 0, metrical structure is constructed as follows: start at the right end of the word, and skip the last syllable. Then group the preceding two syllables into a foot. Since there is only one light syllable left over, it does not get put into a foot. In (b), *Manitoba*, we skip the last syllable. Since the second to last syllable is heavy, we do not group it with the preceding syllable but it makes a foot by itself. The two light syllables of *Mani* are grouped into a second foot.

An *x* on line 1 indicates a stress, which is the head of a foot. Since English feet have their heads on the left, in (a) the *x* goes on the second syllable. In (b) there are two feet, (Ma ni) and (to:), and they are grouped together on line 1. Which foot is stronger? In English, main

stress goes on the rightmost foot, so on line 2 we put an *x* over the third syllable in *Manitoba* — this is the main stress in the word.

In the learning model proposed by Dresher and Kaye (1990), the representations of *Figure 1* are transformed gradually into those of *Figure 2*, as learners set the metrical parameters that generate these representations. In this learning algorithm, called YOUPIE and modelled on a computer in PROLOG, the learner puts these representations together in the order given in (5). In (6) I list what each parameter is, and what cues the learner looks for to set these parameters.

(5)   Order in which parameters must be set

1.	Syllable Quantity	Establishes whether feet are quantity insensitive (default, henceforth "QI") or quantity sensitive ("QS") (and type of QS).
2.	Extrametricality	Establishes edge of domain; can only exclude it at this point.
3.	Foot size:	If QI, only bounded feet available; if QS, unbounded is default.
4.	Main stress:	Depends on correct settings of all the above.
5a.	Headedness:	Sometimes depends on having set main stress.
5b.	Directionality:	Cannot be determined apart from headedness.
6.	Destressing:	Determined by comparing stresses predicted by above parameter settings with actual stresses.

(6)   Parameters and cues

1. Syllable Quantity

a.	Parameter:	The language {does not/does} distinguish between light and heavy syllables (a heavy syllable may not be a dependent in a foot).
b.	Cue:	Words of *n* syllables, conflicting stress contours, indicates QS.

2. Extrametricality
    a. Parameters: A syllable on the {right/left} {is not/is} extrametrical.
    b. Cue: Stress on a peripheral syllable rules out extrametricality on that side.

3. Bounded constituent construction
    a. Parameter: Line 0 constituents are bounded.
    b. Cue: The presence of a stressed non-edge *L* indicates bounded constituents.

4. Main stress
    a. Parameter: Project the {left/right}-most element of the line 1 constituent.
    b. Cue: Scan a constituent-sized window at the edge of a word. Main stress should consistently appear in either the left or right window.

5. Headedness and directionality of feet
    a. Parameters: {Left/right}-headed feet are constructed from the {left/right}.
    b. Cue: Scanning from the {left/right}, a light syllable {following/preceding} any other syllable must be unstressed.
    c. Example: Scanning from the left, if for all (*X L*), *L* is unstressed, then direction = Left, Headedness = Left. If for all (*L X*) *L* is unstressed, then headedness = Right.

6. Destressing (conflates a number of separate parameters)
    a. Parameters: {Various types of} feet are destressed in {various situations}.
    b. Main Cue: The absence of stress on a foot.
    c. Example: The lack of stress on the first syllable of *agénda*, with intermediate acquired foot structure (à)(gén)<da>, shows that this foot is destressed (further cues reveal the conditions under which this occurs).

Space does not allow us to go through all of these parameters and cues in detail, but we can look at at the first one to get some sense of how this works.

The first parameter the learner tries to set is syllable quantity: does the language treat all syllables the same with respect to stress, or is there a distinction between heavy and light syllables? If all syllables are equal, then we say that the stress system is *quantity insensitive*, or QI. If the quantity or weight of a syllable is important, then the stress system is *quantity sensitive*, or QS.

How can a learner set this parameter? Even if you don't know anything about the stress system, you can still keep track of how many syllables words have and where stresses fall. In a language in which syllables are all treated equally, then every word of *n* syllables should be stressed in the same way. But if stress is quantity sensitive, then heavy syllables will not be treated the same as light syllables, and words of the same length can have different stress patterns. So we propose that this is the cue for quantity sensitivity: if you find conflicting stress patterns in words of the same length, you have QS; if you do not find this, you stick with QI, which we assume is the default setting. English is QS, because words of the same length are not all stressed the same way, as we have seen.

Once this parameter has been determined, the learner can use information about syllable quantity to set further parameters, and proceeding in this way, can arrive at the final representations in *Figure 2*. See Dresher and Kaye 1990 and Dresher 1999 for a detailed description of the cues and parameters assumed here, as well as the order in which they are set.

This approach has something of a Piagetian flavour, with later stages depending on and building on what was acquired in earlier stages; but whereas Piaget supposed that later stages are literally invented out of earlier ones, without being innately specified — a position that creates a seemingly intractable mystery (see Piattelli-Palmarini 1980) — in the view sketched here the whole sequence is innately specified .

If this approach is correct, there is no parameter-independent learning algorithm.

5.2. The Triggering Learning Algorithm (Gibson and Wexler 1994)

A different approach is taken by Gibson and Wexler (1994: 409–410), who characterise what they call the Triggering Learning Algorithm (TLA) as follows:

(7)   Triggering Learning Algorithm (Gibson and Wexler 1994)

"Given an initial set of values for *n* binary-valued parameters, the learner attempts to syntactically analyze an incoming sentence *S*. If *S* can be successfully analyzed, then the learner's hypothesis regarding the target grammar is left unchanged. If, however, the learner cannot analyze *S*, then the learner uniformly selects a parameter *P* (with probability $1/n$ for each parameter), changes the value associated with *P*, and tries to reprocess *S* using the new parameter value. If analysis is now possible, then the parameter value change is adopted. Otherwise, the original parameter value is retained."

We can illustrate how this learning algorithm is supposed to work by looking at the diagram in Figure 3, where each square represents a setting of two syntactic parameters. The first parameter determines whether the head of Spec X′ is initial (value 1) or final (0). In this case, the head is the verb (*V*) and its specifier is the subject (*S*). The second parameter encodes whether the head of a complement is initial or final, here exemplified by the relation between a verb and its object (*O*). These two parameters define a space with four states.
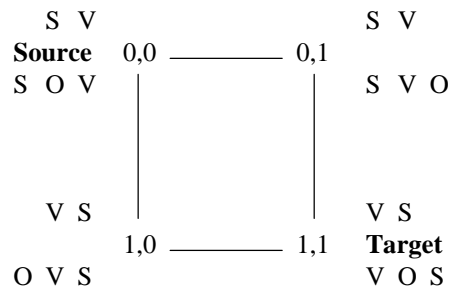


```
   S  V                    S  V
 Source   0,0 ──────── 0,1
   S  O  V      │           │      S  V  O
                │           │
                │           │
                │           │
   V  S         │           │      V  S
            1,0 ──────── 1,1   Target
   O  V  S                    V  O  S
```

Figure 3. Parameter space (Spec-Head f/i, Comp-Head f/i): final = 0, initial = 1

Assume now that the target language is VOS (1,1), and that the learner's current hypothesis is SOV (0,0). Suppose the learner hears a sentence of the form *V O S*. This sentence is not parsable by the learner, who now determines that the current state is not correct.

Even though there is only one setting of parameters that corresponds to *V O S*, it would take a change of both parameters for the learner to reach it. This is not allowed by the Triggering Learning Algorithm, which makes available only the two neighbouring spaces. Neither space yields the target *V O S*. Therefore, the learner cannot move. Thus, the sentence *V O S* is not a *trigger* to a learner at (0,0). Fortunately, in this case there is another type of sentence from the target that the learner will eventually hear, namely *V S*. *V S* is a trigger to a learner at (0,0), since there is a neighbouring space which parses it, namely (1,0). So the learner moves to there. From there, a further presentation of *V O S*, which is a trigger to a learner at (1,0), will take the learner to the target.

A learner following the Triggering Learning Algorithm does not know what any parameter does, it simply tries to match input forms by moving to a parameter space that can parse a given input form. The Triggering Learning Algorithm runs into a number of serious problems:

1. The Triggering Learning Algorithm cannot handle parameters in subset relations.
2. The learner can fall into incorrect grammars from which it cannot escape.
3. Or the learner can thrash around indefinitely, repeatedly revisiting the same incorrect grammars.
4. Learning follows a sequence dictated by accidents of input data; there is no notion of development toward greater complexity.
5. The learner will be unable to match the input perfectly at early stages of acquisition, so the learning process cannot get off the ground.

To see how the Triggering Learning Algorithm would apply to more than three parameters, I generated a set of schematic languages using six metrical parameters. There were thus 2^6 (64) languages. Each language was assigned four two-syllable words, eight three-syllable words, sixteen four-syllable words, and ten five-syllable words. Thus, each language has 38 words. Four pairs of languages are extensionally equivalent: their surface stress patterns are identical, though their grammars assign different structures. Since a learner would have no evidence to decide which grammar is correct, these languages are excluded as target grammars from the following discussion. An analysis

of how the Triggering Learning Algorithm would apply to the remaining 56 languages yields the results in *Table 1*.

Table 1. 64-state Triggering Learning Algorithm, 6 metrical parameters (56 target states tested)

Number Of targets	Safe states	Problem states		
		Local maxima	Cul-de-sacs	Dangerous states
2	16	4	8	36
2	40	13	9	2
2	42	3	1	18
2	48	1	11	4
2	48	1	1	14
2	48	0	6	10
2	52	12	0	0
6	55	9	0	0
8	55	6	0	3
2	56	2	0	6
26	64	0	0	0

In this table, *local maxima* are states (excluding the target itself) from which the learner cannot exit, and *cul-de-sacs* are states that do not connect to the target, though exit is possible to one or more dead-end states. A learner who arrives at any of these states is guaranteed to fail to reach the target. A *dangerous* state is a state that connects to a local maximum or cul-de-sac, as well as to the target. Although a learner in a dangerous state has a chance of reaching the target, success is not guaranteed. In terms of the goal of a learning theory for language, all of these states are *problem* states. *Safe* states are states that do not connect to any problem states; assuming that each triggered transition from a safe state has some probability greater than zero, arrival at the target is guaranteed in the limit. We find that 26 languages have Triggering Learning Algorithms with no problem states, whereas 30 languages have between 8 and 48 problem states. In other words, even though there are no subset relations in the data set, and all languages have the same number of words, nearly one half of the languages cannot be guaranteed to be learnable by the Triggering Learning Algorithm.

5.3. A genetic algorithm (Clark and Roberts 1993)

A different type of learning algorithm was proposed by Robin Clark, and applied in Clark and Roberts (1993) in connection with the loss of V2 in French. On this model, parameter setting proceeds by way of a genetic algorithm that enacts a Darwinian competition of survival of the fittest. A learner simultaneously considers a number of competing hypotheses. Each candidate hypothesis is exposed to input which it attempts to parse. At the end of a round of parsing, the learner assesses how well each candidate did. The candidates are ranked according to their relative fitness. The fittest go on to reproduce candidates in the next generation, the least fit die out. Through successive iterations of this procedure, the candidate set presumably becomes increasingly fit, and converges toward the correct grammar.

There are three main problems with this model:

1. It requires an accurate fitness measure, but none has been proposed.
2. Any such measure requires that the learning space be *smooth*, i.e., that closeness in surface resemblance reflects closeness in parameter space; but this assumption is incorrect.
3. As with the Triggering Learning Algorithm, the developmental sequence is dictated by the input forms encountered.

To get some idea of the problems facing this model, consider *Table 2*.

Table 2. Effects of parameter settings: Selkup (number and percent correct)

	Parameters of 10	in %	Words of 8	in %	Syllables of 20	in %	Main stress of 8	in %
a.	4	40	2	25	7	35	3	37.5
b.	6	60	1	12.5	7	35	5	62.5
c.	7	70	4	50	12	60	4	50
d.	8	80	5	62.5	14	70	5	62.5
e.	9	90	5	62.5	14	70	5	62.5
f.	9	90	3	37.5	10	50	3	37.5

I simulated a genetic algorithm attempting to learn ten metrical parameters for Selkup (Halle and Clements 1983: 189). The chart shows a sample of the results. In the first column are listed the number of parameters correct. This gives an indication of how close the grammar is to the target. In the other columns I listed a number of possible

surface indicators that one might try to use to show goodness-of-fit: number of words correctly stressed (there were eight words in the test), number of syllables correctly stressed, and number of main stresses on the correct syllable. Though the grammars get progressively closer to the target as we proceed down the columns, no single surface measure reflects a monotonic improvement. This is illustrated most dramatically by grammar (f), which has only one parameter wrong and yet does worse than grammar (c), with three parameters wrong, and not much better than grammars (a) or (b). A more systematic simulation of 64 grammars using six metrical parameters confirms the general unreliability of surface measures as keys to the goodness of the grammar. In other words, closeness in extensional space (i.e., the surface data) is unreliably correlated with closeness in intensional space (i.e., the grammar).

Conclusion

Steven Crain (1991) has compared language acquisition to a scavenger hunt: learners are given a list of things to get — a green sock, an old muffler, a banjo — and they run around looking for these things, and collect them as they find them. But if the scenario I have sketched is correct, language acquisition is not a scavenger hunt but a treasure hunt. In a treasure hunt, you have to find a sequence of numbered clues in order, where one clue leads you to the next. Clue #15, for example, might tell you to look for #16 using information you collected at #11 and #14. If you accidentally stumble on this clue before you have reached the earlier ones, it may well be meaningless or misleading to you. Like Meno's slave boy, we gradually construct the solution to the puzzle that is our native language by proceeding systematically, answering a series of questions put by our own inner Socrates.

Selected References

Cheney, Dorothy L. and Robert M. Seyfarth
    1990   *How Monkeys See the World: Inside the Mind of Another Species*.
                Chicago, IL: University of Chicago Press.

Chomsky, Noam
    1975   *Reflections on Language*. New York: Pantheon Books.

Chomsky, Noam
    1981   Principles and parameters in syntactic theory. In: Norbert Hornstein
                and David Lightfoot (eds.), *Explanation In Linguistics: The Logical
                Problem of Language Acquisition*, 32-75. London: Longman.

Chomsky, Noam
    1988   *Language and Problems of Knowledge: The Managua Lectures.*
                Cambridge, MA: MIT Press.

Clark, Robin and Ian Roberts
    1993   A computational model of language learnability and language
                change. *Linguistic Inquiry* 24: 299-345.

Crain, Stephen
    1991   Language acquisition in the absence of experience (with commen-
                taries). *Behavioral and Brain Sciences* 14: 597-650.

Crain, Stephen and Mineharu Nakayama
    1987   Structure dependence in grammar formation. *Language* 63: 522-543.

Dresher, B. Elan
    1999   Charting the learning path: Cues to parameter setting. *Linguistic
                Inquiry* 30: 27-67.

Dresher, B. Elan and Jonathan Kaye
    1990   A computational learning model for metrical phonology. *Cognition*
                34: 137-195.

Gibson, Edward and Kenneth Wexler
    1994   Triggers. *Linguistic Inquiry* 25: 407-454.

Halle, Morris, and G. N. Clements
    1983   *Problem Book in Phonology.* Cambridge, MA: MIT Press.

Kaye, Jonathan
    1989   *Phonology: A Cognitive View.* Hillsdale, NJ: Erlbaum.

Lightfoot, David
    1989   The child's trigger experience: Degree-0 learnability (with commen-
                taries). *Behavioral and Brain Sciences* 12: 321-375.

Piattelli-Palmarini, Massimo (ed.)
    1980   *Language and Learning: The Debate between Jean Piaget and Noam
                Chomsky*. Cambridge, MA: Harvard University Press.

Plato
   Meno. Various editions.
Quine, Willard Van Orman
 1960 *Word and Object.* Cambridge, MA: MIT Press.
Russell, Bertrand
 1948 *Human Knowledge: Its Scope and Limits.* New York: Simon & Schuster.
Stowell, Tim
 1979 Stress systems of the world, unite! *MIT Working Papers in Linguistics* 1: 51-76.