Waiting for Universal Grammar

Geoffrey K. Pullum

School of Philosophy, Psychology and Language Sciences
University of Edinburgh

Pietroski and Hornstein (this volume; henceforth P&H) see linguists as explorers of the component of the human mind that is responsible for the unfailing success of normal human babies in in achieving first language acquisition. They favor a view that is often known as **linguistic nativism**, fashionable among linguists who closely follow the work of Chomsky. Its thesis is that certain innate linguistic prerequisites, possessed by all humans at birth, render language acquisition feasible. P&H group these innate mental characteristics together under the heading of Universal Grammar (UG).[1] But the properties of UG tend to be more boasted of than empirically validated, and P&H supply no new details. This chapter warns readers to heed the warning of Scholz and Pullum (2006) about "irrational nativist exuberance," and draws attention to interesting emergent lines of recent work that P&H do not mention.

Logical possibilities

P&H follow Chomsky (1965) in regarding a human infant as essentially analogous to a device that, on being exposed to an indefinitely long but finite stream of utterances from some human language, constructs an internal representation of a generative grammar for that language. A generative grammar is a finite system of sentence-building procedures capable of building exactly the sentences of the input language and no others. P&H contend that the task of constructing a generative grammar from the input a child gets would be impossible for an unaided intelligence, but being in possession of the information formalized in the theory of UG makes the task feasible or even straightforward.

I want to concede at the outset that it is certainly possible to imagine a way of responding to a finite input corpus of unprocessed utterances from some language by automatically outputting a correct generative grammar for that language. Imagine a device that internally stores representations of generative grammars for English (G_E), Hawaiian (G_H), and Turkish (G_T), and operates by scanning the acoustic form of input utterances. If the acoustic signature of utterance-final consonants is never encountered, then after a reasonable delay for confirmation it outputs G_H. However, if clear evidence of utterance-final consonants is encountered, G_H is ruled out (since in Hawaiian every syllable ends in a vowel), and thereafter if the characteristic signature

[1] P&H introduce the term 'Universal Grammar' in the first section of their chapter, apologising for its ambiguity, but then talk about a 'Faculty of Language' (FL) and a Language Acquisition Device (LAD), returning to introduce the abbreviation 'UG' only near the end. If I understand their intent correctly, UG is the theory of what is in the FL and thus constrains the LAD.

of close front rounded vowels and close back unrounded vowels are observed with reasonable frequency, it outputs G_T (since Turkish does feature those vowel types). Otherwise, if both of these vowels are lacking, after some reasonable time it outputs G_E.

The device unfailingly produces a correct grammar for the right language, after some exposure to utterances. Yet nothing about grammatical properties has to be learned: grammars are selected automatically on detection of certain physical properties of acoustic stimuli, and nothing about grammar need be observed at all. (The process could of course be below the level of consciousness: the language acquirer would not need to be aware of anything about its operation.)

It should not be thought that I am inventing a straw man here: the idea that only a finite number of languages need to be considered is not mine. Chomsky (1981:10–11) proposed very seriously that the learnability of human languages could be guaranteed if only finitely many grammars were allowed by UG, and the idea is referred to as "attractive" by Hornstein (2009: 167).[2]

It might also seem strange to depict the infant as never really learning from properties of utterances, but simply jumping involuntarily to certain conclusions under the influence of trigger stimuli. But this too is explicit in defenses of the sort of UG that P&H espouse (Lightfoot 1989; Gibson and Wexler 1994; Fodor 1993). Language acquisition is claimed to involve internally scheduled leaps of biological growth, which the environment merely triggers in some cases. Note the remarks of Chomsky (1980: 134–136):

> I would like to suggest that in certain fundamental respects we do not really learn language; rather, grammar grows in the mind... There are certain processes that one thinks of in connection with learning: association, induction, conditioning, hypothesis-formation and confirmation, abstraction and generalization, and so on. It is not clear that these processes play a significant role in the acquisition of language. Therefore, if language is characterized in terms of its distinctive processes, it may well be that language is not learned. … It is open to question whether there is much in the natural world that falls under 'learning'.

Logically, it is conceivable that the mental development of both humans and other animals is almost entirely a matter of biologically built-in scheduling, prompted only in some minor respects by sensory experiences, so that "learning" is a folk term with very little applicability. But scientists who believe this need to tell us something about the actual neural architecture of

---

[2] See Pullum (1983) for a detailed argument against trying to define UG in a way that limits the class of grammars (hence languages) to a finite set. The suggestions for how a finite bound might be achieved would certainly allow for an astronomically huge number of distinct grammars, too large for finiteness to be of any use. The finitely-many-languages idea is seldom mentioned in the contemporary literature.

the internally-driven growth capacity, and the ways in which experience triggers it. In P&H's chapter we search for that in vain.

What must be learned

One consideration militating against the empirical plausibility of P&H's view is our planet's linguistic diversity, about which they say nothing at all. Human languages turn out to be so diverse in grammatical terms that a tight set of true universal principles governing them all can hardly be imagined. Some have word formation and inflection processes of extreme complexity: whole sentences can often be expressed as single words in Eskimoan languages. Others (like Vietnamese) have virtually no word-building. Some (like English) maintain fairly strict constituent order, while others (like Sanskrit, and many aboriginal languages of Australia) have remarkably free word order.

Languages differ, for instance, in the order of Subject (S), Verb (V), and Object (O), in every way they logically could. There are only seven logical possibilities for the normal order for simple, stylistically neutral, declarative clauses, we find all seven favored in at least some languages: SVO (English, Swahili); SOV (Turkish, Japanese); VSO (Hawaiian, Irish); VOS (Malagasy, Tzotzil); OVS (Hixkaryana, Urarina); OSV (Apurinã, Nadëb), and no strong preference (Sanskrit, Walbiri). Many other syntactic facts also have to be learned without any discernible possibility of significant help from UG: whether there are prepositionsor postpositions (English *in India*, Hindi *Bharat mẽẽ*); modifying adjectives before the noun or after (English *white wine*, French *vin blanc*); determiners before the noun or after (English *the house*, Danish *hus-et*); and so on.

Such differences cannot be brushed aside as minor divergences from a single human language template.[3] Children clearly have to figure out many parochial syntactic facts on the basis of linguistic experience. Since they manage to do it with virtually 100% success, they could surely learn a large array of other facts about normal syntax at the same time, by the same methods of observation, comparison, and familiarization.

Numerous other aspects of a language must clearly be learned from the evidence of experience, since they are so obviously parochial and idiosyncratic. Most obviously, the properties of individual words have to be learned simply by listening to people use them and seeing what happens in the interaction. The learner has to become acquainted with tens of thousands of words, each having properties of many kinds:

- phonology: the plural suffix on *cats* is an entirely different sound from the one on *dogs*; in *insect* the most heavily stressed syllable is the first, but in *infect* it's the second;

[3] Chomsky remarks that "even down to fine detail, languages are cast to the same mold" and an unbiased scientist from Mars "might reasonably conclude that there is a single human language, with differences only at the margins" (2000:7). Current knowledge about the remarkable typological diversity of human languages makes that look extremely implausible to me, for a Martian investigator even minimally attentive to word and sentence structure.

- inflection: *write* has the past participle *written* (not **writed*); *we* has the accusative form *us* and the genitive form *our*;
- derivational relationships: *ignorance* denotes the property of being ignorant, but *instance* doesn't denote the property of being instant; *terrified* and *terror* are related in meaning but *rectified* and *rector* are not;
- syntactic properties: *eat* can have a direct object (*Let's eat it*/*Let's eat*), *devour* must have one (*Let's devour it*/**Let's devour*), *dine* mustn't have one (**Let's dine it*/*Let's dine*); *likely* takes infinitival complements (*He's likely to be late*) but *probable* does not (**He's probable to be late*); *damn* occurs as a modifier before a noun in a noun phrase;
- literal meaning: *likely* is synonymous with *probable*; *eager* denotes a property that only a mind-possessing entity can exhibit; *damn* adds no truth-conditional meaning;
- conventional implicatures: *lurking outside* hints at furtiveness or ulterior motive, while *waiting outside* does not; *damn* signals irritation on the utterer's part;
- overtones and associations: *ain't* is markedly nonstandard and colloquial; *fuck* is coarse and offensive; *whilst* is old-fashioned; *whom* is distinctly formal; and so on.

UG cannot help in any substantive way with any of this. There is almost nothing universal about the properties of words: some of their properties differ dialectally, and even idiolectally (from one speaker to another). For further evidence of the plethora of aspects of human language that cannot plausibly be universalized, see Evans and Levinson (2009)[4] and, with respect to syntax, Culicover (1999, esp. ch. 3).

The Fawlty strategy

There is a vast literature on approaches to language acquisition with goals other than P&H's (Dąbrowska 2015 offers a very useful survey). But the attitude that P&H seem to maintain toward such alternative literature, and toward research programs that disagree with linguistic nativism, could be called Fawltyism, after the the belief of the fictional bigoted British hotelier Basil Fawlty[5] about the key to getting along with Germans: "Don't mention the war!"

One remarkable failure of mention relates to the details of infants' actual linguistic input. P&H point out that we are not interested in "how a suitably clever child *could* acquire an English grammar ... given an *ideal* sample of English discourse," but rather, "how a typical child *does* acquire an English grammar given a *typical* sample of English discourse—or more precisely, the temporally unfolding subset of any such sample that corresponds to what a typical child might

[4] Evans and Levinson's title ('The myth of language universals') is ill-chosen: their central point is that the sheer diversity of human languages may be more interesting for cognitive scientists than whatever properties languages turn out to share.

[5] In the 1975 BBC TV situation comedy 'Fawlty Towers,' series 1, episode 6: 'The Germans'.

attend to and represent for purposes of formulating and evaluating grammars." So what sort of sample do children typically get? P&H show absolutely no interest in this question.

Estimates of how much language children hear in their early years run to a million utterances a year or more (Pullum and Scholz 2002:44–45, citing Hart and Risley 1995). And language acquisition begins before a child's first birthday and continues into adolescence (Dąbrowska 2015:3 and references there). Since some arguments for UG turn crucially on claims about what children almost never encounter (Pullum and Scholz 2002:19–23), the quantity and content of this input is highly relevant. Yet P&H ignore it, focusing instead on a small range of invented English utterances illustrating selected syntactic and semantic points.

They argue, for example, that generalizing from *Grover washed himself* (where *himself* must be Grover) and *Grover washed him* (where *him* cannot be Grover) would encourage the child to adopt a false generalization, namely that *himself* must have an antecedent close to it, while *him* is not allowed to. They then show falsity of the generalization with these examples:

Kermit expected to feed Grover and wash himself.	(*himself* = Kermit, not Grover)
Kermit expected to feed Grover and wash him.	(*him* = Grover, not Kermit)

This would indeed be puzzling if you adopted the practice of looking at nothing more than positions of pronouns and permitted antecedents in word sequences: *himself* now has its antecedent further away, and *him* refers to a nearby noun phrase — the opposite of what we find in shorter sentences. Yet P&H cite Crain (2012) as having shown experimentally that children as young as 3 can understand such sentences correctly.

But consider how different things are once we assume that children learn not just from overheard word sequences but from the meanings that the context suggests they must have. Reflexive acts like washing oneself are conceptually very different from transitive acts of washing somebody else. Suppose we adopt the neologism *autoablution* for self-washing, to highlight the difference. The first sentence can be paraphrased as "Kermit expected to engage in feeding Grover and autoablution." This does not in any way tempt us to think that washing Grover might be implied: Kermit is the agent, and both the expected activities are his.

P&H's peculiar view is shaped by (i) a refusal to take meaning into account, and (ii) the idiosyncratic fact of English syntax that it expresses reflexive actions like autoablution by using the verb *wash* with a special object pronoun ending in *-self*. Some languages have invariant non-pronouns serving as indicators of reflexivity; many use a form that is also assigned duties like signaling impersonal or a reciprocal meaning (e.g. Spanish *se*); some use affixes on the verb to mark reflexive meanings; some simply mark the verb as detransitivized; some have different constituent orders so that the reflexive element regularly precedes its antecedent; and so on.

What might be universal here is not anything about the syntactic positions of pronouns, but rather the cognitive distinction between reflexive actions and transitive actions — as in the distinction between autoablution and washing an entity distinct from one's own body. Once meanings are considered, not just the linear sequence of words with coreference indices, the different interpretations of *wash himself* and *wash him* become much less puzzling.

Of course, the developmental psychological question of how babies reach the stage of distinguishing utterances about autoablution from those about washing someone still deserves

study; babies are known to develop an awareness of agency and transitivity — their own ability to act in a way that affects an object external to themselves — extremely early in life. But it is a characteristic practice of UG defenders to ignore such genuinely psychological facts. P&H prefer to focus attention on bare word sequences (perhaps annotated with coreference indicators), saying little or nothing about interpretation in context, and simply assert that it is unclear how their contrasting properties could be learned.

P&H allude to "a common mechanism for segregating the permissible from illicit interpretations," about which they can say nothing explicit except that "Whatever this mechanism is, it seems to be operative in young children, at least many of whom have presumably not inspected the relevant linguistic data." It is tautologous to say that there is some kind of mechanism; but the claim about children not inspecting relevant data is impossible to evaluate without intensive study of their actual experience of noticing contextually situated utterances and appreciating the meanings that are being expressed. Most developmental psycholinguists are interested in that kind of study. P&H do not even mention it as a possibility.

Misframed generalizations

The foregoing example shows us that a puzzle about how some generalization could be learned can depend on the way the generalization is framed. There are indefinitely many ways to describe any given set of facts; we should beware of arguments for nativism that depend on perversely fashioned descriptions. They are common in the nativist literature.

Anderson and Lightfoot (2002:18–21) provide a particularly clear case, based on facts first noted by King (1970). Several English auxiliary verbs have reduced or 'clitic' forms (phonologically, single consonants; standard spelling writes *'d* for *had* or *would*, *'ll* for *will*, *'m* for *am*, *'re* for *are*, *'s* for *is* or *has*, *'ve* for *have*). These cannot be used everywhere the full forms occur. Anderson and Lightfoot note that maybe you could learn "*is* may be pronounced [z]" just from noting that people sometimes say *Jay is here* and sometimes *Jay's here*, but you would not be able to learn from experience that *I wonder where Jay is* lacks the clitic form ***I wonder where Jay's*. In their view:

> The problem is that, because we were not informed about what **cannot** occur, our childhood experience provided no evidence for the "except" clause(s), the cases in which the reduced form is impossible... That is, we had evidence for generalizations like ''*is* may be pronounced [z]" ... but no evidence for where these generalizations break down.

Their emphasis on our not being "informed about what **cannot** occur" suggests it is impossible to learn that something never occurs without being explicitly told that. I will return to this patently false epistemological claim later. But they also rely on casting the description in a way that makes learnability maximally puzzling, yielding maximal support for a putative innate device of unknown structure and function. Interestingly, though, in this case the description assumed is actually mistaken.

Anderson and Lightfoot attribute the failure of reduced forms to appear in some contexts to the fact that the following piece of the sentence is empty or missing. In *I wonder where Jay is*,

the location-denoting phrase that would normally follow *is*, the word *where*, is missing because (as English syntax requires) it is an interrogative word and has to appear at the beginning of its clause. In *Bob isn't interested but Jay is*, a phrase with the meaning "interested" is missing from the second clause. They claim (p. 29) that "where a phrase consists of an auxiliary verb together with a following element such as an adjective, and that following element is unpronounced, the use of a full (rather than clitic) form of the auxiliary is necessary to provide the phrase with enough phonetic substance" to stand alone.

The true generalization, however, has nothing to do with unpronounced following phrases. Zwicky and Pullum (1997) pointed out that the supposed generalization the blocking of reduced forms is found in certain cases where the complement of the auxiliary is present and pronounced. This is illustrated in the utterances by B below, which are examples of what Zwicky and Pullum call Rejoinder Emphasis:

A: *You won't go through with that.*
B: *I will tóo go through with it!* B: **I'll tóo go through with it!*

A: *She isn't really gonna call the cops.*
B: *She is tóo gonna call the cops.* B: **She's tóo gonna call the cops.*

Here the material following *will* and *is* does get pronounced, but still *I'll* and *she's* are not permitted. An alternative analysis is needed to cover these facts. Providing it turns out to make the learnability problem much easier.

The clitic forms of auxiliary verbs, being single consonants, cannot bear even weak stress. But certain constructions require weak stress on an auxiliary verb; for example, Post-Auxiliary Ellipsis, as in *I will* or *She is*, allows a clause to consist of a subject and a weakly stressed auxiliary (and to be understood with an implicit verb phrase meaning following the auxiliary). Rejoinder Emphasis also has such a requirement: it involves a clause consisting of a subject, a weakly stressed auxiliary, a heavily stressed *too* (or sometimes *so*), and (optionally) the appropriate complement for the auxiliary. Because of the weak stress requirement, neither construction can permit clitic auxiliaries.

The learnability problem now melts away: learners hear *I wíll!*, and *She is tóo gonna*, etc., and use them roughly as they have heard others use them, complete with the weak stress. They never hear sentences like **I'll!* or **She's tóo gonna*, and unsurprisingly they never produce them.

Attempting to draw a moral about UG from facts that have in fact been misdescribed is a common mistake. Take P&H's discussion of *with*-phrases. They assert that a *with*-phrase at the end of a transitive clause may be interpreted as modifying the object (as in *Karen saw her husband with another woman*) or as an instrumental modifier of the verb (as in *Seymour sliced the salami with a knife*). For some sentences either is possible: *The woman saw the man with the telescope* is ambiguous between "saw the man in possession of the telescope" and "saw the man by using the telescope." However, P&H claim, a third meaning, "The woman saw the man and was in possession of the telescope," where the *with*-phrase modifies the subject, is impossible.

They do not state the generalization they are assuming, but it amounts to a claim that a *with*-phrase at the end of a transitive clause cannot be interpreted as applying to the subject. This

would entail that *Karen saw her husband with another woman* cannot be made true by a situation where Karen spots her own husband while she is in the company of her friend Maureen. This might seem plausible enough. But the generalization P&H tacitly advocate cannot possibly be imposed by UG; even a tiny extension of the range of data refutes it. Consider just one new case: *Karen watched the eclipse with another woman.* There we do understand the *with*-phrase as modifying the subject of the clause (*Karen*), and not the direct object (*the eclipse*). That is exactly the modification relation that P&H implicitly suggest UG forbids.

What's actually going on has to do with something P&H say nothing about: contextual plausibility of entailments has a huge influence on what interpretations seem acceptable. Being in the company of a woman is plausible and natural; being in the company of an eclipse is bizarre or impossible. Charity demands that we select the sensible interpretation. (See Gualmini and Schwarz 2009 for other cases where spurious learnability puzzles are dissolved once pragmatics and charitable interpretation come into the picture.)

Misdirection and legerdemain

P&H repeatedly point to particular semantic facts about miscellaneous snippets of invented sentences in contemporary standard English, and suggest that it is hard to imagine how infants could learn such facts from experience. They offer not a hint of any general theory of how specific innate cognitive states, mechanisms, or whatever would actually help. Instead they conceal the fact that their explanatory hat contains no rabbit by using smoke, mirrors, and misdirection.

It is certainly true that infants, when exposed to the experience of seeing and hearing utterances in context every day for a year or two, reliably form some generalizations and eschew others. The same is true even when we are older and learn certain less common construction types from reading. When we encounter sentences like (i) and (ii), we form generalizations implying that a sentence like (iii) might also be grammatical:

i. *The more I see of him the less I like him*.
ii. *The more they eat the fatter they get*.
iii. *The more it rains the deeper the puddles get.*

But we don't assume that every word in such sentences must contain the letter *e*, or that every such sentence must have an even number of words — both generalizations that happen to hold of (i) and (ii) but not of (iii). The question is not whether indefinitely many useless generalizations go unnoticed when we are forming our general impressions about what we can count on in future experience, but **what are the specific principles or practices that permit us to learn as fast as we do**. That is what we would expect P&H to tell us something about. But they don't.

Indeed, they dismiss or ignore aspects of the acquisition situation that clearly could be of explanatory value. Take the idea that children might tend to produce utterances by combining subparts of utterances they have heard. Far from being carefully rebutted, this is never even mentioned. P&H make the correct observation that *we may have been* is grammatical but versus **we have may been* is not, as if there were some mystery that babies learn such things. But take any collection of English text you like, and count the occurrences of *may have*, *have been*, *have*

may, and *may been*. The first two are extremely common, and the other two basically never occur at all.[6] Thus even if you just made up sentences by randomly stringing together adjacent word pairs that you have definitely heard (which of course is absurdly inadequate as a method for achieving grammatical correctness), you would never construct **we have may been*.

I'm not suggesting that children ever go through a stage of randomly stringing together adjacent word pairs (though in phonology such random trying-out of possibilities does seem to occur; developmental psycholinguists call it the 'babbling' phase). I'm simply pointing to the rather extraordinary fact that P&H misdirect our attention away from the possibility that what children learn might be influenced by the frequency of word sequences found in what they hear.

P&H also neglect the potential relevance of lexical meaning. They say in footnote 3 that "it's hard to see how an English procedure could pair the pronunciation of 'revolutionary new ideas occur infrequently' with a meaning, but not do the same for 'colorless green ideas sleep furiously'." This does not mention the content of word senses at all. Notice that being revolutionary is compatible with being novel, whereas being colorless is incompatible with being colored green. P&H treat such facts as if they could not possibly be relevant to assigning meanings to word strings. It seems to me more natural to assume the opposite.

Likewise, when contrasting the ambiguity of *The goat is ready to eat* with the lack of ambiguity in *The goat is eager to eat*, P&H do not mention the fact that readiness is a property of either thinking beings or inanimate objects, while eagerness can only be exhibited by a mind-possessing entity. So your dry-cleaning can be ready for collection but it can't be eager for collection. Thus the relation of readiness can hold between an entity and an eating situation in two ways, but the same is not true for the relation of eagerness.

That doesn't settle the business of how the syntax and semantics of the *ready to eat* construction might be learned (for a terse overview of the rather complex facts, see Huddleston and Pullum 2002:1256–1259). I merely note that P&H never mention that words like *ready* and *eager* (or for that matter *goat* and *eat*) have meanings. Yet no child learns the syntax of infinitival complements of adjectives without paying attention to the meanings of those adjectives. It is reasonable to think that the latter might have some bearing on the former.

Grueness and induction

Section 4 of P&H's chapter digresses into the philosophy of induction, and the anomalous predicate *grue*, deliberately gerrymandered by Goodman (1954) to be induction-defeating (grueness coincides extensionally with greenness up to a certain arbitrary time point and with blueness thereafter). But all that emerges from this digression is the suggestion that "confirming

[6] I did a quick check on 44 million words of newspaper text and about a million words of classic novels. The frequency of *have may* sequences is literally zero. The sequence can occur accidentally in sentences of more than one clause, of course (*What they have may not suit you*); but such sentences are in practice amazingly rare.

generalizations by empirical induction requires a certain kind of vocabulary; and acquiring grammars, in response to ordinary experience, evidently requires a different kind of vocabulary."

The reader might expect that some hint of this special vocabulary for grammar induction would follow, but it does not. The hint that induction of grammatical generalizations might succeed because of a constrained vocabulary of induction-friendly predicates is dropped in favor of a different metaphor, one suggestive of either gambling or religion: "children make leaps of faith when they settle on grammars that constrain homophony in the specific ways noted." This too is promptly dropped, in favor of a new analogy, about the difference between guessing whether there is an 'O' in a bag of Scrabble tiles and knowing in advance that there is. This leads to a remark about children's grammar acquisition "revealing the character of the Universal Grammar (UG) that both provides and constrains their options."

The metaphors and analogies shift, but it is all metaphors and analogies. P&H give us no glimpse of what UG might actually say. Generative linguists have been talking for half a century (since Chomsky 1965) about the development of a precise and fully articulated theory of UG, tacit knowledge of which somehow converts an impossible learning task into a feasible triggered acquisition task. P&H do not offer any convincing illustrative example of how this might work.

Computational psycholinguists, meanwhile, are conducting interesting experiments on what is called "unsupervised learning," relying on very few prior assumptions about the structure of languages. There are built-in assumptions, of course — every learning program must incorporate some assumptions about what it has to do — but not up-front assumptions about the structure of sentences or the syntactic and semantic principles shared by all human languages and by no nonhuman languages, which is what the UG program was always supposed to be about.

Unsupervised learning of human languages proves very hard. Progress has been slow. But the crucial thing is that the rhetorical project of defending UG is not coming up with any components, principles, or insights that help.

Ideas reminiscent of 1940s empiricism turn out to be more useful than anything emerging from UG. Consider an example of a problem so elementary as to be almost trivial by comparison with figuring out the intricacies of syntax and semantics: identifying the boundaries of words in an utterance assumed to be already analysed into discrete speech sounds (a generous assumption, since the infant's actual input is a continuous stream of unlabeled noise). The structuralist linguist Zellig Harris (Noam Chomsky's undergraduate and graduate mentor) suggested in 1955 that there was a way to take a string of symbols and find out automatically where the boundaries between meaningful elements are, without appealing to anything about meaning or grammar.

The idea is simple enough. Make a line graph with the symbols of the utterance in temporal sequence along the horizontal axis (without spaces, of course) and positive integers up the vertical axis. Above each transition between segments plot (using statistical data gleaned from observed sequences in other texts) a point indicating how many segments have a fairly high probability of coming next. The graph will have troughs just before any highly predictable symbol (for example, in written English the letter *q* is almost always followed by *u*) and peaks where there are many possibilities for the next symbol (after *ck*, any letter can occur). Harris pointed out that segmenting the utterance at the peaks of the graph gives a remarkably reliable guide to where the word boundaries are (or even the boundaries of the meaningful elements called morphemes, if you target somewhat lower peaks).

In 1955, Harris's idea was just an intuition supported perhaps by a little paper-and-pencil experiments, but computers are now fast and powerful enough to confirm over large quantities of data that the method works quite well. Harris's intuition has been confirmed several times (see e.g. Tanaka-Ishii and Jin 2008, Griffiths et al. 2015). A stream of symbols can be broken into plausible candidates for words or even morphemes on the basis of nothing more than probabilities of specific units at given points in naturally occurring sequences.

Whether or not infants identify the words in their language by computing the positional freedom of occurrence for the phonological units in the sequences they hear is not the issue here (though it should be noted that neonates are in fact sensitive to statistical patterns of frequency in meaningless sound sequences: Saffran et al. 1996). My point is merely that one would expect a progressive UG program to at least have a specific proposal for specific built-in assumptions about identifying of wordlike units in utterances, preferably embodying a way in which a learner could identify the meaningful units much faster and with much less work. But that is not happening. UG enthusiasts have proposed no explicit universal word-identification modules to be tested by computational psycholinguists, whereas implementations of Zellig Harris's method have been shown to work.

On the learning of vastly subtler and more complex things in domains like syntax and semantics we are largely ignorant. The mysterious UG module that P&H want us to believe in has not yet contributed anything toward lessening our ignorance.

Abduction and curiosity

P&H's revisiting of Goodman's "new riddle of induction" yields hardly any payoff. They might have done better to look at some of the modern work on Bayesian abduction — though their Fawltyism biases them against this, since the research program in question is (notoriously) a rival. Bayesian statistical learning is certainly not (in my view) the answer to everything, but I do think that anyone interested in the possibly innate prerequisites of human cognition and language acquisition should be drawing on insights from that literature.

The Bayes-Price rule in probability theory (known as Bayes' Theorem after the Rev. Thomas Bayes, 1701–1761), can be interpreted as licensing abductive arguments by relating statistical properties of evidence to plausibility of current beliefs (see Chapters 8 and 9 of Nola and Sankey 2007 for an accessible philosophical introduction). How likely is it that a hypothesized grammatical generalization *G* is accurate, given some body of evidence about sentences *E*? Bayes' Theorem entails that the probability of *G* being correct, given *E*, is proportional to the probability that the evidence would look like *E* if *G* were correct. We don't need to get more technical than this to appreciate the crucially important fact about it: it shows that there is a way in which absence of evidence can yield evidence of absence.

It would be fallacious to reason that because you have never observed any instances of some linguistic expression type, it is therefore grammatically impermissible. But it is not fallacious to reason along these lines regarding a type of expression:

i. if expressions like this were allowed, many instances would have turned up by now;

ii. instead there have been none (or negligibly few); therefore (by Bayes' Theorem),

iii. the estimated probability that some grammatical generalization excludes expressions like this should be raised.

Such probabilistic reasoning works even if there are errors in the data. If you have heard thousands of occurrences of *afraid of it* and only very occasional instances of *afraid from it*, that doesn't just confirm the idea that *afraid of it* is allowed by the grammar; it also provides grounds for upping the probability that **afraid from it* is ungrammatical, only occurring sporadically in imperfect English written by foreign speakers.[7] Your statistical experience of not hearing word sequences of certain types provides rational support for the hypothesis that they are grammatically forbidden.

One interesting recent paper, Friston et al. (2017), uses computer-simulation experiments on rule discovery by Bayesian active inference, and relates it to broader aspects of human cognition, notably insight and curiosity. It contains some particularly interesting remarks about curiosity and the drive to acquire knowledge and find order. Knowledge resolves or reduces uncertainty about the world, which is an imperative for any intelligent entity: "any sentient creature must minimize the entropy of its sensory exchanges with the world," where entropy is "uncertainty or expected surprise, where surprise can be expressed as a free energy function of sensations and (Bayesian) beliefs about their causes." And "resolving different sorts of uncertainty furnishes principled explanations for different sorts of behavior."

Through Bayesian model reduction Friston and colleagues then show, in effect, how learning is enormously facilitated by the assumption that **there are rules or symmetries** — the assumption that the universe is not random. This might be seen as echoing the final remarks of Herbert Feigl in his classic 1934 paper on the logic of the principle of induction:

> The attempt to know, to grasp an order, to adjust ourselves to the world in which we are embedded, is just as genuine as, indeed, is identical with, the attempt to live. Confronted with a totally different universe, we would nonetheless try again and again to generalize from the known to the unknown. Only if extended and strenuous efforts led invariably to complete failure, would we abandon the hope of finding order. And even that would be an induction.

Friston et al. and Feigl both seem to see human beings, even more than any other animals, as constantly (albeit subconsciously) striving to reduce uncertainty not just by gathering evidence about the world of their experience but by actively developing a formulation of it that permits explanation of what occurs and why. This applies to linguistic experience just as much as to other kinds of experience: we search for an explanatory understanding of the verbal behavior that we observe, one that permits us to see why other people's utterances have the form they do — and don't have the form they don't.

---

[7] Google reports some 50 million web hits for *afraid of it*, but only 35,000 for *afraid from it* (and for the web, that is close to zero). Most of the latter are mistakes; one token reads: "All of us thinking about death maybe because we afraid from it or because we thinking that our death more good than our current life."

P&H's claim is that data-driven hypothesizing could not possibly work for language: some hypotheses must be ruled out a priori for humans, not on grounds not fitting the facts but because UG simply won't permit them. An imperative on the part of all sentient beings to seek regularities in experience would not be enough, in P&H's view, to account for what happens: human languages must be designed in a certain special way that would be undiscoverable without a special component of mind supplying crucial clues about what generalizations to adopt.

That view overlooks another whole research program that P&H's Fawltyism precludes them from citing. The work of Simon Kirby and his collaborators over the past two decades has suggested that we should "concentrate less on the way in which we as a species have adapted to the task of using language and more on the ways in which languages adapt to being better passed on by us" (Kirby 2001:110). In a long series of ingenious experiments on human subjects' ability to learn invented languages (see Kirby 2017 for a recent overview), Kirby and colleagues have developed a view very different from the UG one. For example, when human subjects attempt to learn a mini-language with random word/meaning associations, and the errors made by each 'generation' of learners are built into the version of the language handed to the next batch of subjects, regularities in the relation between word structure and meaning components begin to emerge over the 'generations'. The UG mindset sees our progress as assisted by our being restricted to a class of languages that UG deems fit and proper for us. The findings of Kirby and colleagues suggest rather than our cognitive shortcomings (like our limited memory for random irregularity) slowly sculpt languages to make them more learnable by creatures like us.

Of course it is still logically possible that some of what we have yet to discover about language acquisition will turn out to involve inbuilt features of human brains specifically tied to language and nothing else. However, imagining that such features have already been discovered, and that linguistic research has mapped their structure and set out a theory of UG that reveals how they function to aid the acquisition process, would be a major misunderstanding of the state of the art in the cognitive and linguistic sciences.

Acknowledgment

I am grateful to Bob Borsley, Jenny Culbertson, Jim Donaldson, Karl Friston, and Mark Steedman for very useful critical comments. They agree neither with each other nor with me, so what I say here must be regarded as purely my responsibility, not theirs.

References

Anderson, Steven and David W. Lightfoot (2002) *The Language Organ: Linguistics as Cognitive Physiology*. Cambridge: Cambridge University Press.

Chomsky, Noam (1965) *Aspects of the Theory of Syntax*. Cambridge, MA: MIT Press.

Chomsky, Noam (1980) *Rules and Representations*. New York: Columbia University Press.

Chomsky, Noam (1981) *Lectures on Government and Binding*. Dordrecht: Foris.

Chomsky, Noam (2000) *New Horizons in the Study of Language and Mind*. Cambridge: Cambridge University Press.

Culicover, Peter W. (1999) *Syntactic Nuts: Hard Cases, Syntactic Theory, and Language Acquisition*. Oxford University Press.

Dąbrowska, Ewa (2015) What exactly is Universal Grammar, and has anyone seen it? *Frontiers in Psychology* **6**, article 852.

Evans, Nicholas and Stephen Levinson (2009) The myth of language universals. *Behavioral and Brain Sciences* **32**, 429–492. DOI: 10.107/S0140525X0999094X

Feigl, Herbert (1934) The logical character of the principle of induction. *Philosophy of Science* 1, 20–29. DOI: 10.1086/286303 Reprinted in Herbert Feigl and Wilfrid Sellars (Eds.), *Readings in Philosophical Analysis*, 297–304 (Appleton-Century-Crofts, New York, 1949).

Fodor, Janet Dean (1998) Unambiguous triggers. *Linguistic Inquiry* **29**, 1–36.

Friston, Karl J., Marco Lin, Christopher D. Frith, Giovanni Pezzulo, J. Allan Hobson, and Sasha Ondobaka (2017) Active inference, curiosity, and insight. *Neural Computation* **29**, 1–29.

Gibson, E., and K. Wexler (1994) Triggers. *Linguistic Inquiry* **25**, 407–454.

Goodman, N. (1954) *Fact, Fiction, and Forecast.* Cambridge, MA: Harvard University Press.

Griffiths, Sascha, M. Purver, and G. A. Wiggins (2015). From phoneme to morpheme: A computational model. In H. Baayen, G. Jäger, M. Köllner, J. Wahle, & A. Baayen-Oudshoorn (Eds.), *6th Quantitative Investigations in Theoretical Linguistics Conference*. Tübingen, Germany. http://doi.org/10.15496/publikation-8639

Gualmini, Andrea and Bernard Schwarz (2009) Solving learnability problems in the acquisition of semantics, *Journal of Semantics* 26: 185-215.

Harris, Zellig (1955) From phoneme to morpheme. *Language* **31**, 190–222.

Hart, Betty and Todd Risley (1995) *Meaningful Differences in the Everyday Experiences of Young Children*. Baltimore: Paul H. Brookes.

Hornstein, Norbert (2009) *A Theory of Syntax*. Cambridge University Press.

Huddleston, Rodney and Geoffrey K. Pullum et al. (2002) *The Cambridge Grammar of the English Language*. Cambridge: Cambridge University Press.

King, Harold V. (1970) On blocking the rules for contraction in English. *Linguistic Inquiry* 1: 134–6.

Kirby, Simon (2001) Spontaneous evolution of linguistic structure: An iterated learning model of the emergence of regularity and irregularity. *IEEE Transactions on Evolutionary Computation* **5**, 102-110.

Kirby, Simon (2017) Culture and biology in the origins of linguistic structure. *Psychonomic Bulletin and Review* **24**, 118–137

Lightfoot, David (1989) The child's trigger experience: Degree-0 learnability. *Behavioral and Brain Sciences* **12**, 321–334.

Nola, Robert and Howard Sankey (2007) *Theories of Scientific Method: An Introduction*. Montreal: McGill-Queen's University Press.

Pullum, Geoffrey K. (1983) How many possible human languages are there? *Linguistic Inquiry* **14**, 447-467.

Pullum, Geoffrey K. and Arnold M. Zwicky (1997) Licensing of prosodic features by syntactic rules: the key to auxiliary reduction. Presented at the annual meeting of the Linguistic Society of America. Online at http://www.lel.ed.ac.uk/~gpullum/PZ1997.pdf

Pullum, Geoffrey K. and Barbara C. Scholz (2002) Empirical assessment of stimulus poverty arguments. *The Linguistic Review* **19**, 9-50.

Saffran, Jenny R., Richard N. Aslin, and Elissa L. Newport (1996) Statistical learning by 8-month-old infants. *Science* **274** (13 December), 1926-1928.

Scholz, Barbara C. and Geoffrey K. Pullum (2006) Irrational nativist exuberance. In Robert Stainton (Ed.), *Contemporary Debates in Cognitive Science*, 59-80. Oxford: Basil Blackwell.

Tanaka-Ishii, K. and Jin, Z. (2008) From phoneme to morpheme — another verification in English and Chinese using corpora. *Studia Linguistica* **62**, 224–248. doi:10.1111/j.1467-9582.2007.00138.x