Syntax, part 2

Metatheoretical[1] Interlude

The last section was intended to display some tools for the description of languages, but perhaps more importantly, the discussion there was meant to encourage a certain attitude towards languages. People quite naturally think of languages, especially natural languages like English and Dakota, as rich and vibrant systems, consisting of elaborate and protean networks of associations able, in the hands of skilled practitioners, to sustain drama, comedy, alarm, affection, as well as the more mundane and commonplace features of our interactions with the world. The language we speak and how we speak it lies at the core of our being, and is inextricably interwoven with our personalities and our very natures, and determines in large part who we are and what we do.

But our conception in the last section has been austere and prosaic: languages are sets of strings of objects (or, perhaps, sets of structures). Before applying our tool kit to natural languages, it is perhaps worth pausing to outline the assumptions that lead us to the conclusion that such an attitude is far from banal, but rather on the contrary will set the stage for an inquiry that will give substance and rigor to an analysis of one important aspect of what is it to be a human being.

We will observe later that words have internal structure, but for now let us suppose that each word of, say, English is an atomic unit. It's of course hard to say how many words there are in the English language, and we may not even be sure that the number of words is finite, but for the time being let's assume that this is so. Following the suggestion of the previous section, let's call the set of words in English the *lexicon*. Sentences of English are structured strings of words drawn from the English lexicon.

The lexicon for English is analogous to the very simple vocabulary for some of the artificial languages lately considered, such as {a,b}. And as we did for those languages, we can ponder the free monoid over the lexicon for English. Call it **M**. **M** will consist of the infinite set of all finite strings of English words. Of course, **M** is not English, for it contains many strings, infinitely many in fact, which are not sentences in English. For example, the string *ointment peculiar any* is in **M** but not in English. However every string in English is obviously in **M**. Hence, English is a proper subset of **M**.

We know lots of ways to pick out proper subsets of free monoids, and we might try to write one that picks out English. We'll take some steps in that direction in the next

[1] Like metalanguage is a language about a language, metatheory is a theory about a theory. In this case we're pausing to ask questions like, "What properties should our theory of language have?"

section. But first, we'd better get a clearer characterization of the set we're trying to identify, so that we can know when we're making progress.

There have been many approaches to this question, a question that we might reformulate a bit more dramatically like this: What is English (or for that matter, any other natural language)? Our goal of learning something about human nature (making linguistics in a sense a branch of psychology) will in large part determine how we characterize English pretheoretically, and it is perhaps worth considering alternative characterizations in order to bring out the salient aspects of the approach we will take.

We might, for example, take the set of English sentences to be the set of sentences that have ever been uttered by speakers of the language. But this would surely be an unwise choice, for a number of reasons. First of all, to avoid being circular, we would require an independent characterization of the speakers of the language, which of course cannot make reference to the language itself, since this is what we are trying to define. Second, the vast majority of these utterances, such as the remarks of Benjamin Franklin to himself in his bath, are forever lost, and it would be quite perverse to try to account for data we in principle can never get. More seriously, perhaps, such a conception of the data runs counter to our intuitions about what English is. There are no doubt sentences which are uttered today, and others which will be uttered tomorrow, which we would certainly include in English, but which have never been uttered before in the history of the world. Thus within the net of English we will want to include specimens haven't yet been uttered but might well have been, and be able to distinguish these from strings in **M** (such as the one mentioned above) which are just not sentences in English.

Another problem with the above approach is that we have strong intuitions that not every utterance of an English speaker is a sentence of English, even when that speaker is intending to speak an English sentence. Speakers not only sometimes say something they don't mean, but they also utter strings that on careful reflection they themselves would regard as gibberish. We make so-called "slips of the tongue" and false starts, we get distracted in mid-stream, and occasionally sneezes, burps, snorts, and giggles will erupt inside the sentences of the best of us. Presented with a careful transcript of our ordinary conversations, most of us would be able to identify numerous cases in which what we might call "a production error" has occurred. The ability to do this, to judge our own speech against some standard, suggests that we might do well, as theoreticians, to make a distinction between what actually comes out of the mouths of people we observe, and what would have come out in ideal circumstances, i.e., circumstances in which the standard is adhered to without the complications introduced by the production mechanism.

Our goal is consequently leading us to a rather abstract theory, in the sense that what the theory will be a description of is an idealization, much like a physicist finds it helpful to describe what happens on nonexistent frictionless planes and in perfect vacuums. To introduce terminology suggested by Noam Chomsky, our theory will be a description of the *competence* of the *ideal speaker-hearer*.[2] No one is the realization of this ideal, but we do have reason to believe that the description of the ideal case will be quite useful in illuminating an aspect of what speakers of languages have in their heads. To take Chomsky's famous example, both (1) and (2) below are quite nonsensical, yet every English speaker feels that (1) has a property that (2) lacks, namely, that (1) is English, while (2) is not.

(1) Colorless green ideas sleep furiously

(2) Furiously sleep ideas green colorless

The idea is that our grammatical competence induces us to assign a quite different status to (1) as opposed to (2), both of which are completely novel, and this difference must be formal, that is, something having to do with the form of the string, and not its content, much like strings in the MIU system were well-formed quite independent of any meaning or “sense” that could be associated with them. In other words, it appears languages have systematic properties that cannot be explained on the basis of the uses to which they are put.

We might then choose to attempt to write a description of the competence of the ideal speaker-hearer, hoping that this enterprise will lead to insights about human nature. As in all cases of pretheoretical characterizations of the domain of the theory, we are making an informed guess at this point. It might turn out that this idealization will lead nowhere, that it encourages no insight into our fundamental question. For the time being, however, we'll adopt an optimistic stance and assume that our idealization will in the end be fruitful.

So we set out to describe the competence of the ideal speaker-hearer. We said above that no one realizes this ideal, but there is another sense in which *every* speaker of English instantiates the ideal, since each such adult speaker has the ability to observe how their own performance diverges from this ideal. This implies that each speaker has some representation of the ideal, which, though perhaps not called upon in every day conversation, can be accessed to render judgments about sentences, some of which will be quite novel.

---

[2] In *Aspects of the Theory of Syntax,* Chomsky (1965.3) writes: “Linguistic theory is concerned primarily with an ideal speaker-listener, in a completely homogenized speech-community, who knows its language perfectly and is unaffected by such grammatically irrelevant conditions as memory limitations, distractions, shifts of attention and interest, and errors (random or characteristic) in applying his knowledge of the language in actual performance.” This book in one of the earliest statements of the goals of the then rather new framework of generative grammar, which among other things shifted the focus of inquiry somewhat away from languages and towards the people who speak them, or, to put this another way, away from language and towards the grammars of those languages. For a sometimes amusing, sometimes irritating, but, in my opinion, quite accurate discussion of the tangle of abstractions involved here, see Botha (1989).

The abstraction to underlying competence

These two senses of the ideal indicate that there are two distinct abstractions involved in our characterization of the object of inquiry. The first of these I will call **the abstraction to underlying competence**. The idea here is that we choose to set aside factors which mask in some way what speakers of languages actually know about their language.

For example, Northfield, Minnesota is in the middle of pretty flat farm country, and therefore (3) is a very odd sentence.

(3) Jesse James camped in the mountains surrounding Northfield the night before his gang tried to rob the bank there.[3]

But we clearly would not want to say (3) is not in English, in spite of the fact that it's a weird thing to say. The reason presumably is that geography is not subbranch of linguistics, or to put this another way, the oddness of (3) has nothing to do with the language, but rather with extralinguistic (and hence, for our purposes irrelevant) factors. So we want our account to recognize (3), and account for speakers' "bizarreness reactions" with independent mechanisms.

Here's another example. It's not difficult to construct a perfectly grammatical sentence that is 1200 words or so long. Speakers may find it boring, even rude, but these things are not relevant to linguistic theory. If we develop a grammar that generates such sentences, we should not be in the least alarmed, everything else being equal. In fact, recalling the moral of Exercise 6 in the previous section (this was the *this book has many many...many pages* case), it seems likely that an elegant grammar of English will recognize infinitely many sentences (each finite), and if so there will be infinitely many sentences in English that no speaker has ever uttered and no speaker ever will. This is not any more of a problem than rather pedestrian observation that there are infinitely many even integers, though this means that there are infinitely many even integers that no human ever has or ever will contemplate. Once we recognize that it's easy to characterize infinitely many objects finitely, there's no intrinsic barrier to supposing that human have knowledge of such characterizations.

Some cases are a bit trickier. Consider a recipe for forming a kind of relative clause in English. (We obviously will want to make this more precise in the long run, but this first approximation will do for our purposes.) Take two sentences like those in (4).

(4) a. cats eat mice
b. dogs chase cats

[3] This is depicted in the film "The Great Northfield Bank Robbery" starring, among other people, Robert Duvall.

If you have identical noun phrases in each of the sentences, you can insert the second sentence to the right of one of those noun phrases, append a *that* in the middle, and delete the second of the identical noun phrases:

(5) a. cats [dogs chase cats] eat mice -->
b. cats that [dogs chase cats] eat mice -->
c. cats that dogs chase eat mice

Notice that, in this case at least, the *that* can also be left out:

(6) cats dogs chase eat mice

The relative clause is sort of an answer to the question "Which cats eat mice?" Consider now (7)

(7) a. mice die
b. cats eat mice

If we perform the above operation we get either (8a) or (8b):

(8) a. mice that cats eat die
b. mice cats eat die

Now, try to stay with me here, instead of embedding (7b) in (7a), I will embed (5c) in (7a):

(9) a. mice die
b. cats that dogs chase eat mice

Once again, I take (9b), stick it to the right of *mice* and insert *that*:

(10) mice that cats that dogs chase eat die

(10) will likely sound quite bizarre to most speakers, to say nothing of (11) in which I've deleted the *that*s:

(11) mice cats dogs chase eat die

If you can't quite see what (11) means, consider (12) in which I've changed the relative clauses into the passive voice,

(12) mice that are eaten by cats that are chased by dogs die

(Notice, by the way, that deleting the *that*s in (12) leads to chaos. It's interesting to ask why this is so.) Some people can work themselves into processing (11) by lengthening the pronunciation of *mice*, inserting a pause after *eat*, and pronouncing the intervening material rather quickly. For a lot of people it will seem to snap into focus for at least a brief instant, and then fuzz up again, a remarkable feeling, a rather rare instance where you can almost *feel* your grammar working.

Now, what do we want to say about the bizarreness of (11)? We have two choices. First we could say that (11) is not in English, i.e. it is ungrammatical, and design our account so that it fails to be recognized. Second, we could say that (10) is perfectly grammatical, and that its difficulty is to be attributed to nonlinguistic factors. On this scenario, we want to design our grammar so that (11) is recognized.

If you expect here a clear and compelling argument one way or the other you will be disappointed, for this is a very complicated case. However, most linguists would tend to favor the second option, i.e. recognizing (11) and trying to explain away its difficulty by appeal to factors independent of grammar. Here are some reasons to go this route.

First, people often can with patience and perhaps pencil and paper bring themselves to judge (11) grammatical. Second, there appear to be examples with the identical form which are much easier to process, such as (13)[4].

(13) The game those boys I met invented resembles chess.

Third, there doesn't appear to be a coherent and plausible account in terms of grammar that accounts for the oddness of (11) while allowing cases like (8b). Fourth, and perhaps most important, there does appear to be a promising account of the difficulty of (10) in terms of an appeal to on-line processing mechanisms. (See Chomsky and Miller (1963).)

The distinction we're drawing here is very important, and it's perhaps worth lingering on this point a bit. We're guessing that focusing on what people know about the language they speak will lead to interesting observations about human nature. Furthermore, we're guessing that an illuminating description of this knowledge will include a set of principles and rules that characterize the notion of "possible sentence of L", where L in the language in question. We're further supposing that speakers' judgments of possible sentences are sometimes distorted by factors irrelevant to knowledge of L. We're trying now to see what sort of factors might contaminate our data so that we can take these into account when considering speakers' judgments.

I'm suggesting here that one such contaminating factor may be the mental mechanism that tries to assign a structure to strings of words it encounters. (11) is meant to illustrate a case in which it seems likely that the underlying competence recognizes a sentence that the sentence processor cannot handle, which of course means that there are two interacting but distinct mental faculties, the knowledge of language and the mechanism that implements that knowledge. One reason to think there are two distinct faculties is that they sometimes produce conflicting results. Take for example the sentences in (14), which your knowledge of the grammar of English will tell you are

---

[4] This example is cited in Smith (1989), Chapter 5.

perfectly grammatical, whereas the on-line processor will likely mis-analyze these cases on the first pass at least.

(14) a. the horse ran passed the barn fell
b. while Mary was mending the sock fell[5]
c. fat people eat accumulates[6]

There is no easy way to decide in general when a "bizarreness reaction" is due to one faculty or the other, but scientists try to do all sorts of difficult things, and there's no reason why language scientists shouldn't try to do this one.

Let me put the main point here in a slight different, somewhat more dramatic way. We're interested in characterizing a person's knowledge of language, and we've proposed that this knowledge is quite abstract. So far, I've suggested setting aside factors like "making sense" (as in the distinction between (1) and (2) above), knowledge of geography (and by implication a host of other nonlinguistic domains such as economics, literature, courtship rituals, the rules of baseball, etc.), boredom, death, and the on-line sentence processor. It is possible that, once you take away these factors and those like them, *nothing is left*. That is to say, it is possible that there is no distinguishable coherent body of knowledge that is purely linguistic, and hence it is possible that everything discussed in this chapter is practically pointless.[7] Naturally, there are some reasons to think it isn't so, but it's still worth acknowledging the possibility.

Let's look at another case. Recall strings (14g,h,i) from the last section, here renumbered as (15)-(17):

(15) the pen which I lost was expensive
(16) I don't know where the pen is
(17) *the pen which I don't know where is was expensive

Here, we certainly want the system we construct to recognize (15) and (16), and the question here is whether or not we want it to recognize (17). First, let's display the structure of these examples. (15) contains a relative clause *which I lost* which is a variant of *that I lost* which, according to the recipe of the preceding example in this section, comes from something like (18) by the deletion of *the pen*.

---

[5] This example is from a landmark doctoral dissertation on sentence processing, Frazier (1979).

[6] This example is from Pinker (1994), widely regarded as one of the most enjoyable, accessible, but not uncontroversial books ever written about linguistics. By the way, some brainy Carleton students pointed out to me that *accumulates* is a thing in a science, I think it was geology. They aren't edible, I don't suppose, but I guess people could eat them. But what we're after in this example is accumulation in the sense of gathering in large quantities.

[7] This view is far from uncommon. To mention only two fairly prominent examples, some think it is possible to reduce explanations of language behavior to simple and quite general laws of learning. (Skinner 1957), reviewed (brutally) by Chomsky (1959). Or perhaps we might explain well-formedness purely in semantic terms, for example, Langacker (1987).

(18) the pen [I lost the pen] was expensive

(16) is a slightly different but similar sort of case. So as not to get us distracted by a discussion of the details, let's make the plausible assumption that (16) is underlyingly something like (19).

(19) I don't know [the pen is *somewhere*]

What happens here is that *somewhere* is changed to *where* and this is moved to the front of the embedded sentence, yielding (16). Analogously, the underlying form of the a sentences in (20) and (21) is the b string:

(20) a. I know what Kiki wants ___
b. I know [Kiki wants something]

(21) a. I wonder when the Twins will win ___
b. I wonder [the Twins will win sometime]

The point here is that both relative clauses and cases like (16) contains sentences with "gaps" in them. We will focus on this property at length later on, but for now it's only important that you get the gist of the structure of these sentences. As you can now see, (17) is formed by taking (16), inserting it into the place of the nearly synonymous *I lost the pen* in (18) and relativizing on *the pen*:

(22) a. the pen was expensive
b. I don't know where the pen is
c. the pen [I don't know where the pen is] was expensive -->
d. the pen which [I don't know where the pen is] was expensive
e. *the pen which I don't know where is was expensive.

The important point here is that the crashingly bad (22e) (=(17)) is formed from perfectly fine sources by means of a process we have some independent motivation for. Question: why is this string so bad?

Well, if you've followed this example carefully, you know that it makes sense, in fact, it means more or less what *the pen which I lost was expensive* means. It doesn't appear to run afoul of our independent theories of what things like pens, knowledge, costs, etc. Now, it could be a processing problem, analogous to what we proposed for the example above involving mice, dogs and cats. And, once again, it is very difficult to tell pretheoretically.

To sharpen the discussion a bit, let's propose a principle that might start us toward an explanation of the bizarreness of the (17). If you look at (17) next to its near-synonym (15), you see that one difference between them is that the relative clause in the former has two "gaps" while that of the latter has only one (the subscripts indicate which item the gap is related to):

(23) a. the pen$_i$ [which I lost ___$_i$] was expensive
b. the pen$_i$ [which I don't know where$_j$ ___ $_i$ is ___$_j$] was expensive

In the light of this, we propose Principle G ("G" for "Gap"):

Principle G: Relativizing a sentence which has a gap in it is prohibited.

So the relative clause in (23a) comes from (24a), which has no gap, while the relative clause n (23b) comes from (24b), which has a gap, and therefore the relative is prohibited by Principle G.

(24) a. I lost the pen
b. I don't know where the pen is ___.

For the sake of discussion, let's suppose the Principle G is true. We're asking if G is an aspect of speakers' knowledge of English, or is it rather an aspect of the sentence processor. At this point it is very hard to tell, but one reason to suspect that it is part of the knowledge of English is that (17) seems much more resistant to "improving" by pencil and paper analysis, intonation variation and other such factors which implicate a processing difficulty, as in the *mice cats dogs chase eat die.* To the extent that this argument is convincing, we propose that G is one of the principles that comprise English speakers' knowledge of English.

To summarize so far, one abstraction we want to make is **the abstraction to underlying competence**. We want to get a description of what it is that people know when we say they know a language, and it appears that the one good way to get at this knowledge is to inquire about speakers' intuitions about the well-formedness of sentences. However, judgments about well-formedness are sometimes confounded by irrelevant factors. It's not always easy to determine the source of a "bizarreness reaction", but such considerations must always be kept in mind in we want to make sure the data base for our description is relevant and coherent. Henceforward, we will call underlying competence in language **knowledge of grammar**. The idea is that it is this knowledge of grammar which gives each of us the ability to judge our own performances in the light of an internalized standard.

The abstraction to a homogeneous speech community

It's easy to determine that the internalized grammars of people who are said to speak the same language differ from each other. For example, it would be quite surprising to find two English speakers who share exactly the same vocabulary. And if the vocabularies are different, the grammars those vocabularies are a part of will recognize different sets of sentences, that is to say, different languages. Technical terminology is often responsible for such differences. For example, (25) contains technical or obscure terms that are easily recognized by some speakers of English, but it will make the sentence totally incomprehensible to others.

(25) Peyton recognized the blitz and checked to a screen.

What are we to do about such cases? If the purpose of our inquiry is to discover the principles that regulate the syntax of English, we might for convenience sake simply assume (contrary to fact) that every English speakers knows every word in English. This assumption is probably harmless since it simply means we will set aside differences that are due solely to vocabulary differences since, for the time being, we're not interested in vocabulary differences.

However, other cases are not so easily dealt with. Consider (26), a string which in my language[8] is perfectly normal and natural, and means something like "Would you like to accompany me?"

(26) wanna come with?

I'm aware that there are some speakers of English who find such strings very jarring, and of only marginal acceptability. (27) illustrates another case.

(27) a. I don't drink a lot of coffee these days
b. I don't drink a lot of coffee anymore
c. I drink a lot of coffee these days
d. I drink a lot of coffee anymore

Some speakers of English find (27d) just fine, while others reject it. Another such case appears in (28).

(28) a. my car was washed
b. my car got washed
c. my car needs washed

Many people in the Pittsburgh area find (28c) completely normal, but this view is not widely shared elsewhere in the English speaking world. These examples illustrate the point that the dialects of different groups of people are typically partially disjoint. Dialects are related to each other not by the subset relation but rather by a sort of family resemblance. What this means is that "English" (and all other natural languages) is a concept rather like the one associated with the word "pig". Pigs don't have to be grey or pink to be pigs, in fact I suppose (not knowing much about pigs) that in principle a pig could be any of whole range of colors. But no pig is colorless. In other words, the concept associated with "pig" is abstract. Pigs all have a color, but not any particular one. No actual object is pure pig and nothing else. Every real pig has properties not shared by all other pigs.

This doesn't prevent biologists and others from making generalizations about pigs. Notice that it's also not the case the people who study "pig-nature" are interested only in

[8] I was born in Chicago's South Side, and lived in Illinois until I was 18.

properties that all pigs share. For example, I imagine that porcine reproduction is a complicated and fascinating topic, and there are probably many generalizations that apply only to the males or only to the females of the species.

The lesson for us here is that even though English is not instantiated in any particular speaker of the language, there are still likely to be interesting and true generalizations to be stated about the abstract object English. We expect to find differences in the grammars of English speakers, and these differences will result in varying judgements about particular strings. If there is a specifiable object called "English" it is likely to be quite abstract.

Let’s try to find it.

§§§§§§§§§§

Syntax 2

Phrase Structure Rules do two things, simultaneously. First, they specify what we might call the hierarchical organization of phrases, and second, they specify the left-right order of constituents within those phrases. For example, consider the set of phrase structure rules in (1)

(1) S -> B C D
B -> F G

These rules generate the tree in (2).

(2)

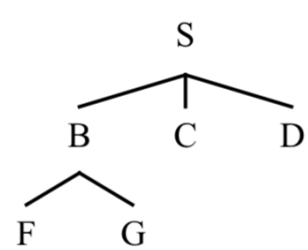

Hierarchically speaking, the rules tell us that F is a part of a B, and B is a part of an S. "Linearly" speaking, the rules tell us that within a B, F precedes G.

Let's define PS rules more precisely. As we did in the last section, we can think of the items that appear at the nodes of trees as falling into two classes, the nonterminal vocabulary (V_N), such as **S** and **1CON** from PL, and the terminal vocabulary (V_T), such as **p** and **q** in PL. We can say then that

(3) A Phrase Structure Rule is a rule of the form

$A \text{ -> } b_1 \ldots b_n$

where A is a single member of V_N and b_1 through b_n are symbols drawn from either V_N or V_T. (For set theory enthusiasts, this is $V_T \cup V_N$, i.e. the union of V_N and V_T.)

A Phrase Structure Grammar, then, is any collection of Phrase Structure Rules. We interpret these rules as we did in section 1, which is to say that we regard them as licensing trees in a way which is by now familiar to you.[9] We now turn to natural languages.

It's not so easy to say how one should begin using PS rules to describe the syntax of natural languages, and so we will follow a time-honored tradition in scientific investigation. We'll make several assumptions, and then see if the resulting theory we construct is illuminating in the way we want it to be. If it isn't, we'll go back and revise the assumptions. If the theory looks promising, we'll regard this as tentative confirmation of the assumptions, celebrate our good fortune in making a good initial guess, and forge ahead.

First, let's have a brief look at an approach that is consistent with the PS Grammar framework, but which we'll certainly want to abandon. This approach looks like this:

(4) S -> Bill left
S -> the man left
S -> Toledo is in Ohio
S -> My brother told me that chocolate is toxic for cats

etc.

We would get trees like:

(5)

[9]There are other ways of interpreting PS rules, which we set aside for now.

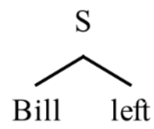

and similarly for the rest of the examples. (4) contains perfectly legitimate PS rules, and (5) is a fine tree, which says that the string *Bill left* is a sentence, which, of course, is true. When we discover a sentence in English that we haven't accounted for, we simply write down that sentence, adjoin an arrow and an S to its left, and add this new rule to the grammar. What's wrong with this approach, then?

There are a number of problems with this "minimalist" approach[10], but I'd like here to concentrate on just one of these. We can give a good argument that this sort of grammar cannot be the sort that is represented in the minds of speakers of English. You probably have anticipated this argument already, but here it is fleshed out a bit.

Speakers of English don't learn sentences of the language one at a time. We can see this by observing that speakers of English have clear intuitions about the grammaticality of strings of English words even though they have never encountered those strings before. We already noticed this in the examples in (14) of section 1, which I repeat here:

(6)
a. babies sleep in cribs
b. *sleep babies cribs in
c. colorless green ideas sleep furiously
d. *furiously sleep ideas green colorless
e. do you often walk to school
f. *walk you often to school
g. the pen which I lost was expensive
h. I don't know where my pen is
i. *the pen which I don't know where is was expensive
j. how Ann Salisbury can claim that Pam Lauder's anger at not receiving her fair share of acclaim for *Mork and Mindy's* success derives from a fragile ego escapes me.
k. *how Ann Salisbury can claim that Pam Lauder's anger at not receiving her fair share of acclaim for *Mork and Mindy's* success derives from a fragile ego escape me.

[10] Notice that this is “minimalism” with a small “m” and is meant in the ordinary way. It should be sharply distinguished from “minimalism” in the sense Chomsky intends in *The Minimalist Program* (2005). You might want to have a look at Chomsky’s book someday, but probably not today.

The point for our current purposes is that speakers of English have no trouble deciding which of these strings is in the language they speak, even though they may not have encountered those strings before. To say then, as the "minimalist" approach does, that learners of English simply add new sentences to their grammars as they hear them is clearly wrong. A more accurate approach would claim that learners of a language construct a rule system that makes predictions for them about many as yet unseen cases. We therefore need to make our PS grammar more elaborate. Here's one way to go about this.

Consider the sentences in (7).

(7) a. men should leave
b. the men should leave
c. the tall men should leave

These strings might be analyzed as having two parts: *should leave* preceded by various strings of words. Let's call these various strings *Noun Phrases* (NP). We therefore adopt the following PS rules:

(8) S -> NP *should leave*
NP -> *men*
NP -> *the men*
NP -> *the tall men*

Now, what about the internal structure of the NP? It appears to consist of one obligatory item, *men*, and some optional items *the* and *tall*. Following the strategy we adopted in version three of PL in section 1, we assign these lexical items to categories. We'll choose N for the noun *men*, ADJ for the adjective *tall* and DET for what is sometimes called the "determiner" *the*. We thus get:

(9) S -> NP *should leave*
NP -> N
NP -> DET N
NP -> DET ADJ N

Lexicon: N: *men*
ADJ: *tall*
DET: *the*

Already we see that we have some of what I will call *projectability*, by which I mean the property some systems have of making predictions about as yet unseen cases. For example, suppose we now encounter the following sentence:

(10) some tall men should leave

This leads us to add *some* to the lexicon as a DET. When we do so, we predict the grammaticality of several new cases. Confirmation of this move comes from the observation that the sentence in (11) is well-formed.

(11) some men should leave

The projected cases multiply rapidly as new items are added to the Lexicon.

Exercise 1

The fact that *Many short men should leave* is grammatical leads to the projection of several other grammatical sentences. List some of these.

The NP rules together say that NPs have an obligatory constituent N (which we can call the *head* of the NP) together with a preceding optional determiner and an optional adjective in that order. We might add a new piece of equipment to the machinery of our PS rules in order to say this explicitly. We can *collapse* the three NP PS rules in (9) by means of parentheses. (These obviously function in a different way than the parentheses did in PL.) Consider then the revised grammar in (12):

(12) S -> NP *should leave*
NP -> (DET) (ADJ) N

Lexicon: N: *men*
ADJ: *tall, short*
DET: *the, some*

The parentheses indicate that those symbols they enclose are optionally selected when expanding the node on the left side of the arrow. This change is not just cosmetic. In fact, it makes a prediction that there is a fourth kind of NP that we haven't examined yet.

Exercise 2

Is this prediction true? Justify your answer.

Exercise 3

For each of the following rules, write out all the unparenthesized rules they abbreviate.

a. A -> B (C) (D E)
b. A -> (B (C)) D
c. A -> (B) C (D)

Suppose now we observe that (13) is grammatical in English.

(13) the men with violins should leave

The question is what to do with the words *with* and *violins*. Let's assume that *with* is a preposition (P) and *violins* is a N. Notice that there are many possibilities as to which rule or rules to add so that our grammar will generate this sentence. Here's some of them:

(14) a. S -> NP (P) (N) *should leave*
b. S -> NP (P NP) *should leave*
c. S -> NP (PP) *should leave*
PP -> P N
d. NP -> (DET) (ADJ) N (P NP)
e. NP -> (DET) (ADJ) N (PP)
PP -> P N

<u>Exercise 4</u>

Draw the trees for (13) that would result from choosing (14c) or (14d).

Notice that the data we have considered so far won't select any of these hypotheses as preferable to any other. We need more data.

First, then, let's consider some other cases concerning the nature of *violin*. Assuming this is an N, we'd like to know if it is also an NP. Well, if we assume that it's an NP, we make lots of predictions. If these work out, we'll be inclined to assign *violin* NP status. Here's some sentences predicted grammatical by (14b) but not by (14a):

(15) a. men with the violins should leave
b. men with expensive violins should leave
c. men with the expensive violins should leave

Since all of these are grammatical, and this would be predicted by saying the *violins* is an NP, we should certainly adopt this. Thus we have so far

(16)

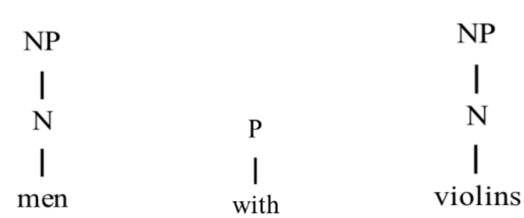

We're still stuck with the question of how the preposition *with* and the NP *violins* gets hooked up to the tree. Now, I know that you have a strong intuition that *with violins* is a prepositional phrase, and that this PP is a part of the NP *the men*. It's good to have that intuition, but we do need to be careful. Sometimes our intuitions are flat wrong. (For example, I have the very strong intuition that the earth is stationary and the sun revolves around it.)

Here's a sensible procedure. We provisionally adopt the analysis that represents the intuition, and then stay on the look out for confirming or disconfirming evidence. So suppose we adopt none of the proposals in (14) but one that is close to the proposal in (14e):

(17) S -> NP *should leave*
NP -> (DET) (ADJ) N (PP)
PP -> P NP

In order to scope out predictions that this decision might make, let's turn to the analysis of the other part of the sentence. Along side sentences we've seen so far, we also find:

(18) a. the men should stay
b. the men might stay
c. the men should buy a book

Once again, I'll make some unmotivated assumptions here for the purposes of pursuing our main question, and as usual, we should keep in mind that I've made them and watch out for relevant evidence. It looks like *should* and *might* are the same sorts of things, we'll assign them to the category AUX (for auxiliary verb), and that *leave* and *stay* are the same sort of thing, which we'll call V (for verb). We thus have:

(19) S -> NP AUX V
NP -> (DET) (ADJ) N (PP)
PP -> P NP

with the obvious changes in the lexicon

Consider now (18c). Once again, there are many ways we could modify our grammar in order to generate this string. We might in this case follow the general strategy we've adopted and propose another phrase following the auxiliary verb, namely a verb phrase (VP). VPs will then consist obligatorily of verbs with optional NP *complements*. We therefore modify the grammar in (19) to (20):

(20) S -> NP AUX VP
NP -> (DET) (ADJ) N (PP)
VP -> V (NP)
PP -> P NP

As usual, a move like this makes many predictions, in particular, that we ought to find all kinds of NPs following verbs, not just a DET N NP like *a book* which motivated the rule in the first place.

Exercise 5

Are these predictions confirmed? Justify your answer.

Let's investigate a bit further the structure of the English VP. We also find sentences such as

(21) Kiki lives in Northfield

We immediately suspect that *in Northfield* is a PP, and that it's inside the VP. This would lead us to write another VP rule as in (22). ((21) also reveals that auxiliary verbs are optional constituents of sentences, so we modify our grammar accordingly by enclosing AUX in parentheses.)

(22) S -> NP (AUX) VP
NP -> (DET) (ADJ) N (PP)
VP -> V (NP)
VP -> V (PP)
PP -> P NP

But now we might wonder whether or not these two VP rules should be collapsed into one. Were we to do so, say along the lines of (23), we would make several predictions.

(23) S -> NP (AUX) VP
NP -> (DET) (ADJ) N (PP)
VP -> V (NP) (PP)
PP -> P NP

For example, we're predicting here that we ought to find sentences in which the verb is followed by both an NP and a PP, and as you can see, this is confirmed by sentences such as those in (24)

(24) Sally put the book on the table
John sent a letter to Sue

Now, here's an interesting observation: Our grammar predicts that there should be cases of structural ambiguity in the VP. Consider, for example, a sentence like

(25) Mary kissed the children in the kitchen

Exercise 6: Our PS rules will assign this two structures. Give them. Why is this good news for our theory?

To summarize what we've done so far, here's our current proposal for the grammar of English:

(26)

S -> NP (AUX) VP
NP -> (DET) (ADJ) N (PP)
VP -> V (NP) (PP)
PP -> P NP

plus a Lexicon of the appropriate sort (We'll return to this later on.)

This grammar generates a wide array of trees, and together with the Lexicon, generates many sentences. We still, of course, have some problems. For one thing, our grammar will *undergenerate*. This means that there are sentences in the target that our grammar won't generate, such as

(27) do you often walk to school
the pen which I lost was expensive

But we are getting *Babies sleep in cribs* and lots of other perfectly well-formed strings, so we ought to be provisionally pleased with our progress. Our grammar is also *overgenerating*, which is to say that it recognizes strings that are not in the target, such as

(28) *John admires
* Bill like Mary

We'll try to fix these problems shortly.

Exercise 7

How many distinct trees (without lexical items) will our grammar generate? Justify your answer.

§§§§§§§§§§

Before facing the difficulties our PS rules encounter, let's explore some other kinds of operations that might be useful in describing English. Consider the sentences in (29).

(29) a. Snow falls in Northfield
b. Snow fell in Northfield

Our current grammar generates both of these sentences. But speakers of English have a clear intuition that these sentences are closely related to each other, and our proposal does not represent this. An English speaker might say,"(29a) is a present tense sentence, and (29b) is the past tense version of the same sentence."[11] Our strategy so far has been to respect intuitions like this (keeping in mind that they might be mistaken) and design our grammar so that it displays them. Now, there are several ways to set up the display in this case. One way would be to give some structure to our Lexicon, by figuring out a way to say, in this case, that *fall* and *fell* are different versions of the same verb. We're going to have to do this anyway (since it's a fact about English that we'd want our description to represent). But (for pedagogical reasons and anticipating future developments) I'm going to propose a slightly more elaborate solution.

Here's all the forms of the verb *fall*:

(30) fall (falls), fell, falling, fallen

For now, we won't worry about the *fall-falls* alternation, and we'll consider these two as being one form, *fall*. *Fall* and *fell* are different from *falling* and *fallen* in that they can occur in sentences that don't contain an auxiliary verb:

(31) a. Snow falls
b. Snow fell
c. *Snow falling
d. *Snow fallen

In general we find this pattern holds of every verb in the language, that is to say, when no auxiliary verbs are present, verbs can take two different forms. This is part of the generalization we want to represent.

It's convenient to have a name for these two forms of verbs. We could choose any names we like. For example, we could call *fall* "the Carleton form" and *fell* "the St. Olaf form". However, even though it's a bit misleading (see fn. 2), we'll follow traditional analyses and call *fall* (*admire, swim,* etc.) "the Present tense form" and *fell* (*admired,*

---

[11]Of course, the simple present tense is not ordinarily used to refer to an event happening simultaneously with the utterance of the sentence, but rather to a characteristic or habitual property. To talk about an on going event, English speakers often revert to what we will call the *progressive form* (which in this case would be *Snow is falling in Northfield*). We'll take up the description of cases like this soon.

swam, etc.) "the Past tense form". We also call *falling* "the Progressive form" and *fallen* "the Perfective form". So verbs appear in four different forms. Present and Past forms appear without auxiliaries, while Progressive and Perfective forms require auxiliaries.

To represent this, I make the following changes in our grammar. First, I remove the parentheses from AUX in the PS rule that expands S, and write a rule that expands AUX into either PRES or PAST. (This much encodes the generalization that every sentence is either in what we are calling the present or past forms.)

(32)

S -> NP AUX VP
NP -> (DET) (ADJ) N (PP)
AUX -> PAST
AUX -> PRES
VP -> V (NP) (PP)
PP -> P NP

plus a Lexicon of the appropriate sort

We can introduce here a bit of notation that will make our grammar more compact. When items are enclosed by curly braces { and } on the right side of a PS rule, we interpret this to mean that exactly one of these items must be chosen. For example, the rule in (33) abbreviates the rules in (34).

(33) A -> B {C,D,E}

(34) A -> B C
A -> B D
A -> B E

With this handy device, we revise our grammar like so:

(35)

S -> NP AUX VP
NP -> (DET) (ADJ) N (PP)
AUX -> {PRES, PAST}
VP -> V (NP) (PP)
PP -> P NP

plus a Lexicon of the appropriate sort

Now I elaborate on the Lexicon. I want to divide the Lexicon into two parts. The first, which for reasons that will become clear in a moment, I'll call the *d-lexicon*. The d-lexicon contains all the words in the language, except that it only contains what is known as the "citation form" of verbs. This is what is sometimes called the "infinitive form", for example *dream* as in *To dream is to enter into a certain kind of brain state*. All other forms of verbs are relegated to the *s-lexicon*, which I'll describe more fully soon. When one has built a tree according to the PS rules, lexical insertion is done by the d-lexicon only.

So now I'm getting trees like (36).

(36)

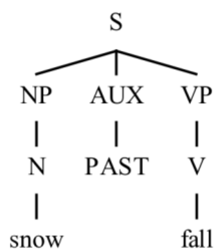

Now what I want to do is write a rule that transfers to the verb the information that PAST is in AUX. I want to do this in a way that will apply to all relevant cases in the language. Here's one way to do it. First, since I might remember from earlier cases (like (18): *the men should stay, the men might stay,* etc.) that eventually I'll have more items in AUX, I invent a category label for PRES and PAST. For this I'll bow to tradition, and select T (for "tense").[12] Thus we have:

(37) AUX -> T
T -> {PRES, PAST}

This gives us tree like (38):

(38)

[12] I'm uncomfortable with this label because this classification is *not* directly related to the time reference of the sentence, as the label might indicate to the unwary. The "tenses" we're talking about here are formal, syntactic tenses, and are not reliable indicators of the semantic tense of the sentence as this is normally understood. Sometimes tradition is a powerful thing.

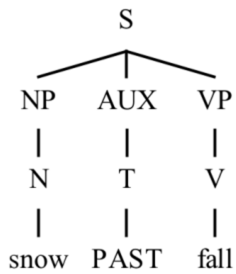

Now what I'll do is write a rule that is reminiscent of the rules in the MIU system. You'll recall that that kind of rule has an "if-then" format, consisting of a structural description (SD) which specifies the class of objects to which the rule applies, and a structural change (SC) which specifies the operation to be performed on the object. A general term for this kind of rule is *transformation*. Most if not all of the transformations we will write take a class of trees as inputs and deliver a class of trees as outputs, like rules in the MIU system take a class of strings as input and give a class of strings as output. I'll first state the rule and name it (again, for reasons that will become clearer later) Affix Hopping. Then I'll explain how it is to be interpreted.

Affix Hopping:

SD:	X	{PRES, PAST}	V	Y
	1	2	3	4
SC:	1	0	3	4
			[2]	

As in the MIU system, in order to see whether or not Affix Hopping (or any other transformation, for that matter) applies to a given tree, I must first factor the tree in order to see if there's a way in which it satisfies the structural description. In this case, I see that the SD has four parts: the variable X, either PRES or PAST (braces are interpreted much like they are in PS rules), V, and a variable Y. The question now is whether I can divide up the tree so that the SD will fit, template-like, over it. It's possible to very precisely state what counts as a factorization of a tree, but we can make use of our convention that we always draw trees heading down the page with the root at the top to make this easier to see. To factor a tree, draw a line through it which never crosses or goes underneath a branch, nor over the root. The factors are the nodes the line goes through. Here's some possible factorizations of the tree in (38):

(39)

a.

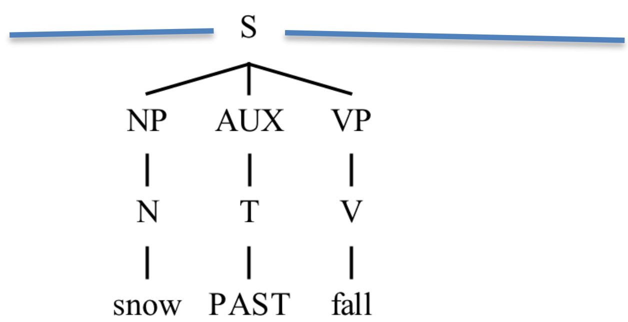

b.

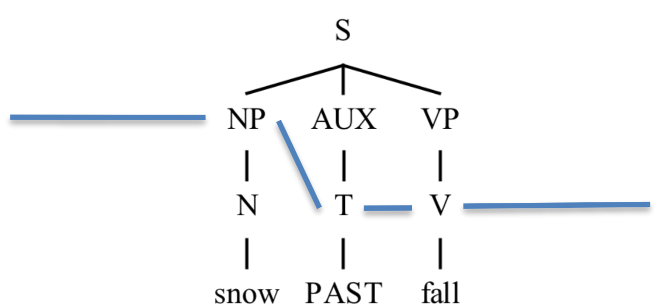

c.

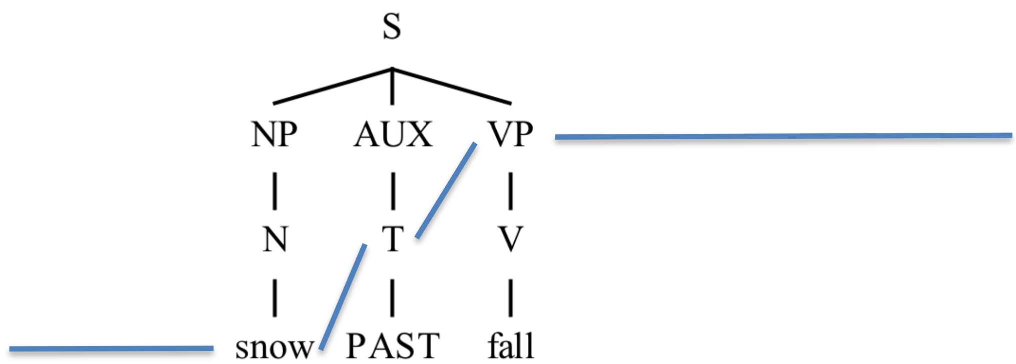

(39b) tells us that that one can move around the "levels" so long as one obeys the other rules, thus giving us NP-T-V is a possible factorization. (39a) reveals that S all by itself is a possible factorization. (39c) shows *snow*-T-VP as a possible factorization.

(By the way, if you think of PS rules as specifying lines in a derivation (rather than trees) as we did with our first rendition of PL back in part one, every legal line in a possible derivation would count as a possible factorization of the tree. For example, consider a line-by-line derivation *snow PAST fall* according to our current grammar:

S
NP AUX VP
N AUX VP
N T VP
N T V
snow T V
snow PAST V
snow PAST fall

Each one of these lines is a possible factorization. There are of course several other equivalent derivations, such as one that expands VP before NP, and every line in those derivations would be possible factorizations as well.)

The following are not possible factorizations:

(40)

a.

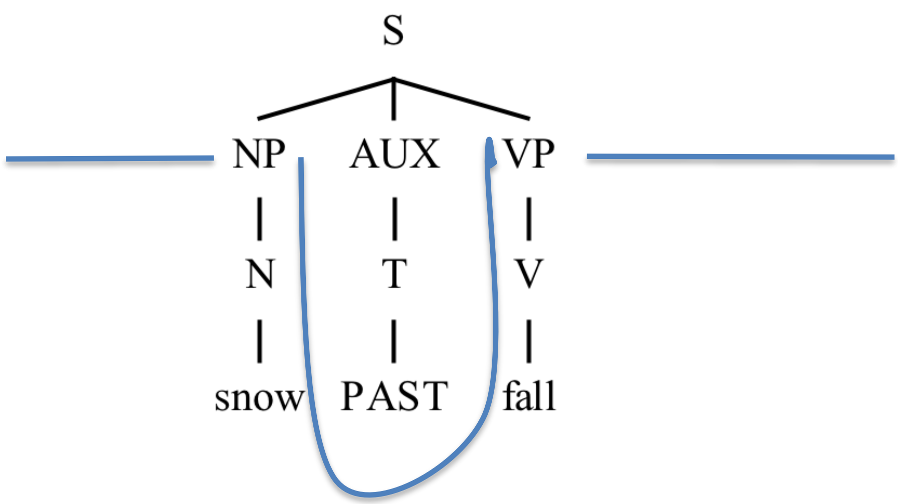

b.

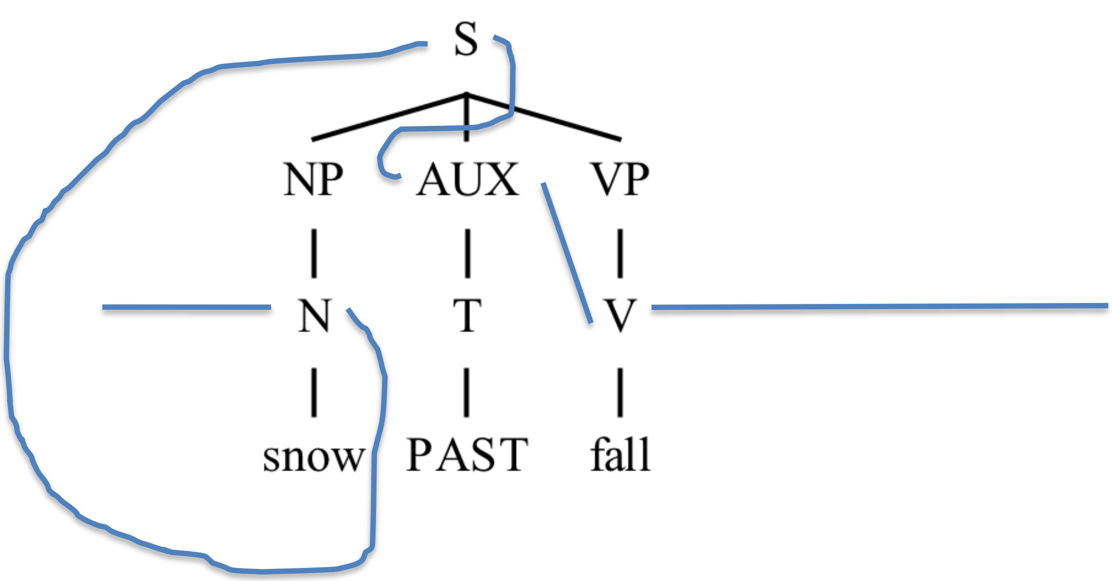

Therefore, neither NP-VP nor the preposterous N-S-AUX-V are possible factorizations. (These will not show up as lines in any legal derivation in the sense of the above parenthetical remark either.)

Returning now to Affix Hopping and its application to the tree in (38), we see that the SD is indeed met:

(40)

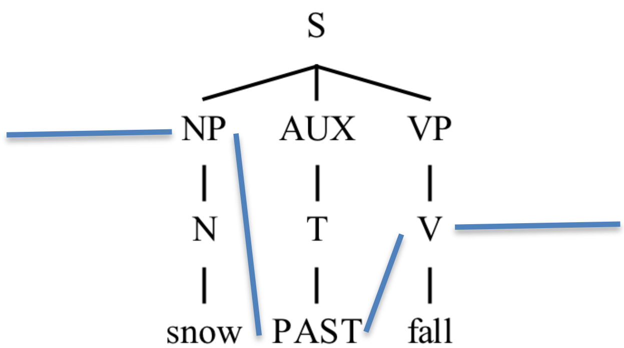

As in the MIU System, variables in the SDs can stand for any stretch of tree (it needn't be a constituent) or nothing at all. In the present case, the X variable picks up the NP, and Y is null.

So we know that Affix Hopping will apply to the tree in (40). Now let's turn to what the rule does to the tree. You'll notice that each of the factors in the SD are numbered. This is done purely for convenience. The SC tells me what to do with each factor. In this case, I'm to leave factor 1 alone, and likewise with factor 4. The 0 under factor 2 tells me to delete it. The operation performed on factor 3 is new for us, and I'll now explain what it is.

Recall that what we want to do is get the information about the contents of T to the verb. Formally, what the rule says to do is mark factor 3 with a "feature" consisting of factor 2. Here's what this looks like:

(41)

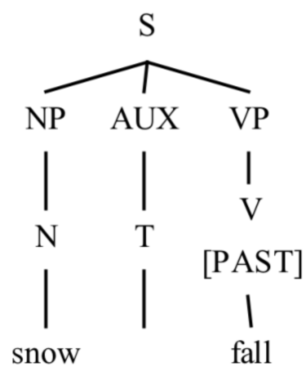

We'll discuss features later on, but for now we only need to note that a feature is anything in the tree that is enclosed is square brackets.

We need one more convention:

(42) Feature Percolation

When a transformation assigns a feature to a node A, that feature automatically appears on every node A dominates.

Application of Feature Percolation gives us (43).

(43)

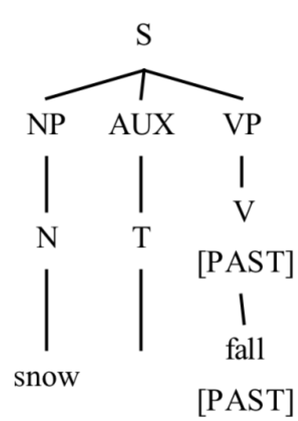

At this point, the s-lexicon applies to look up what the PAST form of *fall* is. We'll call this operation Lexical Lookup. The result is (44).

(44)

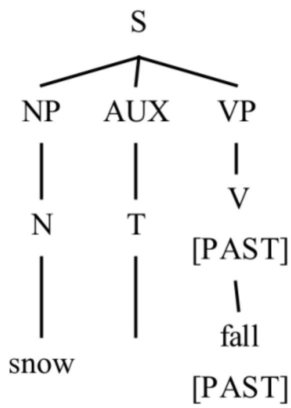

You may have noticed that our theory is getting a bit complicated. Here is its overall structure:

(45)

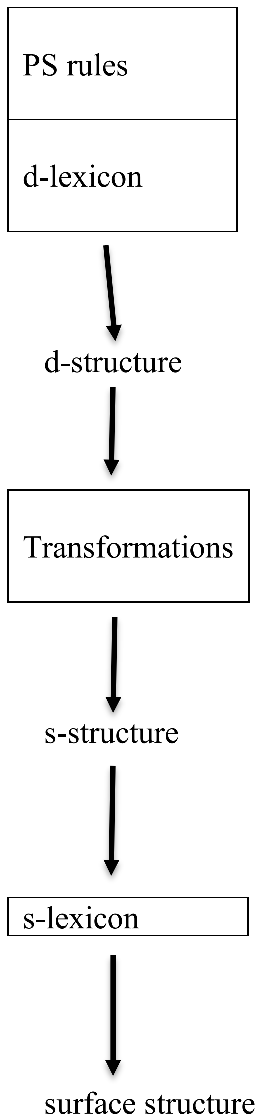

General Principles:

We can call the items listed in the boxes "components". The PS Component and d-lexicon together are often called the "base component". The output of this component is class of trees, each of which is called a "d-structure".[13] The transformational component consists of a set of rules that modify tree structures, applying to the output of the base, and resulting in "s-structure", another class of trees. The s-lexicon applies to s-structures, producing what we will call "surface structures". "General Principles" are just that. They

[13]In the early days of generative grammar, this level of analysis was called the "deep structure", but this term has since been discarded because it invited unwelcome implications that the level was profound in some sense. "D-structure" is felt not to have such misleading connotations.

apply so as to regulate the kinds of derivations produced. One of our goals is to specify a grammar so that the concatenated leaves of every surface structure is a sentence of English, and every sentence of English is the concatenated leaves of at least one surface structure.

Here, then, is our grammar so far, which can be seen as an instantiation of the general theory in (45):

(46)

PS Rules:

S -> NP AUX VP
NP -> (DET) (ADJ) N (PP)
AUX -> T
T -> {PRES, PAST}
VP -> V (NP) (PP)
PP -> P NP

a **d-lexicon** which contains verbs only in their citation form

Transformations:

Affix Hopping: (obligatory)

SD:	X	{PRES, PAST}	V	Y
	1	2	3	4
SC:	1	0	3 [2]	4

an **s-lexicon** that adjusts verb morphology

General Principles:

Feature Percolation

When a transformation assigns a feature to a node A, that feature automatically appears on every node A dominates.

This theory licenses derivations, as in the following derivation of *Snow fell*:

d-structure:

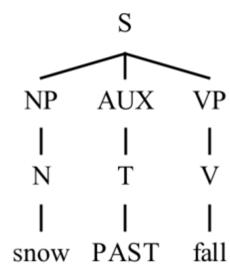

Affix Hopping:

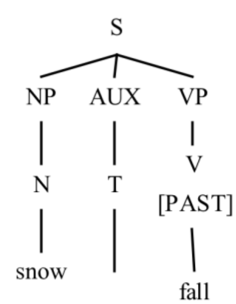

Feature Percolation:

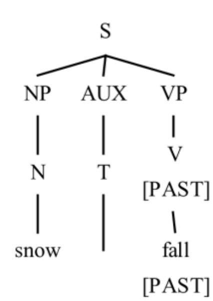

s-structure

Lexical Lookup:

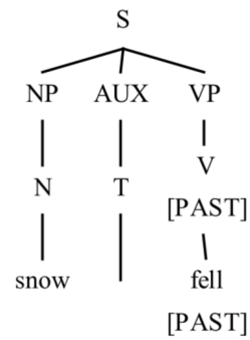

surface structure

<u>Exercise 8</u>

Give derivations for the following:

a. Bill ate a peach
b. Mary swims in the pool

<u>Exercise 9</u>

I made Affix Hopping an obligatory rule. What this means is that it *must* apply if its structural description is met. Why is this, as Martha Stewart would say, a good thing? (Hint: consider the alternative.)

Let me try to summarize where we are, and make an observation. We are building a model of the rather abstractly characterized knowledge that an ideal speaker/listener has of English. We hope this model will provide some insight into what exactly it is that people know when they know a language, and will also shed light on what it is to be the sort of organism that can learn and use such things, i.e. what it is to be a human being.

We proceed, basically, by making some guesses, formulating them as precisely as we can, and testing them out. We might call the guesses “hypotheses” to convey an air of sophistication about them, but whatever we call them, it worth noting that the guesses come before data acquisition, both logically and temporally. The collection of hypotheses is known as a theory. In the case of syntax, the theory has two levels. First, we build a theory of a language, called a grammar. Grammars make predictions about what data we should find when we look, and they also implicitly suggest further hypotheses. Second, on a more abstract level, we can ask what our theory of English suggests about what we might expect to find when we look at other languages, such as Vietnamese or Chamorro or Xhosa. Looking at it the other way, we might have an educated guess about what languages are like, and then we would build models, i.e. grammars, that conform to that hypothesis and see if they explain the data in an illuminating way. If so, we might tentatively put forward a suggestion about what human languages are like, and therefore, what sort of mental characteristics human beings have.

But we're a long way from this. Our grammar at this point is a baby one. It recognizes a handful of sentences, but it woefully both under- and over-generates. You'll spend the next few weeks (at least) trying to make it better, hopefully with imagination, insight, and rigor.

The guesses we've made, or rather that I've forced upon you, aren't exactly random but they shouldn't be taken as received opinion, either. I daresay every feature of our grammar is controversial. But the good news is that we're off the ground. Our goal has been to focus the problem and give you a method for working on a solution and some preliminary suggestions about how to go about testing and defending your ideas. If we've done that, we're in good shape.

One final observation. Perhaps you have the intuition that our method of getting the verbs tensed in English is far more complicated than it needs to be, taking as it does PS rules, a transformation, features, a general principle on features, and two distinct lexica. You may think that the "real" way we do it has got to be a lot simpler. This is a good intuition to have. But the gauntlet is down. We have a theory. If you think you have a better one (one that is *more true*), you know, I hope, what to do.