

## 1 What we're trying to accomplish

There are lots of things that we as native speakers know about our native language. We can make very fine-grained judgments on the acceptability and interpretation of sentences in our native language that indicates a relatively complex system of knowledge. We in essence compute in some way whether a sentence is grammatical or not, and what interpretation it has. Even given a sentence we've never heard before, we know how it can and can't be arranged and what it would mean.

The theory we've been developing over the semester is an attempt to model and understand the system that underlies these judgments of grammaticality and interpretation. It is an attempt to answer questions like *Why isn't "Students the books read" English?* *Why can't "she" refer to Mary in "Which book of Mary's did she read?"*? *How is it that English speakers all know this?*

The basic model we're working with is one in which a set of component parts (words, at least in a certain sense) are provided to a system (a computational procedure of some kind) that assembles them into a structure. Sometimes it will succeed, sometimes (for example, if it is given a set of incompatible words) it will fail. When it succeeds, the structure will come out in a certain (deterministic) way. Our goal is to define a system that is capable of assembling *all*—and *only*—structures for the sentences that we intuitively judge to be grammatical. If the system achieves this, then we have gained a certain degree of understanding of the sort of thing our language knowledge is—that is, we are discovering what it *is* to know a (native) language.

A second consideration is a sort of philosophical one: Children exposed to language data will rapidly come to have the same language system as the adults in their environment. They have reliable success in reaching this point, and many aspects of the system appear to be too complicated to have been deducible from the data the children receive. Given that children seem to “go beyond their input” in this way, we assume that the children must also bring to the task of language acquisition something that constrains or guides the process of language acquisition, some kind of “language instinct” that children have simply by virtue of being human.

In addition to being able to describe, say, the system that underlies a native English speaker's judgments about English sentences, we would like also to understand the nature of this “language instinct.” What constraints does it place on the kinds of languages a child can learn, how is it possible that kids can learn so quickly and succeed so reliably?

The hypothesis we have been exploring is that the language instinct, or *Universal Grammar*, consists of the building blocks of a language system—essentially already in place at the point where children are acquiring language—with a narrow range of options that can differ from language to language. The child acquiring a language then has only the task of determining which options are appropriate for the language spoken in their environment. This model of language acquisition (and of syntax) is often referred to as a *Principles and Parameters* model, the idea being that there are a set of basic principles of syntax that are shared by all humans—by virtue of being human—and languages differ on a relatively small number of “settings” or parameters.

We have taken English as our model language for the purpose of exploration and we have looked at it in some depth, but under the Principles and Parameters hypothesis, most of what we learn by looking at English in such detail is equally applicable to other languages as well. Yet, it is still important to consider how things work in other languages too, because this helps us understand what the options are.

The *minimalist program* is a more specific approach to the Principles and Parameters model of language, which takes as its basic hypothesis that syntax is nothing really more than a means of mediating

between sound and meaning. Under this view, the mind is understood as segregated into different “modules” responsible for different things. One (at least) module is responsible for understanding, thinking, reasoning, and so forth—the *conceptual-intensional* system. Another module is responsible for running our articulation and perception systems. Each module is separate from syntax, but syntax serves as a point of contact between the modules, generating structures that each module can interpret. Syntax’s job is to assemble representations that are in a form that these other systems (the conceptual-intensional and articulatory-perceptual systems) can use. We assume that there is a certain class of things (structures, properties) that these modules are capable of interpreting—those are “interpretable” and everything else is “uninterpretable.”

So, the design of our computational system for syntax is one where the building blocks (lexical items, vocabulary items from the language) are assembled in such a way that they are interpretable—nothing that is uninterpretable is left over. If something uninterpretable is left over, the structure cannot be interpreted—we have an ungrammatical structure.

The goal of this model of syntax is to explain the (apparent) complexity we find within and across languages in terms of the procedure for rendering assemblies of lexical items interpretable.

## 2 The properties of words

When we think about the words from which sentences are formed, we quickly see that they come in different classes. There are major classes of words, those we call nouns and verbs, adjectives and prepositions. Within those classes, there are also subclasses of words, such as the transitive verbs or the intransitive verbs.

So, individual words have properties, like being a verb, or being a transitive verb. This is something we know about the words individually—not necessarily as part of the system, but as part of what we learned when we learned the vocabulary of the language.

As a way for us to indicate the properties of words, we list them as *features*—where “features” is really just a technical term for “property.” A noun is said to have the feature [N] or perhaps [category:N] by virtue of being a noun, a verb is said to have the feature [V] or [category:V]. To distinguish transitive verbs from intransitive verbs, we need to rely on a concept of “needing something”—transitive verbs are those that “need” two noun-category things, and intransitive verbs “need” only one. Some verbs don’t seem to need any, such as *rain* or *snow*. So, we need to have a way to record the “needs” of a verb.

Moreover, we want to be sure that the needs of the verb are met before the structure is interpreted—we suppose that a transitive verb isn’t really interpretable without having both of its nouns. So, we say that a transitive verb has two uninterpretable features, one for each of the nouns it needs, and once we provide the verb with the nouns it needs, then the verb becomes interpretable.

This leads naturally into a formalization of needs and means by which those needs can be satisfied in terms of “uninterpretable features” that need to be “checked” (or removed) before the interpretation system can use the resulting structure. We write uninterpretable features like [*u*N], meaning that the lexical item that carries this feature cannot be interpreted until it is “given” something that matches the need, in this case something with the category feature [N].

Words have many other properties as well. In many languages, some words seem to carry information about things like plurality (number), semi-arbitrary sub-class membership (gender), and relationship to the discourse participants (person). In order to describe the English pronoun *them*, we need to specify that it is 3rd person (neither speaker nor hearer) and plural.

### 3 Constituents and the hierarchical structure of sentences

It is readily apparent that sentences have a certain kind of hierarchical structure. A string of words like *the President's book about a goat* turns out to be able to go in most sentences in the same place that a word like *it* can. So, at some level *it* and *the President's book about a goat* are the same kind of thing. On the other hand, *the President's book about a goat* clearly is made up of a bunch of other things. Among them are strings of words that themselves can be replaced by words like *it* or *his* (*The President's book about it*, *His book about it*). We find that there are groups of words in sentences that act together as a unit, which we call *constituents* (parts of a sentence) and that often we can find constituents inside other constituents.

A natural way to view the arrangement of words into hierarchical structures is as a type of (repeated) combination. One word is combined with another to form a combination that can itself be combined with another word (or combination of words), and so forth until the entire structure is built.

Generally, within a *phrase* (a string of words that behaves as a constituent), we find that one word within it seems to serve as its “center,” as the most important word in it, that seems to define the properties of the entire phrase. So, the verb phrase *kicked buckets* is at its core the verb *kick*. It seems to be a relatively stable fact about phrases that they have within them such a “center” from which the properties of the phrase as a whole are drawn. One such property is the type of word that the phrase can substitute in for (that is, the category: noun, verb, etc.). So, a “verb phrase” is a phrase whose center, or “head” is a verb. A verb phrase itself acts essentially as if it were a verb, and so we suppose that the head's properties become the properties of the combination of this head with other material. We say then that the features of the head *project* to become the features of the phrase.

From this comes the idea that combinations occur by taking two constituents (each being either a single word or a combination that has already been created) and creating a combination from them, where one of the two things is the *head* and determines the features of the combination.

### 4 Verbs, events, and $\theta$ -roles

A verb (or predicate more generally) can be thought of as describing a kind of event or state, where different events/states have varying numbers of participants. A verb, by virtue of its meaning, must be marked with some information as to the number of participants its event/state involves. So, *kick* involves two participants, for example, the one doing the kicking, and the one getting kicked. *Put* involves three, the one doing the putting, the one getting put, and the one describing the location where the putting terminates. *Laugh* involves one, the laugher. The participants are said to get *thematic roles*, or technically  $\theta$ -roles. There is one  $\theta$ -role per participant.

There seem to be a relatively small number of  $\theta$ -roles, although exactly how many depends a lot on how much you abstract away from the predicates' own meaning. The most basic set of  $\theta$ -roles appears to include Agent, Experiencer, Theme, Goal, Possessor, and Proposition. We've also posited a Possessee  $\theta$ -role for the special case of the verb *have*.

Of these  $\theta$ -roles, Theme is the most general. It is in a broad sense the participant affected by the event/state. An Agent is generally the instigator of the event/state, and is generally a participant to which a mental state and even intent can be ascribed. An Experiencer is consciously affected by an event/state. Because having conscious mental states seems to be part of being an Experiencer or an Agent, a sentence will generally sound anomalous if a non-volitional participant is assigned one of these  $\theta$ -roles (or be interpreted as if the participant had such mental states, perhaps metaphorically): *The sunlight danced across the surface of the lake*, *The rock kicked John*, *The tree fears lumberjacks*. A Goal generally indicates the termination of a path involved in an event (generally, for events involving motion). A Possessor is a

participant who stands in a possession relation to another participant. Proposition is a role that is assigned to entire clauses, generally something that could in principle be true or false.

The verb/predicate is thought to be basically “in charge” of what  $\theta$ -roles it assigns. Part of our knowledge of any given verb is the knowledge of which  $\theta$ -roles it assigns, which participants it “needs.” As mentioned earlier, this can be formalized in terms of interpretability: A verb is not interpretable until it has been given as many participants as it needs. We formalize this by putting uninterpretable features on the verb that state the requirements, what would be necessary to render the verb interpretable.

We assume that, in principle, a simpler computational procedure is better (and one can make the case for this on basic principles of scientific investigation generally—the simplest solution is the best solution until shown to be inadequate to deal with the complexity of the data, and the simplest solution is also the one that is most vulnerable to being shown to be inadequate, allowing progress in understanding and predictive power: *Ockham’s razor*).

Given the simplicity assumption and the assumption the verbs are uninterpretable until they are provided with their arguments, we hypothesize that what it means to “give” a verb its arguments is to combine the verb with its arguments by means of the procedure (*Merge*) that takes two syntactic objects and forms a composite syntactic object from them, with the features of the composite object being the features of one of the two objects. When a verb that has an uninterpretable feature (indicating that it needs an argument) is Merged with something that constitutes such an argument, the uninterpretable feature is taken care of, and, if that is the last one, the verb becomes interpretable (at least with respect to its argument requirements).

This view entails that the satisfaction of  $\theta$ -role requirements is done *locally*. The arguments of a verb all start off right next to the verb, within the composite object headed by (i.e., that has the features of) the verb. Strictly speaking, it could have been otherwise, but it turns out that as the study of language has progressed over the past 50 years or so, various types of evidence have been unearthed that point to the fact that this locality of  $\theta$ -role assignment is in fact a property of the system.

Another observation that has been made over that time is that the  $\theta$ -roles appear to be assigned to quite consistent positions in the structure. Themes are always in more or less the same place, Agents are always in more or less the same place, etc. Post hoc, we can understand why this should be so: for the purpose of learning the words of a language and their meanings, it is quite advantageous to be able to deduce structural information from the basic meaning of the verb. If a verb has an Agent, we know how that is represented in the structure. If we (or a child) hear a new verb *blick* in the following context *John blicked until he got bored and stopped*, we can deduce that John is a Agent, and that therefore *John* started in the place in the structure where Agents start.

Seeing this tendency, a strong hypothesis was put forward: There is a strict correlation between the  $\theta$ -role that a participant gets and where it is first introduced into the structure of a sentence. This is the Uniformity of Theta Assignment Hypothesis (or UTAH). Discussion of UTAH continues in section 6.

## 5 Case and structural position

Another property that one can observe in many languages is that noun phrases (the type of phrase that can generally be referential and/or serve as participants in events/states) are marked morphologically, depending on where in the sentence they are. In English, this is evident in the pronoun system: *I saw John. John saw me. My brother saw John*. The varying form a noun phrase takes based on its position is case marking.

An important fact about case marking is that it is not correlated with the  $\theta$ -role that an argument gets. So, the nominative (subject) case can appear on Agents (*I laughed*), or Themes (*I fell, I was kicked*). Likewise, for participants inside a noun phrase, the genitive case can appear on Agents (*my closing of the*

door), Themes (*my destruction*), or Possessors (*my house*). There seems to be a designated position for each of these cases, and that position seems to relate to where the noun phrase appears in the structure (and not where the noun phrase may have started—that is, not which  $\theta$ -role the noun phrase gets).

Combining the idea of the UTAH with the observation that case tells us where in a sentence a noun phrase is, we are already led to the idea that a participant can occupy more than one position in a syntactic structure. The subject position, which receives nominative case, is the one that precedes auxiliaries and modals: *I will laugh, I will fall, I will be kicked*. This appears to be a constant position in the structure of a sentence. Yet because the  $\theta$ -role assigned to *I* in each instance is different, it must also be associated with the position within the verb phrase designated by the UTAH. This association with two positions, one lower in the structure, and one higher in the structure, is thought of as a kind of “movement,” where a participant is first combined with the verb within the verb phrase, and then moved to a higher position in the structure that corresponds to the case it receives.

There are other relatively clear cases of movement as well. One example is the passive, in which an underlying Theme, which surfaces in the active in a lower position appears in the passive in the higher subject position: *I have seen him. He was seen*. We think of the 3rd person pronoun as moving from the position it holds in the first sentence into the position it holds in the second sentence when the verb is passive.<sup>1</sup>

## 6 Ditransitive verbs and UTAH

It seems that the largest number of  $\theta$ -roles any verb can require is three, an example of such a verb being *put*.<sup>2</sup> In order for the UTAH to be able to designate these positions in an unambiguous way, we need to posit a two-level verb phrase. In a sense, verb phrases are always made of two distinct phrases, which we have called the phrase headed by “Big V” (the lower of the two phrases) and “little *v*” (the higher of the two phrases). We can also tell by the word order in sentence like *Mary gave a book to John* that the verb is pronounced in the position of the “little *v*,” taken to be an indication that big V has moved to little *v*.

In the interests of simplicity (in the form of structural uniformity), we assume that verbs always (whether ditransitive or not) have this two-stage structure (both a little *v* and a big V, with little *v* above big V in the structure), and that big V always moves to little *v*. With this much in place, we can now state the UTAH more precisely. The full set we will use is listed below, though some aspects of it have yet to be discussed.

Uniformity of Theta Assignment Hypothesis (UTAH)	
AGENT	DP daughter of $v_{agent}$ P, DP daughter of $n_{agent}$ P
EXPERIENCER	DP daughter of $v_{exp}$ P, DP daughter of $n_{exp}$ P
THEME	DP daughter of VP, DP daughter of NP
GOAL	PP sister of V, PP sister of N
PROPOSITION	CP/TP sister of V
POSSESSEE	$n$ P sister of Poss
POSSESSOR	DP daughter of PossP
EXPERIENCER	PP daughter of VP

Looking at the statement above, it hasn’t quite managed to embody the goal state of having an unambiguous one-to-one relation between structural position and  $\theta$ -role. Looking at the first two lines, AGENT

<sup>1</sup>Another example is *wh*-movement, covered later, where an interrogative phrase appears at the beginning of a sentence, before even the subject: *Who have I seen?*

<sup>2</sup>Probably. Arguably, *bet* has four, and someday it would be nice to figure out how that could be. I bet you \$10 that it will not be straightforward.

and EXPERIENCER each have two environments listed, one within  $vP$  and another within  $nP$ . We might be able to explain this away by supposing that  $n$  and  $v$  themselves have something in common (like, being a category-former) and that when viewed this way, the agent might really be just the DP daughter of the maximal projection of any category-former of the AGENT-introducing type. A second thing to note there is that “DP daughter of  $vP$ ” is not by itself sufficient to determine the  $\theta$ -role of the DP—it matters what “flavor” of  $v$  we are talking about. We take the category of  $v$  to be capable of introducing an argument, but the  $\theta$ -role that argument gets (AGENT, or EXPERIENCER) depends on which kind of  $v$  is involved. This can be considered to be at least somewhat analogous to types of T (could be past, nonpast, infinitive, etc.) or of C (could be interrogative, declarative, exclamative, etc.). A possibly more thorny issue involves where you can find EXPERIENCERS. It seems that there are at least two positions where an EXPERIENCER might appear, one for EXPERIENCERS like *(to) Chris* in *Pat seems to Chris to be drunk*, and the other for EXPERIENCERS like *Chris* in *Chris fears spiders*. Yet a third (so far unpredicted) place where EXPERIENCERS seem to be found is where *Chris* is in *Spiders scare Chris*. So, there are some open questions about exactly how EXPERIENCERS work, although most of the time they act syntactically the same way that AGENTS act. And therefore perhaps some open questions about the validity of the broadest form of the UTAH. But it seems—if not right—at least *close* to right.

<p>Transitive verb (with THEME) <i>kick, eat, see</i></p> <pre>       vP      /  \     DP   v'   AGENT /  \        V+v  VP            /  \           &lt;V&gt;  DP                 THEME           </pre>	<p>Transitive verb (with PROPOSITION) <i>say, hear, claim</i></p> <pre>       vP      /  \     DP   v'   AGENT /  \        V+v  VP            /  \           &lt;V&gt;  CP/TP                 PROPOSITION           </pre>
<p>Unergative verb <i>dance, jog, shout</i></p> <pre>       vP      /  \     DP   v'   AGENT /  \        V+v  &lt;VP&gt;           </pre>	<p>Unaccusative verb <i>fall, sink, melt</i></p> <pre>       vP      /  \     V+v  VP         /  \        &lt;V&gt;  DP             THEME           </pre>
<p>Ditransitive verb (with GOAL) <i>put, send, give</i></p> <pre>       vP      /  \     DP   v'   AGENT /  \        V+v  VP            /  \           DP   V'         THEME /  \               &lt;V&gt;  PP                   GOAL           </pre>	<p>Ditransitive verb (with PROPOSITION) <i>tell, persuade, promise</i></p> <pre>       vP      /  \     DP   v'   AGENT /  \        V+v  VP            /  \           DP   V'         THEME /  \               &lt;V&gt;  CP/TP                   PROPOSITION           </pre>
<p>Raising verb <i>seem, appear</i></p> <pre>       vP      /  \     V+v  VP         /  \        &lt;V&gt;  CP/TP             PROPOSITION           </pre>	<p>Raising verb (with EXPERIENCER) <i>seem (to Pat)</i></p> <pre>       vP      /  \     V+v  VP         /  \        PP   V'     EXPERIENCER /  \                 &lt;V&gt;  CP/TP                     PROPOSITION           </pre>

Thanks to UTAH, we know, for every argument, where it started out. That's important, let me say that again. **Thanks to UTAH, we know, for every argument, where it started out.** So, one of the most important things you can do when trying to determine the structure underlying a sentence you are trying to analyze is work out what the Agents, Themes, Goals, etc., of the verbs involved seem to be.

## 7 Merge, Move, Adjoin

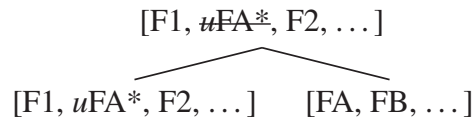
The engine of our model of syntax are the operations that put things together to form a larger syntactic object (*Merge* and *Adjoin*), and the way that these contribute to the elimination of the uninterpretable features of the structure being built (*Agree*).

We can think of this process metaphorically as being a sort of workbench, where all of the lexical items that are going to be used to form a sentence are scattered about. One by one, we pick them up and combine them with another lexical item (or some already-built object) to make a new object (which we then put back on the workbench).

When we put two things together (*Merge*), we create a combination that itself has certain properties. The properties that the combined object has seems to be essentially the properties of one of the two things that were combined. So, of two things that are put together by *Merge*, one of them “projects” its properties to the combined object.

The other thing we assume is that we don’t gratuitously or randomly put things together (at least not with *Merge*). We suppose that we put things together when we *need* to put them together. One of the reasons that we might need to put two objects together is if one object is uninterpretable without the other. (Another reason we might do this is following the Hierarchy of Projections, which will be explored a bit more in the next section).

The determination of which object projects its features to the combination of the two objects seems to be determined by which of the two objects “motivated” the *Merge*—which of the two objects had a need for the other one. This can be shown very abstractly as follows:



Here, an object that needed the feature [FA] was combined with something that had the feature [FA], and the features of the combined object are the features from the object that needed the feature [FA].

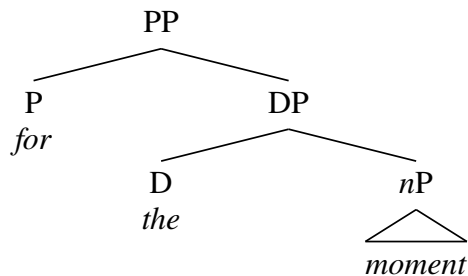
We have adopted some notational conventions and terminology for drawing these (recursively) combined objects into tree structures, based partly on traditional ways of labeling these objects.

A **head** is generally a single lexical item, taken from the lexicon, the “atoms” of the structure, not made up of anything else. A standard way to notate a head in a tree is by writing its category feature (N, V, Adj, P, C, T) and below that its phonological realization and/or some other means of uniquely identifying which lexical item it is, and sometimes below that (some of the) additional features that the lexical item has:

C	P	D
<i>that</i>	<i>for</i>	<i>the</i>
[clause-type:decl]	[ $\mu D^*$ , $\mu case:acc$ ]	[ $\mu case:$ ]

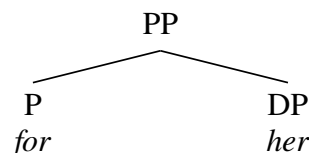
When we combine two objects, one of the two projects its features to the combined object. The one that *didn’t* project its features is—by virtue of not projecting its features—a **maximal projection**. Its features have reached the end of the road. Since its features will not be the features of the combined object, its features will project no further. A maximal projection is often labeled with a P after its category feature (traditionally, the “P” stands for “phrase”). So, below, the features of the D characterize the combination of the D and the  $\mu P$ , hence the category of the constituent *the moment* is D. Because it is combined with P and its features do not project further (the whole combined object is of category P), the constituent *the moment* has projected its features as far as it will, and is therefore a maximal projection, written **DP**.



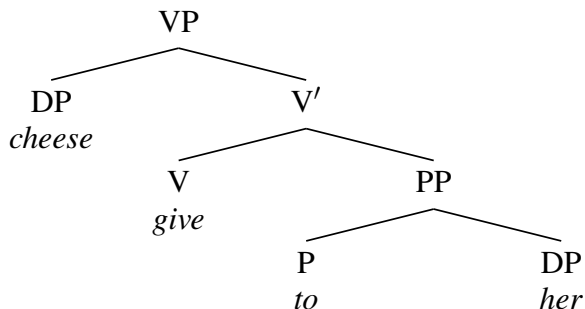


The constituent that is a sister to a head like this is usually called the **complement** of the head, and is—by necessity—a maximal projection. It is standard practice to *always* write the complement as a P, even in light of the comment about heads and maximal projections made immediately below.

Because the determination of whether something is a head and the determination of whether something is a maximal projection are made on the basis of different considerations (something is a head if it comes out of the lexicon as an atom, and is not a combination of any smaller units; something is a maximal projection if its features do not project to the object of which it is a part), it is perfectly possible for a single item to be both a head and a maximal projection. A perfect example of this is the English pronoun. The pronoun *her* has category D, but it also has no needs, it is interpretable (with respect to things it might need to Merge with) as it stands. So, it will not motivate Merge and will not project.



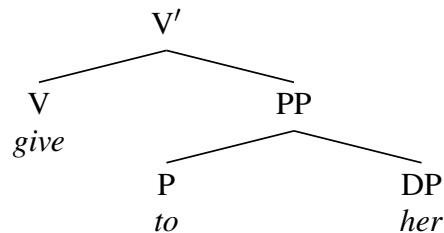
In the tree above, *her* is a head and also a maximal projection. In such situations, it is not clear whether it should be labeled as a head (labeled “D”) or as a maximal projection (labeled “DP”). By convention, it is always labeled as DP, but take note that this is a relatively arbitrary decision and plays no role in the function or predictions of the system. Because a head’s features can project more than once (a head can have more than one need that must be satisfied by Merge), the situation can arise where a constituent is neither a head (it was created by Merge, its features were projected from one of its component objects) nor a maximal projection (because it is subsequently Merged with another object, but its own features project). In such a case, the label is written with a “bar” (which was, in the early days, an actual horizontal bar written over the category label (e.g.,  $\bar{V}$ ), although these days is usually written with a “prime” symbol:  $V'$ ).



Here, V combines with *to her*, which had category P, and the features of the head V project. The complement’s features do not project, so it is a maximal projection of category P, a PP. The constituent *give to her* is then combined with the pronoun *it*, still motivated by a need of the verb, and the features of

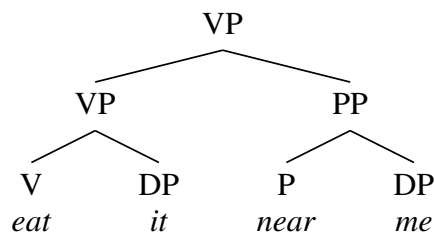
the verb project again to the entire constituent. Thus the whole constituent is the maximal projection (VP), and the *give to her* constituent is neither a head nor a maximal projection (V').

The decision about whether to label something as a maximal projection can also be predicted in part by the features it has. If an object has a strong uninterpretable feature (e.g., [*uD\**]) that has not yet been checked, then the object is not a maximal projection—another Merge will be necessary, and because this object is motivating the Merge, it's going to project its features to the resulting object. For this reason, when you draw a partial tree, you might still draw an **intermediate projection** (an X') label even when the tree has not (yet) been extended further:



Merge is a way to combine two objects where the combination is forced (either by one of the objects or by the Hierarchy of Projections). There is a second way to combine two objects, which is not forced: **Adjoin**. In general, we assume that all of the forced operations are carried out first, and only then are the unforced operations allowed. Therefore, adjunction only occurs to maximal projections (although I will say something about head-movement and head-adjunction in a moment—it turns out to be a somewhat different case).

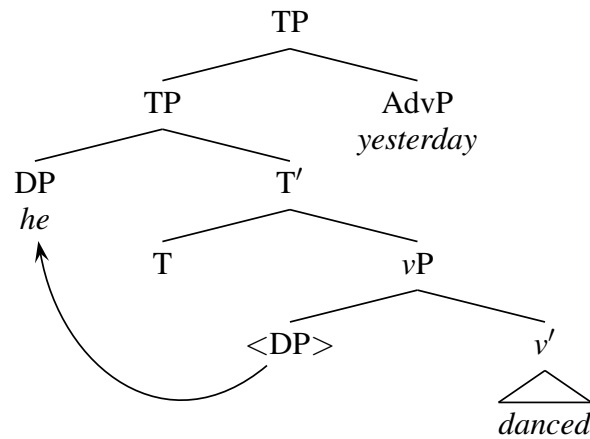
Adjunction is an optional and basically invisible operation as far as the features are concerned. A *vP* is no different whether a PP is adjoined to it or not, a TP with and without an adjunct is exactly the same. As a way to recognize this “invisibility” of adjunction, the standard way of writing adjunction is by simply repeating the label of the object adjoined to as the label of the combined object.<sup>3</sup>



Here, a PP has been combined with a VP, where the VP had already been “finished” and was not going to project its features any further.

Something that is adjoined cannot have any needs that it satisfies by adjoining (at least, not a need of the sort that would have motivated Merge), so something that is adjoined is necessarily a maximal projection—it is not going to project its features further. Again, there is an issue of what label to use when adjoining a head that also happens to be a maximal projection (such as an adverb). By convention, just as with complements discussed above, it is written as being a maximal projection, and—again—this is an arbitrary decision about labeling that doesn’t affect the functioning or predictions of the system.

<sup>3</sup>There is at least one piece of evidence for a certain “reality” to this labeling convention for adjunction: The phenomenon of *one*-replacement in English. The word *one* appears to be able to replace any NP in a structure (although, at least in many people’s judgments, it cannot substitute for an N that is not an NP). An NP with a PP adjoined to it can be replaced entirely with the word *one*, and the “inner” NP (to which the PP was adjoined) can also be replaced by *one*, leaving the PP adjoined to *one*. That is: *The book of poetry with the red cover on the shelf* can be reduced as *The one on the shelf* or *The one with the red cover on the shelf*.



Alongside Merge and Adjoin, it is also often necessary to **Move** something from one part of the structure to another. This is understood as in fact being a process of making a **Copy** of the thing that is to “move” and then incorporating it into the structure a second time, for example Merging it again.<sup>4</sup> When we make a copy of a maximal projection and re-Merge it higher in the structure, the copy is generally labeled by writing the label of the maximal projection that was moved within angled brackets. This was also shown above. It is also very useful (particularly in complicated structures) to **draw an arrow** that indicates what has moved and from where.

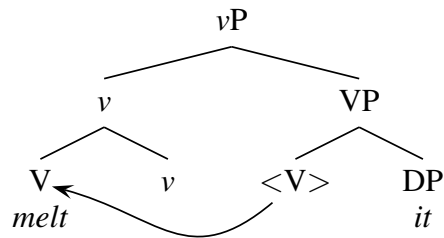
You may also notice (above) that, again, the convention of labeling as “XP” objects that are both heads and maximal projections is used with respect to the pronoun *he* in the specifier of TP. The logic here is the same—anything that goes there would necessarily be a maximal projection (and of course it’s possible to have things there that aren’t heads, e.g., *the children*), so for consistency anything in a specifier (along with anything in a complement, and anything adjoined to a maximal projection) is labeled as a maximal projection even if it also happens to be a head.

The other case of movement that we’ve seen is **head-movement** where a head moves up to another head. We understand this as a process that creates a “complex head” by in a certain sense fusing the two heads. This is drawn in much the same way as an adjunction structure is drawn, but that obscures a couple of points of difference, so it is worth spending a moment on the characteristics of head-movement.

When a head moves up to another head, e.g., V moves to *v*, this doesn’t change the fact that the higher head is still the lexical item whose features have projected in the constituent it heads. So, below, the features of *v* are still the features of *v'*, regardless of whether V has moved up to combine with *v* or not. Therefore, the label of the complex head remains *v*. It is for this reason that the combination of two heads via head-movement is thought of as a kind of adjunction: the moving head is attached to the higher head in a way that doesn’t alter the essential properties of the higher head.

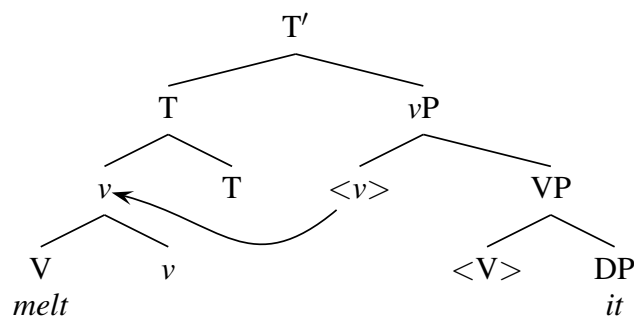
On the other hand, head-movement, *unlike* XP-adjunction discussed above, *is* motivated by a need of the higher head, in much the same way as Merge is motivated. The higher head is only interpretable when the lower head has moved up to it—that is, an uninterpretable feature is checked by head-movement.

<sup>4</sup>An alternative, and possibly better, way to look at this is that—rather than making a copy and Merging the copy—we just pick some object inside and Merge that same object again. The result of this in a tree would be that this object has two different mother nodes. This makes for weirder-looking tree structures, but it resolves some questions about what happens to uninterpretable features when you move something for the purpose of checking an uninterpretable feature on the moving object. And it also means that we wouldn’t need to posit a “copy” operation. We probably will not pursue this any further here, but it might be worth keeping in the back of our minds when doing future analyses.



By convention (and in fact in the past people have also proposed this as a substantive aspect of the theory of head-movement), the moving head is attached to the left of the higher head.

Once combined in this way, the complex head is relatively “opaque”—it is still seen by the syntax as an “atom” of sorts, and so if a yet higher head requires a lower complex head to move to it, the entire complex head moves together to form an even-more-complex head higher up. (Note that the structure below would not actually occur in English, because *v* doesn’t move to T, but it would in the French version of the sentence *it melted*.)



It’s worth mentioning that head movement doesn’t fit all that well in our system, insofar as it is not simply assembling things together but actually takes an already-formed object and changes something inside of it (replacing one member of a Merged pair with something else). There are a couple of ways one could go with this. One radical possibility is that we might want to hypothesize that there is no head movement at all, and try to explain the effects that head movement seems to have in some other way. Since a lot of what has been taken to be evidence of head movement concerns the order in which things are pronounced, or the degree to which several heads seem to be pronounced together, we might try to enhance the morphological or pronunciation system to yield those effects rather than moving things around in the syntax. Certainly some, including Chomsky himself, have suggested this. Another possibility might be to reimagine what is happening in head movement; one fairly elegant version of such an account is discussed by Matushansky (2006, *Linguistic Inquiry*). In Matushansky’s account, heads move into regular specifier positions, but then get morphologically readjusted later to yield the structures we’re familiar with (instead of creating those structures directly through movement). In the interest of making progress, we will just continue to suppose that head-movement works as outlined above: Moving a head X to the position of another head Y entails replacing Y in the structure with a Y that has X adjoined to it. But there are certainly places here where different decisions could be made.

That pretty much covers the types of movement and tree-drawing conventions in our system, so next we can turn to a discussion of motivations for movement and the mechanics of feature checking.

## 8 Motivations, Agree, feature checking, and the Hierarchy of Projections

The way that we express the “needs” of a lexical item is by designating the lexical item as uninterpretable unless those needs are met. In the cases already mentioned above, these are situations where, e.g., a preposition needs an object, so we designate P as having an uninterpretable feature [ $uD^*$ ], which will be eliminated when an object of category D is Merged with it.

As the system has developed, we have discovered many uses for uninterpretable features in explaining the phenomena in syntax, over and above these “selectional” features.

One could accurately say that the engine that makes this system run is basically the existence of uninterpretable features and the specification of how such features can be eliminated. Through the semester, the notion of feature checking (the elimination of uninterpretable features) has been refined in various ways: First, we considered only checking under sisterhood (essentially, checking by Merge), but further exploration suggested that “long-distance” feature checking is possible, and so “local” feature checking needed to be distinguished from “long-distance” feature checking.

The elimination of uninterpretable features (feature checking) is considered to be a two-step process. First, two features that could participate in checking are identified (and valued, where a value is needed) by the operation **Agree**. Second, if the feature is far away but yet can only be checked “locally,” a movement is effected that would bring the two features close enough together to allow the uninterpretable one(s) to be checked.

Agree can relate any two features that are in a c-command relationship (at least up to phases, covered later). If an uninterpretable feature c-commands a feature that “matches” it (that could check the uninterpretable feature), then those features can be related by Agree. Similarly, if a matching feature c-commands an uninterpretable feature that it matches, those features can be related by Agree. We can informally think of the c-command relation as a requirement for the features to “see each other.”

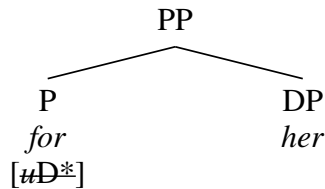
There are two different dimensions on which uninterpretable features can differ from one another. One dimension is the valued–unvalued dimension. An example of a valued uninterpretable feature is the [ $uD^*$ ] feature on P: this expresses a need for a very specific kind of matching feature, a category feature [D]. An example of an unvalued uninterpretable feature is the [ $u\phi$ :] feature of T—this expresses a need for some  $\phi$ -features, whatever they may be.

The other dimension is the strong–weak dimension. A weak uninterpretable feature can be checked “at a distance.” So long as it “can see” a matching feature, Agree can relate them and the uninterpretable feature is eliminated/checked. A strong uninterpretable feature requires locality—not only must the two features related by Agree be able to “see each other” but they must also be able to “touch each other.”

The practical consequence of strong uninterpretable features is that they motivate movement (unless they can be satisfied directly by Merge). Weak uninterpretable features don’t motivate movement. We indicate strong uninterpretable features with a “\*”.

As we find in other parts of the system, Agree is essentially “lazy” and will only relate a feature with another if there was not a closer one. Thus, the [ $u\phi$ :] feature of T will Agree with the  $\phi$ -features of the Agent and not the Theme in a sentence with a transitive verb (such as *The rhinoceros is extinguishing fires*—it isn’t *\*The rhinoceros are extinguishing fires*).

There are few different “reactions” that we find to a strong feature. Which one is chosen depends partly on the interaction with the Hierarchy of Projections, discussed shortly. If a head has a strong category feature, Merge is one response:



The sisterhood relation between P and DP here is close enough to allow for the strong  $[uD^*]$  feature of P to be checked. This kind of response to a strong feature is the one that is most closely tied to the assignment of  $\theta$ -roles.

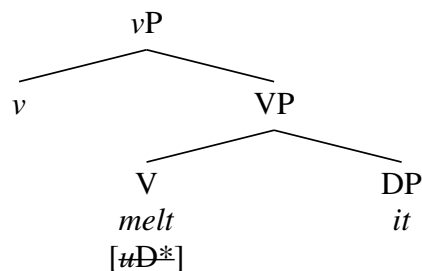
When the Hierarchy of Projections is involved, the Hierarchy of Projections “takes priority.” It’s not exactly clear why it does, but it nevertheless does. So, let’s pause for a moment to review the Hierarchy of Projections.

The **Hierarchy of Projections** (“HoP”) is a stipulation that we hope ultimately to be able to derive from the properties of the interpretive component of grammar. It is basically a specification of the order in which lexical heads are Merged into the tree.

Hierarchy of Projections
C > T > Neg > M > Perf > Prog > Pass > $v$ > V
D > Poss > $n$ > N

The Hierarchy of Projections tells us the order in which things occur in the spine of the tree (or the spine of a DP). These projections are in general optional—if there is no Prog, it’s not a problem. But if there is a Prog, it will occur lower than Perf (if there is one), higher than Pass (if there is one) and  $v$ , etc. A couple of things do not seem to be optional, like  $v$  and  $n$ . When there is a V, there will also be a  $v$ , and when there is an N, there will also be an  $n$  and a D. The hypothesis is that something about the semantic interpretation is responsible for that. For  $n$  and  $v$ , this might be a consequence of these being the “category makers”—many people have supposed that N and V are not actually distinct categories but are both essentially “roots,” which would be nouns if the complement of  $n$  and verbs if the complement of  $v$ . If that is true, then the semantic system needs  $n$  or  $v$  in order to know what category the thing is. Similarly, D might be required to establish reference. If semantics can filter out structures that lack D,  $v$ ,  $n$ , then we do not need to encode their non-optionality as part of the syntactic system.

The Hierarchy of Projections is a second way to motivate Merge (the other being the checking of a [strong] uninterpretable category feature). The way this works is that once a constituent is “finished” (there are no strong uninterpretable category features to check, and whatever adjunction there is to be done has been done), the next step is to Merge the finished constituent with the next head up on the HoP.



The motivation for Merging  $v$  with VP above is the Hierarchy of Projections. VP had no further uninterpretable category features to check and was therefore “finished.”

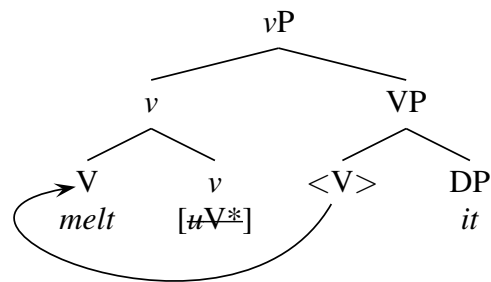
One peculiar property of Merging in order to satisfy the HoP is that it creates a configuration in which Agree cannot apply between the two things that Merge. In the tree above,  $v$  has a  $[uV^*]$  feature (because

$v$  always does, this motivates the head-movement that we turn to next), but Merging the VP (which has a V category feature) with  $v$  does not succeed in checking the  $[uV^*]$  feature of  $v$ , despite the fact that  $v$  and VP are sisters.

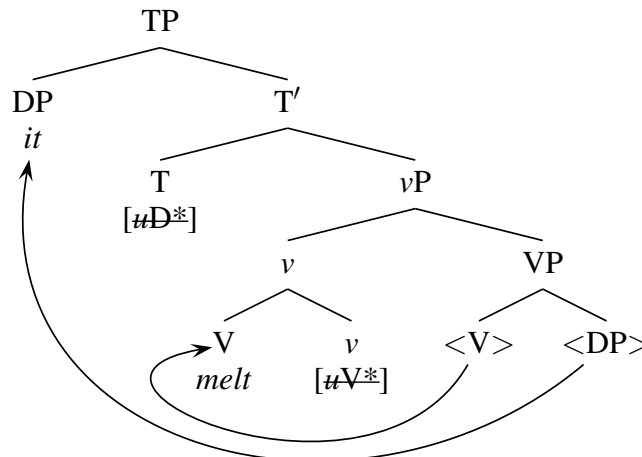
So: Even if there is another reason to Merge, the HoP takes priority, and when Merge happens due to the HoP, the two things Merged are not eligible to check each other's features. (Note, however, that features inside the two things Merged are still eligible for checking. So, T is Merged into the structure as a response to the HoP, but when that happens, it can then see  $\phi$ -features within the  $v$ P. It just wouldn't be possible for T to check a  $[uv^*]$  feature).

Returning now to how the system responds to strong uninterpretable category features, we saw that one response was to simply Merge a constituent with the right category (e.g., satisfying the  $[uD^*]$  feature of P). A second way that such a feature can be checked is by head-movement of a lower head of the right category up to the head that has the strong uninterpretable feature (e.g., checking the  $[uV^*]$  feature of  $v$ ).

Generally, the head-movement option is taken when the sister of the head is of the right category to do the checking, but is prevented from checking due to having been Merged for the HoP. So, as just discussed,  $v$  and VP are Merged in order to satisfy the HoP and so the  $[V]$  feature of VP cannot check the  $[uV^*]$  feature of  $v$ . In that situation, the lower head moves up to the higher head (which puts the features close enough that the strong feature can be checked):<sup>5</sup>



The final possible response to a strong uninterpretable category feature is movement of a constituent of the appropriate category into the specifier. This option is generally taken when DPs and PPs are involved, and occurs when a potential checking configuration wasn't blocked by a satisfaction of the HoP. An example of this is the movement of the subject into the specifier of TP.



*There are a couple of things in the following paragraph that anticipate things that will be covered post-midterm.*

<sup>5</sup>It is possible that something like the Head Movement Constraint could be derived from this, if the decision about whether to move a head or a whole XP is determined by whether the XP is the sister of the motivating head or not. Maybe also worth noting here that for our interpretation of negative sentences, we need to assume that the HMC does not hold, since we will have auxiliaries moving up past NegP to reach T.

There is a special case that can lead to head-movement, which is where an unvalued uninterpretable feature is “valued as strong.” These are the cases where auxiliaries move to T, or when T moves to C. When an interrogative C is Merged with a TP (in a main clause, in English), the [clause-type:Q] feature of C values the [*u*clause-type:] feature of T as strong ([*u*clause-type:Q\*]). That is, it becomes strong as a *result* of Agree. The now-strong feature of T cannot be checked in place, not so much because TP and C are not close enough, but rather because it only became strong “too late” (it wasn’t strong initially, and it wouldn’t have been strong if C had valued it as [*u*clause-type:decl]). In that case, the head moves up: T moves to C.<sup>6</sup>

## 9 Functional structure

It is possible to divide the heads in a sentence’s structure into “lexical” and “functional” heads (and we have many more functional heads than lexical heads). The lexical heads are those that contribute a kind of substantive meaning (nouns, verbs, adjectives, adverbs, possibly prepositions), and the functional heads are the others. The functional heads in a structure each has an important role to play.

- C is considered to be where the clause type is determined—it distinguishes questions from statements.
- T is where tense information is, determining whether a clause is infinitive, past, or present. T also plays a kind of facilitating role in agreement between the subject and the highest verbal element. Finally, T has the “EPP” property (encoded as a [*u*D\*] feature on T) that forces all sentences to have a subject.
- *v* is the means for assigning Agent and Experiencer  $\theta$ -roles, and—given that it is required for every verb—may also be what determines that a verb *is* a verb. There have been proposals, for example, that suggest that the “big V” isn’t really a verb at all, but just a lexical root that becomes a verb by virtue of having a *v* above it. This would take us partway toward understanding the relationship between *dance* the verb and *dance* the noun, for example.
- *n* is the counterpart of *v* except for nouns, also responsible for assigning Agent and Experiencer  $\theta$ -roles.
- D is the uppermost category of the “noun phrase” and is the category that requires case and is often considered to be the interpretive location for concepts like definiteness.
- Poss is a category whose sole purpose is to introduce a Possessor (and assign the  $\theta$ -role) within a DP.
- Perf marks perfective aspect (“completed”).
- Prog marks progressive aspect (“ongoing”).

---

<sup>6</sup>It is reasonable to consider this “valuing as strong” mechanism to be kind of... clunky. But the reality is that we have an observed pattern we need to account for, and this explanation at least has the right “shape.” The relationship between T and the closest auxiliary is special, and can motivate a movement even when other similar relationships don’t. There’s something special about the agreement relationship between T and the head it values inflectional features for, such that it can only be checked locally and motivates a movement. Since strong features are how we represent a requirement of local checking, we need these features to be strong but only in this one environment. It might be able to be stated in a different way, but the complexity will surely be equivalent, since at the end of the day, it is the pattern in the data that is complex.



- Pass forms passive verbs, which it does by imposing a requirement on its  $vP$  complement that the  $v$  not assign any  $\theta$ -roles (Agent or Experiencer) nor check accusative case (essentially, requiring that its complement  $vP$  be the same  $vP$  that appears in unaccusative structures).
- M are modals, including *shall*, *should*, *will*, *would*, *can*, *could*. The infinitive marker in English (*to*) is also considered to be a modal.

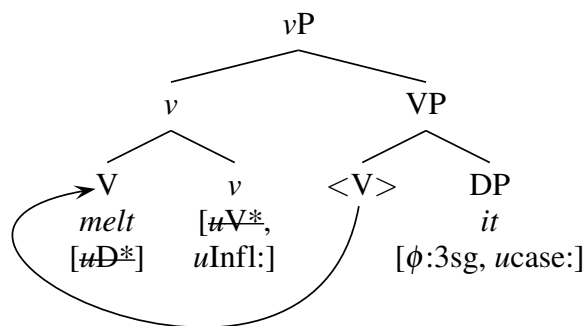
## 10 Agreement, inflection, and the auxiliary system

It is probably worth taking a special look at the auxiliary system in English and the way in which subject-verb agreement and other inflection is assigned.

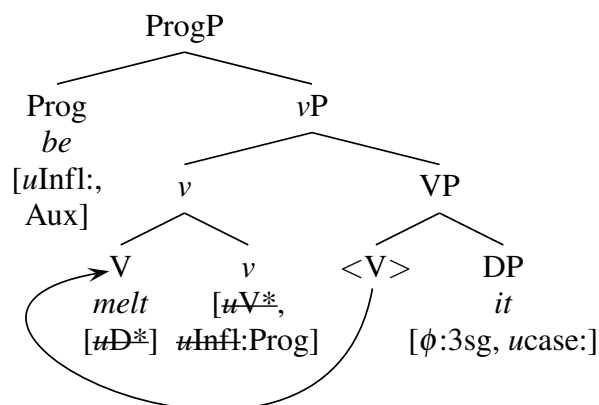
There is a set of things that take inflection, these are things that (at least abstractly) can be inflected. All of the auxiliaries (Perfective *have*, Progressive *be*, Passive *be*, Modals) and the verb ( $v$ ) need to be inflected, which is reflected by having a  $[uInfl:]$  feature.

A  $[uInfl:]$  feature is an unvalued, uninterpretable feature that must get a value from something else. Because we build the trees up from the bottom, we generally introduce something with a  $[uInfl:]$  feature first, and then, later, add something to the tree that will value the  $[uInfl:]$  feature. That is,  $[uInfl:]$  features are valued by the next thing up the tree that is capable of assigning inflection.

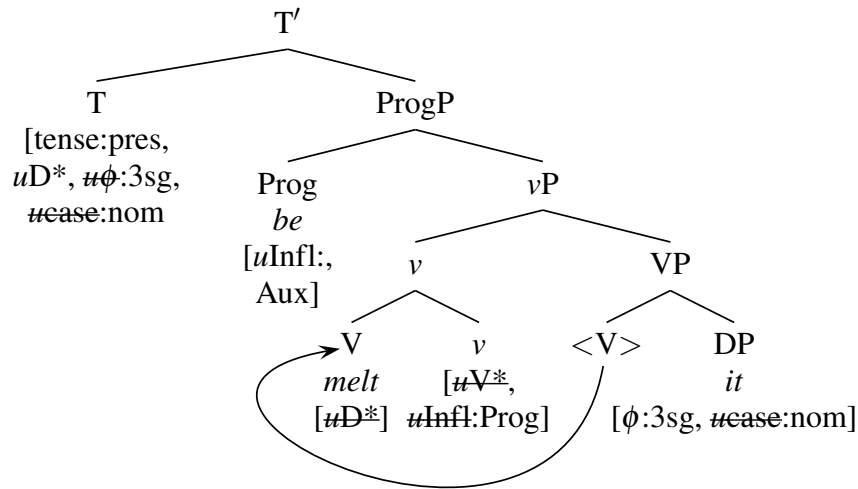
The things that are capable of assigning inflection are T, M, Perf, Prog, and Pass. T is a somewhat special case, so we'll come back to that. For the others, at the point they are Merged into the tree, they can value (and check) a  $[uInfl:]$  feature they c-command. For a simple example with a progressive auxiliary (*It is melting*), we first have the following structure (where I am now explicitly indicating the  $[uInfl:]$  feature of  $v$  as well as the  $\phi$ -features and case feature of the DP and the selectional feature of the V *melt*):



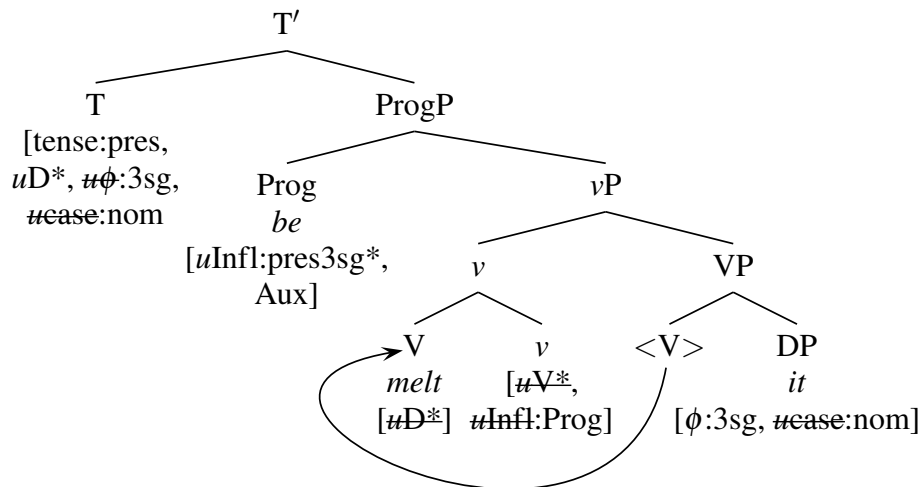
Assuming that there is a Prog in the numeration (that is, “on the workbench”), the HoP would have us add that next, resulting in the structure below. Once Prog has been Merged with the  $vP$ , Prog (or, more specifically, its category feature) will c-command the  $[uInfl:]$  feature of  $v$  and will be able to value the  $[uInfl:]$  (via Agree). Prog itself needs inflection, so it has a  $[uInfl:]$  feature of its own, but we have now taken care of the inflection on  $v$  (which ultimately will cause the verb to be pronounced as *melting*).



Suppose now that we have reached the point where we merge T with the ProgP. T has a [ $u\phi$ :] feature, and is capable of assigning inflection. When T is Merged, it must first take care of its own [ $u\phi$ :] feature. T will c-command the DP *it*, which has  $\phi$ -features, and so the [ $\phi$ :3sg] feature of *it* can value and check the [ $u\phi$ :] feature of T. At the same time, the [ $u\text{case}$ :] feature of the *it* can be checked by the [ $u\text{case}$ :nom] feature of T. At this moment in the derivation, the structure looks as below.



Then, T (now with  $\phi$ -features valued) can value the [ $u\text{Infl}$ :] feature of Prog, which it does with both its tense value and its  $\phi$ -features. Furthermore, because Prog is an auxiliary (it has an [Aux] feature, as does Perf, Pass, and M), T values the [ $u\text{Infl}$ :] “as strong”: [ $u\text{Infl}$ :pres3sg\*].



The next step will be to move Prog to T (by head-movement) in order to check the [ $u\text{Infl}$ :pres3sg\*] feature of Prog (which, because it is strong, has to be checked very close to T—i.e. with the Prog head-adjoined to T). This happens in the same way that V moves to  $v$ , and then the derivation continues, moving *it* into the specifier of TP, etc.

There are a couple of comments to make about the foregoing discussion of inflection. First: we are supposing that a [ $u\text{Infl}$ :] feature can be valued and checked by any of the following category features: Pass, Perf, Prog, and M. The morphology will interpret a head with, for example, a [ $u\text{Infl}$ :Perf] feature on it as taking the *-en* form (like *eaten*, *been*, *had*). This would seem to be the simplest system capable of capturing the inflection facts we see in the English auxiliary system, though we might need to make it more complex for other languages potentially. Second: When a [ $u\text{Infl}$ :] feature is valued by T, there are a couple of differences. It is not sufficient (to get the morphology right) to value with just the category feature—knowing that a verb’s inflectional ending was determined by T doesn’t differentiate *eat* from *ate*,

so we need to know what *kind* of tense it was, not just that it was tense. But in fact, even just knowing whether the tense was present, past, or infinitive is not enough either, since (at least in the present) we need to be able to distinguish *eat* from *eats* as well. To account for this, we are making the (essentially stipulative) assumption that T is capable of valuing a [*uInfl*:] feature only once it has both tense and  $\phi$ -features to contribute. And since T doesn't have  $\phi$ -features inherently (it needs to collect  $\phi$ -features from something else), the valuing of T's  $\phi$ -features must precede T's valuing of a lower [*uInfl*:] feature.

## 11 DPs and Case

Every D starts off with a [*ucase*:] feature—in order to be interpretable, it must have this feature valued and checked. The function of case in a sentence can be thought of as being an indication of where in the structure the DP is—based on what values its [*ucase*:] feature.

Nominative case (*I, he, she, they*) goes to DPs that wind up in the specifier of a (finite) T. T has a [*ucase:nom*] feature that will match, value, and check on the closest DP to T (which would be the Agent, if there is an Agent). We can think of [*ucase:nom*] as being a feature that means “I need to assign nominative case to something.”

Genitive case (*my, his, her, their, the dog's*) goes to DPs that wind up in the specifier of the silent D that assigns genitive case ( $\emptyset_{GEN}$ ). The  $\emptyset_{GEN}$  D has a [*ucase:gen*] feature that will match, value, and check a DP within the complement of  $\emptyset_{GEN}$ . Note that, since  $\emptyset_{GEN}$  is a D itself, it also has an unvalued [*ucase*:] feature that needs to get valued and checked elsewhere. So, in *My dog left*, the DP *my* has its case checked (genitive) by the  $\emptyset_{GEN}$  that heads the DP *my dog*, and the DP *my dog* has its case checked (nominative) by the finite T.

Accusative case (*me, him, her, them*) goes to DPs that generally stay in place. Prepositions and *v* can have [*ucase:acc*] features that can value such DPs. Whether *v* has a [*ucase:acc*] feature depends on whether it assigns a  $\theta$ -role to a DP in its specifier (“Burzio’s Generalization”); the *v* in a passive, and the *v* for an unaccusative, do not have this [*ucase:acc*] feature and cannot check accusative case.

*Of*-case is a special case that goes to DPs that are Themes within an *nP*. It is the direct analog of accusative case inside *vPs*. The *n* has a [*ucase:of*]. The *of*-case is pronounced on a DP by a prefixed *of*, but this *of* is not part of the syntactic representation, it is just how you pronounce a DP with *of*-case. This means that in *John's gift of cheese to Mary*, *of cheese* is a DP (not a PP).

There is also one special instance in which accusative case can be assigned by a V, which arises for Vs that have a possessive *have*-type meaning. This is specifically required for double object constructions: *John gave me a book* (meaning ‘John caused me to have a book’). Here, *me* gets accusative case from *v*, and *a book* gets accusative case from V.<sup>7</sup>

*Later, we will also discuss null case, in connection with PRO, but we haven't gotten there yet.*

One point to note: If a D (like the  $\emptyset_{GEN}$  D) both assigns case and needs case, it will have (in a grammatical sentence) two checked [~~*ucase:xxx*~~] features. For example, in *Pat's dog left*, the  $\emptyset_{GEN}$  D that serves as the head of *Pat's dog* has both [~~*ucase:gen*~~] (because it successfully assigned genitive case to *Pat*) and [~~*ucase:nom*~~] (because it was successfully assigned nominative case by T). We just assume that the morphology knows what to do with this, that *Pat's dog* as a whole should be pronounced in nominative case and not in genitive case. Maybe it's just because nominative case was the most recent case feature checked, or maybe there is a difference between an “I need case” feature and an “I need to assign case” feature that we haven't represented. But it clearly needs to be determined in some way, and in this direction. Similarly, we can't allow the  $\emptyset_{GEN}$  D, by virtue of the fact that it both needs case and needs to assign case,

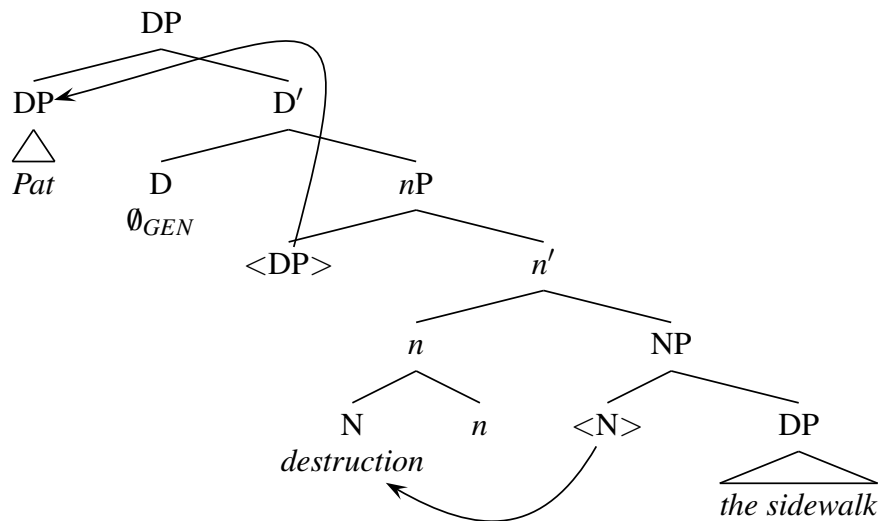
<sup>7</sup>It is near-certain that this is not the best analysis we could have for these “cause-to-have”-type constructions, but so as not to get sidetracked here, we will just stipulate this because it basically works.

to assign case to itself and satisfy both requirements. However this is implemented, it only makes sense if these features cannot value and check each other within a single head.

## 12 DPs, possessors, and clauses

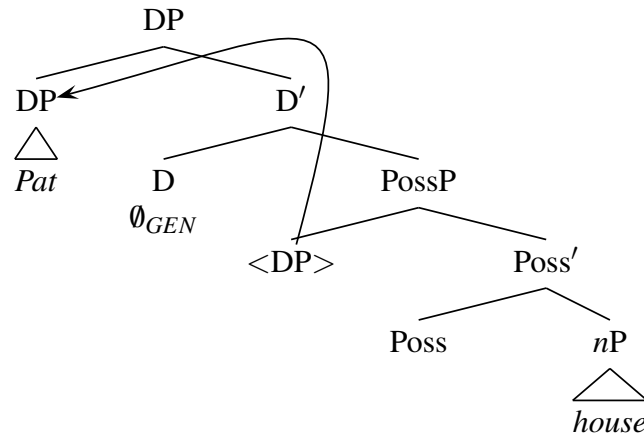
Due to the fact that you can take a sentence like *Pat destroyed the sidewalk* and turn it into a DP (*Pat's destruction of the sidewalk*), we conclude that the DP has the same kind of complexity options that a full clause has. In particular, we need to assume that there is a *nP* (“little *nP*”) inside the DP for the same reasons that we need to assume there is a *vP* (“little *vP*”) in clauses (specifically, for the purpose of having a place for all three  $\theta$ -roles for ditransitive verbs). The structure inside a DP derived from a verbal clause is closely parallel to the verbal clause—instead of *vP* you have *nP*, instead of *V* you have *N*, the case of the “subject” is genitive instead of being nominative, the case of the “object” is “of” instead of being accusative.

The “subject” of a DP (a DP in the specifier of a DP) will surface in genitive case. In those situations, the head of the overall DP is  $\emptyset_{GEN}$ , which has a [*uD\**] feature, that assigns genitive case to a DP and also forces it to move up into the specifier.<sup>8</sup> This genitive DP moves to the specifier of DP and could have come from a couple of different places, depending on what  $\theta$ -role it started with. In a deverbal nominal phrase, the genitive DP might have started in the specifier of *nP* if it was an Agent, or in the complement of NP if it was a Theme. In a possessive construction, the genitive DP will have started in the specifier of PossP.

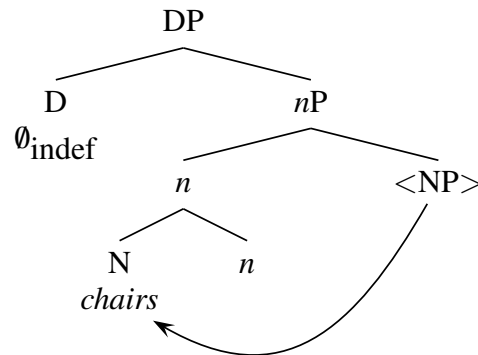


Where a possession relation is indicated (*Pat's house*), the semantics of possession is structurally represented by the presence of Poss. The DP in the specifier of Poss has the  $\theta$ -role of Possessor and the *nP* that is the complement of Poss has the  $\theta$ -role of “Possessee.” If there is a meaning of possession, there is always a Poss.

<sup>8</sup>In Adger’s formulation of this, he used the feature [*ucase:gen\**] to drive movement to the specifier of the higher DP. This removes some of the parallelism between TP and DP, however, and I can not see any differing predictions, so long as we suppose that there is nothing intrinsic in the system that would prevent something of category X from requiring a different thing of the same category X.



All DPs that have an N also have an *nP* within them, and just like with the *v*, the *n* always causes movement of the N up to head-adjoin to it. So, the *n* always has a [*uN\**] feature to make this movement happen. This means that even for simple nominal phrases like *chairs*, with no possessor and not derived from a verb, there is still a relatively complex structure.



Above, I labeled the indefinite article as  $\emptyset_{\text{indef}}$ , but that is a bit misleading insofar as the assumption is that the indefinite D is the same in *chairs* and *a chair*—the difference between the silent determiner for plural indefinites and the pronounced determiner for singular indefinites is one that is handled by the phonology and not one that the syntax cares about. That is, we assume that there is just one indefinite D, but there is number agreement between the N and the D so that we can pronounce the D that “agrees in number” with the N. We implement this by having an agreement feature on D ([*u\phi*:]) that will get the  $\phi$ -features from N via agreement.

### 13 The distribution of $\phi$ -features

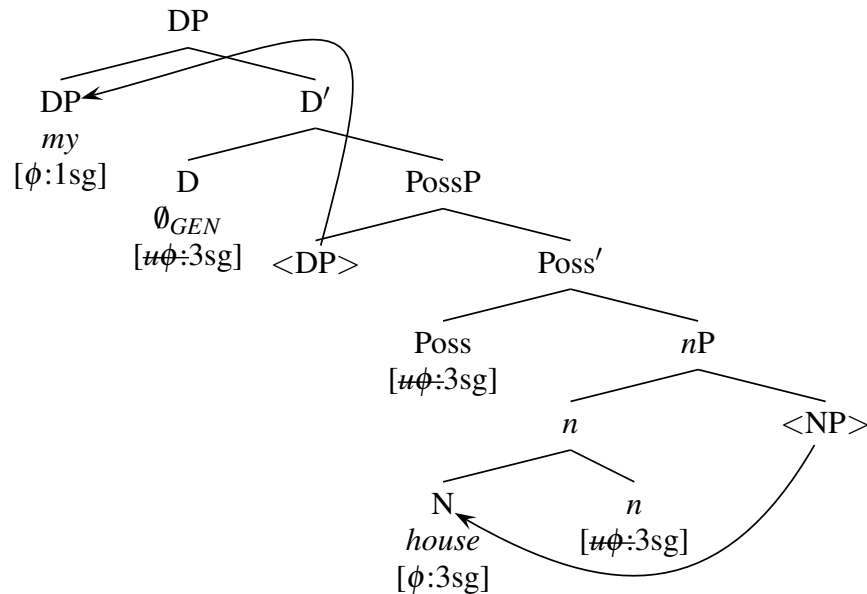
There are a few things to say about how  $\phi$ -features get distributed around in the tree as well. We assume that N is drawn from the lexicon with defined  $\phi$ -features. Having some lexically-determined person and gender, and probably acquiring number specification as it is drawn from the lexicon and placed into the numeration. However, we know that we need the DP containing such an N, if it is in the subject position, to be able to transmit those  $\phi$  features down to a verb for subject-agreement. The N is generally moderately deep within the DP, so we need to consider how these  $\phi$ -features are transmitted.

Another point that fits in this section is the division of nouns into mass and count nouns. Often people speak as if “singular” and “plural” are the only types of number, but there are also mass nouns which are not countable and for which neither singular nor plural makes much sense. It seems appropriate nevertheless

to count this as the same type of feature, even if it means “number” is a bit of a misleading label for it. It’s not really possible to consider mass to be a type of singular or a type of plural, since even in English, mass nouns seem to behave in some ways like singular and in some ways like plural (at least with respect to how the morphology responds). For verb agreement, mass nouns have the same effect as singulars (*The sculptures melt*, *The sculpture melts*, *The ice melts*), but the realization of the indefinite article (as either *a(n)* or  $\emptyset$ ) groups mass and plural (*A sculpture*, *sculptures*, *ice*). So the morphology needs to be able to see mass as its own feature separate from singular and from plural. We will therefore use the labels “sg”, “pl”, and “mass” for the values of number. Some languages also need to represent “dual” somehow as well.

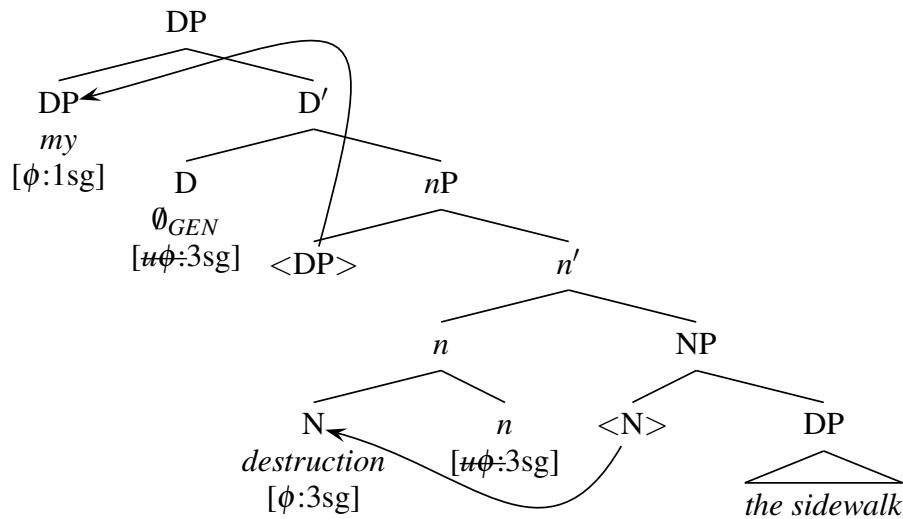
To get the features from N up to D(P) in something like *the sculptures* can be accomplished by positing a  $[u\phi:]$  feature on D. When the structure is built up to the point where D is Merged, D will look down the structure for something with a  $\phi$ -feature and copy its value. This would result in something like a  $[u\phi:3pl]$  feature on D. At this point we must assume that even though the feature is no longer uninterpretable (it is checked), the value is still there and can be used by the morphology to determine the form, and even value another uninterpretable  $\phi$ -feature (like the one on T).

This works for straightforward cases, but things get a bit complicated when a DP has another DP within it, like where there is a possessor or an agent of a deverbal noun. The containing DP does not agree with the genitive DP, but rather with the head noun (*my chair is comfortable*, *\*my chair am comfortable*). So if the uppermost D has a  $[u\phi:]$  feature, we need to make sure that the closest  $\phi$ -features to that uppermost D are those that came from the head noun. If we consider the structure above for *Pat’s house*, it’s clear that if we say nothing further, *Pat* would be the closest DP. But if the PossP itself inherits the  $\phi$ -features from the head noun, then the features of the PossP would be closer to the uppermost D than the features in the specifier of PossP. And that’s what we want. So, we need to say that Poss has a  $[u\phi:]$  feature as well.



The same issue arises with deverbal nouns, but there is no PossP, there is just an agent in the specifier of *nP*. We can deploy the same solution here as well, by saying that *n* also has a  $[u\phi:]$  feature. This should get the agreement right, so long as we make the further assumption that deverbal nouns always have a  $[\phi:3sg]$  feature of their own. That 3sg will propagate up and allow *My destruction of the sidewalk* to come out with 3sg  $\phi$ -features.<sup>9</sup>

<sup>9</sup>Perhaps uncomfortably, it doesn’t seem like we want to put a  $[u\phi:]$  feature on *v*, because that would shield the  $\phi$ -features of an agent or experiencer in ways we don’t want. We’ll have to leave this as an as-yet-unexplained difference between *n* and *v*.

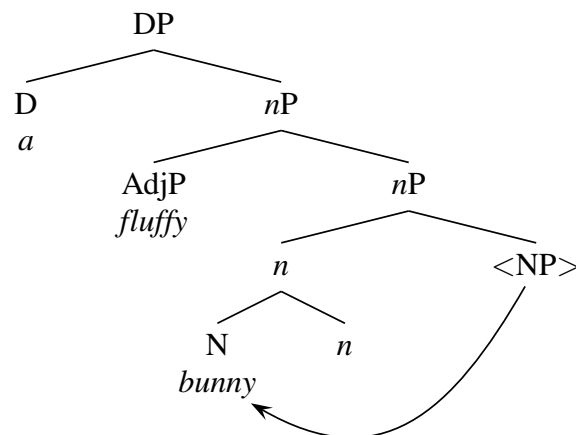


So, in summary, we assume that *n* and Poss each have a  $[u\phi:]$  feature. D also has a  $[u\phi:]$  feature, except for pronouns. A pronoun like *my* would have *interpretable*  $\phi$ -features  $[\phi:1sg]$ .

## 14 Adverbs and adjectives

Adverbs and adjectives are modifiers that attach to the structure via Adjoin. The intuition is that they modify the semantics but do not affect the syntactic status of the thing that they attach to—so if you Adjoin an adverb to a *vP* it was a *vP* before, and it is still a *vP* afterwards. The semantics of adjunction is such that the thing being adjoined to is the thing being modified semantically. So, if you have an adverb that describes the manner of an event, it is it properly adjoined to the *vP*, and if you have an adverb that is more of a “scene-setter” describing time or location, it is properly adjoined higher, to the TP.

Adjectives generally adjoin to *nP*, and adverbs generally adjoin to either *vP* or TP. Although it is true that adjectives cannot adjoin to *vP* (adjectives go with nouns) and adverbs cannot adjoin to *nP* (adverbs go with verbs), this restriction is not encoded in our system anywhere. In particular, do not succumb to the intuitive thought that “an adverb needs a verb” and that there therefore must be a strong selectional feature on adverbs. Adjunction (of phrases) is an operation that does not check features, so we cannot say that adverbs have a feature to check. We must instead assume that what goes wrong if you mix adjectives and verbs or adverbs and nouns is something either in the semantics (it wouldn’t be a semantically interpretable sentence if they were mixed like that) or in the pronunciation (if adverbs and adjectives are actually the same and only agreeing with what they attach to, such that being attached to a verb yields a pronunciation as an adverb and being attached to a noun yields pronunciation as an adjective).



## 15 A list of heads and features

Just as a kind of quick reference, here's a list of heads and features that they would generally occur with. Features that sometimes occur but not in all situations are in parentheses.

*We haven't yet discussed C, and the clause-type features, but I will leave them in here anyway.*

Heads and features	
N	$\phi$ :value
<i>n</i>	$uN^*$ , $u\phi$ : (ucase:of)
Poss	$uD^*$ , $u\phi$ :
D	$u\phi$ : , ucase: ( $uD^*$ , ucase: gen)
P	$uD^*$ , ucase:acc
V	( $uD^*$ , $uP^*$ , $uT^*$ , $uC^*$ )
<i>v</i>	$uV^*$ , $uInfl$ : (ucase:acc, $uD^*$ )
Pass, Perf, Prog, M	Aux, $uInfl$ :
Neg	
T	tense:value, $uD^*$ , $u\phi$ : , uclause-type: (ucase:nom, ucase:ing)
C	clause-type:value ( $uwh^*$ , $utop^*$ , ucase:null)