What are Linguistic Representations?*

David Adger

August, 2021

Abstract

Linguistic representations are taken by some to be representations of something, specifically of Standard Linguistic Entities, such as phonemes, clauses, noun phrases etc. This perspective takes them to be intentional. Rey (2021) further argues that they are inexistent. Here I argue that linguistic representations are simply structures, abstractions of brain states, and hence not intentional, and show how they nevertheless connect to the systems that use them.

Keywords: terminology, ontology, representations, syntax, phonology

1 Introduction

We generative linguists can be cavalier with terminology and we usually cut short our explanations of phenomena before clearly linking

---

them to the phenomenal world. This is a side effect of talking amongst ourselves too much. It causes no end of headaches for philosophers, with a spate of books over the years attempting to work out what on earth we are doing, and whether it is what we should be doing. The latest of these is Georges Rey's (2020) *Representation of Language*, which is a broad defence of some of the philosophical implications of Chomsky's approach to language, as developed over the last half century or more. Rey does, however, take a different direction philosophically from Chomsky in a number of areas. There is much of the book that I agree with wholeheartedly, including the nuanced approach to issues of innateness, the defence of linguistic judgments as data, the arguments against externalist views of linguistic entities, the broadly psychological conception of the part of linguistics that generative grammar encompasses and much more.

A major proposal in Rey's book, though, is that linguistic representations are intentional (that is, they are representations "of" something and hence "about" that something). A good summary of the view is given on page 211. By this point in the book, Rey has identified Standard Linguistic Entities (SLEs—Noun Phrases, phonemes, Tense nodes, etc.), and he says of them (emphasis in the original):

> [Generative linguists] are concerned with an *internal computational explanation* defined over mental *representations* of them [SLEs]—just as vision theorists are concerned not with real colors, but with computations in the visual system over representations of colors, whether or not there really are colors answering to them

Rey's core point is that the computations that generative linguists appeal to to explain linguistic phenomena are over representations and the representations are "of" SLEs. Rey further argues that SLEs are 'intentional inexistents'; they are non-existent entities that nevertheless can be what representations are about. Unicorns, Zeus, and noun phrases are inexistents, and contrast with chairs, Chomsky and quarks, which exist. Mental representations are equally representations of both kinds of thing, and are hence intentional.

In response, I will argue that, in fact, when linguists use the term 'linguistic representation' they usually use that term in a non-relational sense, meaning 'linguistic structure' (where 'linguistic' may be replaced by phonological, syntactic etc.). Linguistic representations are therefore not "of" or "about" anything. They are ontologically grounded not by intentionality but by being abstractions from hypothesized physiological properties of the brain. My claim is narrow. I'm not arguing that our mental representation of Zeus or Chomsky are non-intentional. I'm arguing that the posits of generative linguistics are not representations, and hence no question of intentionality arises.

I don't think I'm saying anything startlingly original here; the views I'll sketch out are not commonly articulated by generative linguists, though they are views that can be discerned through Chomsky's writings. Chomsky (2000) (159-161), for example, explicitly says, when talking of Marr's computational theory of vision that " "representation" is not to be understood relationally, as "representation of,"" extending this approach to the representations involved in language, which "enter into interpretation, thought, and action, but [for which] there is no reason to seek any other relation to the world." Chomsky

mentions the world here, but I think we can be fairly sure that we can interpret this as including inexistent entities of the sort Rey appeals to.

I think what I have to say here is similar also to Collins' (2014) proposal that "the theorist employs dyadic representational talk in order to frame generalisations about monadic, non-relational states." That is, if I understand it correctly, exactly right. The SLEs that linguists talk about with locutions such as "TP", or "c-command", are just structures and mechanisms of the mind (abstracted from unknown physiological structures and mechanisms). Collins, following Goodman (1949), proposes a "method of hyphenation" to achieve the connection between theoretical talk and the actual entities, but I don't think such a method is necessary to understand what linguists are doing. I'll argue that it's simply a fact that generative linguists use the term "representation" as a technical monadic term, and sketch out why we do this[1].

Turning from terminology to the ontology of linguistic representations, Rey takes it to be a "scandal that linguists have not sufficiently pondered what they [=SLEs] are. [p288]." But there is no surprise here. Generative grammarians think of the posits of their theories in much the same way as physicists think of quarks, particles in wave form, etc.: they are the best explanations of the phenomena, though

---

[1]Rey mentions in a footnote to page 50 of his book that in early works Chomsky employs the term "representation" and does so "without any assumption of the (meta-) representations being of the phrases of a natural language". My contention, defended in section 2, will be that this has remained the standard use in generative grammar ever since (agreeing with Collins 2004).

we do not know exactly what they correspond to in the world we can currently observe (an approach Rey explicitly approves of (e.g. p92ff)). Again, this is old hat. As Chomsky notes in his (2003) response to Rey (2003), "corresponding in some manner to such psychological entities [SLEs—DA] there are physiological configurations, also said to be internally represented." (see also Chomsky 1980, passim). The phrase "in some manner" simply keeps open exactly how the structures of the mind are to be related to the structures of the brain (whether by properties of neurons (e.g. Pallier, Devauchelle, and Dehaene 2011), the internal structure of cells (e.g. Gallistel and King 2011), electrical oscillations (e.g. Ding, Melloni, Zhang, Tian, and Poeppel 2016), etc.).

Rey is of course aware of these ways of thinking about the issues, and argues that even if one thinks of generative grammar in this way, the grammar and its outputs must have some connection to the physical world and intentionality must play a role in that connection: the speaker's grammar builds representations of the relevant properties. Rey provides some criticism of remarks in Adger (2019b), where it is suggested that the connection is rather made through transduction. Rey grants a role for transduction but argues that "intentionality is needed where mere transduction gives out. It provides a "common coin" needed at the interface between grammar and, *inter alia*, the perceptual system." [p9]. Intentionality is needed because " [c]hildren must hear ambient noises, exactly as they seem to do, *as* SLEs of the ambient language" [p284, emphasis in the original]. In section 4, I'll try to be more explicit about how I think a generative grammar can both explain core linguistic phenomena, and be incorporated into a model of the speaker, without appeal to intentionality.

2 Terminology

Generative linguists have a bad habit of using non-technical terms technically and marking the difference by minimal orthographic means, or, indeed, by not marking the difference at all. Many a student has been perplexed by why case and Case are so different, or by why Merge doesn't actually involve merging. Terminology proliferates in our field, unfortunately making it a full time job to keep abreast of significant changes. Worse, we sometimes systematize our sloppiness, such as the common use of the term 'grammar', for both the linguist's theory and the natural (mental) object that it is a theory of (Chomsky 1965, p25), as well as the posits of that theory and the objects those posits are about (NPs, copies, features, etc.).

The term 'representation' is one of these words whose use strays from common practice. The first extensive use of the term in generative grammar was in Chomsky's (1955) *Logical Structure of Linguistic Theory* (LSLT), published in mimeographed form in 1955 but not released as a book until twenty years later. LSLT introduced the notion of "Level of Linguistic Representation", where each such level is a system of symbols and a means for combining them. These systems of symbols were intended to provide an analysis, at a particular level, of utterances[2].

At the time Chomsky wrote LSLT the ontology of linguistics was resolutely externalist, following Bloomfield (1933). Speech was the

---

[2]This word, utterance, was used at the time pre-theoretically to refer to something a speaker of a language may say. The term has a slightly different use today, where it is commonly used to mean 'token of a sentence'. Chomsky uses the pre-theoretical notion throughout LSLT and later.

real thing, and linguists represented properties of speech in their work. Linguistic theory was built from the bottom up, starting with observable properties of utterances. Through processes of segmentation and classification, more abstract levels of structure could be constructed (Harris 1951). The more abstract levels are built from less abstract ones, creating a system of levels where each level represents properties of larger parts of the utterance (phonemes, morphemes, words, phrases, sentences, texts). Representations at the different levels were typically ordered sequences (strings) of symbols. A phonemic representation, for example, was a series of symbols that referred to the properties of that utterance that were part of the phonemic structure of the language (where the structure of a language was conceived of in a Saussurian fashion). A morphemic representation was constructed from the phonemic one, and so on. The word "representation" here is used in its usual relational sense: the representations were about properties of the utterance, so Harris (1951) could write, for example (p34), "[T]he representation of speech as a sequence or arrangement of unit elements is intimately connected with the setting up of phonemic distinctions between each pair of non-equivalent utterances".

At the beginning of LSLT, the discussion of lower levels of linguistic structure (phonetic, phonemic, morphophonemic, word, category) clearly connects to this perspective, and Chomsky in fact uses the locution "representation of the utterance." However, the perspective in LSLT is top-down rather than bottom up. Each level is an independently specified concatenation algebra consisting of a set of primes (symbols) and relations at that level, and the algebra specifies certain strings of symbols as well-formed, so the "representations" are not de-

rived from the utterance. Rather, each string of symbols at one level can be converted into a lower level of structure through a specified set of mappings (e.g. at the phonemic level we have a single symbol /p/ for the bilabial stop in *pill* and *spill*, but there is a mapping to the phonetic level where there are distinct symbols, representing an aspiration difference between the stops ([p^h] vs [p])).

The argument for the independent specification of levels as opposed to their derivation from properties of utterances was most clearly made for the higher levels: the principles of sentence structure are obviously quite distinct from those that can be extracted from physical properties of the utterance, and Chomsky highlighted this fact, arguing that phrase structure representations involved sets of strings, as opposed to a single string. The abstract character of syntactic structure was quite distinct from the (essentially finite state) character of lower levels. This provided general motivation for the idea that the "levels of representation" must be independently generated systems of formal symbols, related to each other by mappings. This proposal was bolstered by a number of case studies showing that better analyses of linguistic phenomena could be developed in such a system. In the approach Chomsky developed, purely formal properties of the levels, especially the phrase structure and transformational levels, had real explanatory value, even though they could not be mapped to the physical utterance by procedures of classification, segmentation, etc.

Though the linguistic levels were still called "levels of representation" in LSLT, their symbols don't represent physical properties of utterances or anything that can be directly derived from such properties: they are fundamentally monadic and sui generis. From this

perspective, a "representation of an utterance" at level L should not be understood as the symbols being 'about' the utterance any more than the combination of two hydrogen atoms with one oxygen atom is 'about' water; they are rather an independently generated collection of structures related to each other by mapping functions.

LSLT severed the link between physical utterances and grammatical explanation that had been central to the approach of Bloomfield and his followers. It argued that you cannot directly extract a grammar from the signal; rather the grammar is independently specified and its links to the signal are indirect. This severance underlies the later competence-performance distinction: there is an independent computational system (the grammar) which is accessed by all the other components of cognition that use it (for particular events of speaking, signing, thinking, planning, etc.). The grammar, and the structures it generates, are not representations "of" anything physical.

Chomsky in LSLT kept the word "representation" but changed the theory behind it (without even, as is his usual wont, putting it in capitals). This reconceptualization of the notion of linguistic level, not as a representation of properties of the utterance, but rather as an abstractly specified structure, is, I think, the root of the idea that became the standard monadic view of linguistic representations: each level is given as an autonomous system that is related by mapping functions to other systems. Together, the levels and the mappings function to specify the grammar of the language. Rather than representing linguistic properties of utterances, the grammar generates structures whose properties correspond to them (see below for discussion of how). As the theory developed, this notion of representation as generated struc-

9

ture remained, and it is the main use of that word in Chomsky's work since. We might think of it as the following (shamelessly borrowing from Soare's (1996) Recursion Convention):
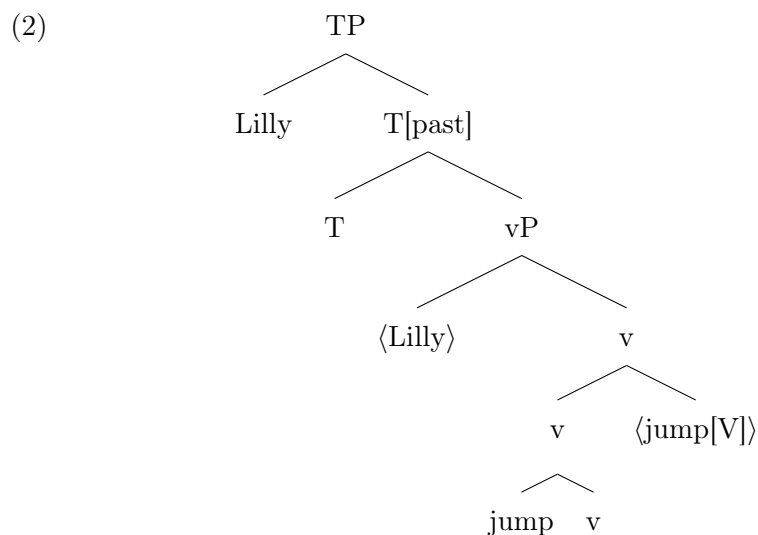
(1)   The Representation Convention
        Use "representation" to mean independently generated linguistic structure as defined by the computational mechanisms of the relevant variant of generative grammar.

As well as using the term "representation" to talk about particular mental structures associated with utterances in some fashion (see below), Chomsky uses the term in a related (non-relational) way to talk about the grammar as a whole. Rey quotes Chomsky in the introduction to the 1975 edition of LSLT writing "The child constructs a mental representation of the grammar of the language" and notes the use of the same term in Chomsky's 1965 account of language acquisition. However, all of these uses of the term should be read in the light of the discussion above. "A mental representation of the grammar of the language" is just the mental structure (brain state) which is, at the relevant level of abstraction from physiological mechanisms, the grammar of the language. Chomsky is particularly clear on this in Chomsky 1987 where he elucidates what knowledge of language is: a generative procedure which is a particular brain state. Possession of that generative procedure allows an individual to speak/sign and understand their language(s). No notion of representation in the relational sense is intended or needed.

Of course the word 'representation' can also be used relationally.

The physical configuration of lines and letters on the page below in (2) is a representation of a syntactic structure.

(2)

```
                        TP
                      /    \
                  Lilly    T[past]
                           /     \
                          T       vP
                                /    \
                          ⟨Lilly⟩     v
                                     /   \
                                    v    ⟨jump[V]⟩
                                   /  \
                                jump   v
```

The syntactic structure which (2) is a representation of, however, is something which is not on this page. That structure is what Adger (2019a) calls "a gesture of the mind". A gesture with my hand is made possible by my anatomy and made actual by my intentions. When I utter or comprehend *Lilly jumped*, I make a gesture of my mind: the physiology of my brain takes on a particular configuration, dependent on what I, the speaker, intend to express (or by what I perceive). (2) represents this gesture of the mind[3].

Linguists, as Rey points out, usually collapse the use/mention distinction, so we talk of T c-commanding v in structures like (2), though what we mean is that T c-commands v in cognitive structures like those

---

[3]Of course, (2) is a representation of a *type* of mental event, as opposed to a particular token of it (see Rey's useful discussion of tokens and types in linguistics in his chapter 3).

represented by (2). Rey is perplexed by the idea that the brain really contains T, and not a representation of T, but that is exactly what linguists intend. "T" in (2) represents T, a putative object of the world that occurs in certain mental events. We say that T c-commands v in a finite clause much as we might say that hydrogen bonds to oxygen in water.

Rey's idea of "representational pretence" is almost the opposite of this. He suggests that generative linguists are committed to computations over representations of SLEs, but not to SLEs themselves. In actuality, we are committed to SLEs, and to the idea that SLEs are built by processes that are computational, but we are only committed to our theories of these being representational (in the standard sense that the theories are about the SLEs etc.).

Summarizing, generative linguists usually use the word "representation" to mean autonomously generated structure. Moreover, we indulge in a general use/mention collapse where we talk of structures by means of talking about their representations. Combining the habit of use/mention collapses with the Representation Convention does, I admit, make penetrating the murky depths of generative syntax less straightforward than one might wish.

3 Ontology

I said above that when a generative linguist says "T c-commands v," we are talking about putative entities in the world, (types of) gestures of the mind that are ultimately to be understood as configurations of the brain. We are not talking about representations. Rey's counter-

argument to this view is as follows:

> Given that phonemes are not in the acoustic stream, does
> anyone really want to insist that they have a "theoretical"
> reality comparable to, say, electrons? If they do not exist in
> the acoustic stream, where are they? Do they have causal
> powers? How? ... Of course, if (as any right minded sci-
> entist likely thinks?) the mind just is the brain, then, it
> follows SLEs are in the brain, no? So SLEs would have the
> causal powers of brain states. Silly philosophical problem
> dissolved. But not quite so fast. The inference is obvi-
> ously invalid. Ghosts, angels, the fountain of youth, colors,
> rotating cubes, ideal geometric shapes: lots of "things we
> think about" are said to be "psychologically real" and "in
> the mind," in that thoughts about them play a real role in
> people's lives. But it would be lunacy to say that the things
> thought about, themselves, are in the brain!   Rey (2021)
> p324.

I think what goes wrong here is the idea that "we think about"
phonemes, T, c-command, and v, in the same way that we think about
ghosts and colours. The linguistic theorist certainly does think about
T etc., and, as I noted above, the theorist is engaged in an activity
that involves representing (in the familiar dyadic sense). If I think
about a ghost, there is a momentary brain state, but, as Rey notes,
that brain state is not a ghost. However, unless I'm being a theorist,
I don't think about phonemes, or T, or c-command. My mind takes
on configurations in which T is involved (different configurations, for

example, if I think to myself the sentences *Lilly jumped* and *Did Lilly jump?*). But though the mental configuration involves T, it is not about T.

T, v, c-command etc are technical terms used to refer to putative objects and relations at a particular level of description (the mental) cashed out as brain states in some as yet unknown way. Again a chemical analogy is apt: hydrogen, oxygen, valence and bonds are technical terms used to refer to objects and relations at a particular level of description (the chemical), cashed out as quantum effects in some way that physicists understand. Aside from our relative levels of ignorance of the way that brains work compared to the way that chemical compounds work, the ontology is of the same type.

4   Connection

If syntactic representations are monadic, how do they connect with the world? They must, as Chomsky notes "be accommodated to the sensory and motor apparatus." (Chomsky 1995, p221). For Rey, this is where intentionality is indispensible. He argues that brains are sensitive only to local physical properties of their environments, and that what he calls "abstruse properties," such as *being a TP* and other SLEs are not local or physical. It is, he suggests, intentionality that bridges the gap.

Adger (2019b) argues, in contrast, in that transduction is involved in the connection between the grammar and the sensory and motor apparatus, but of course an abstruse property is not transducible. If SLEs are abstruse properties, transduction is, indeed, insufficient.

However, here I think Rey is assuming what he wants to argue, since, as I have suggested above, SLEs are not abstruse properties. They are just abstractions over brain states. I think the best kind of counter-argument is a counter-proposal, so in the remainder of this paper, I will sketch out how the grammar connects with the physical world in two important areas: (i) how the grammar explains linguistic phenomena; (ii) how the grammar can be "accommodated to the sensory apparatus" in parsing. Space precludes discussion of production and of relations to meaning, but I take it that Rey's argument, that intentionality provides a common coin for the grammar and the parser, is the crucial one to address here (see also Collins and Rey 2021).

Let's first look at how grammars figure in an account of speakers' linguistic judgments. Linguistic judgments are part of the data (and hence explananda) of generative linguistic theory, and Rey uses their properties to provide strong arguments for doing linguistics in a generative way.

Take an example like the following, a core case of the kind of thing syntactic theory is meant to explain (what Rey calls 'why-nots'):

(3)   Which woman did you wonder when left?

Native speakers of English have a fairly uniform reaction to examples of this sort (they find it unacceptable). If you, a philosopher, ask a syntactician why the reaction is what it is, they will unhelpfully say something like (4), and everyone's eyes will glaze over:

(4)   The sentence is ungrammatical because the wh-subject has
        been extracted from the lower subject position across a filled

15

specifier of CP, which violates a locality condition on its relationship with its trace.

There's a lot to be unpacked here which is relevant to the question of how the grammar actually works to explain the phenomenon, and which is usually left unsaid, but I want to initially focus on a difference in presuppositions about the philosopher's question. The syntactician is answering a question about the grammar, but the philosopher is asking a question about the speaker.

What the syntactician could have said was the following.

(5)   A speaker encountering (3) will judge it as unacceptable because they go through the following process: they *connect* elements of the physical signal they hear (or see) to internal grammatical structures, which we call lexical items, given in a particular order. The speaker's grammar (their generative procedure) is unable to generate a structure the terminal nodes of which have the lexical items the speaker has identified in the presented order: specifically the structural relationship between the subparts of structure *which woman* and *left* are disallowed by the grammar. This is why the speaker's grammar fails to generate a structure for (3), and explains their reaction.

I'll return to the crucial term "connect" in the discussion of transduction below, so for the moment will leave it unanalysed.

Let's first see how this kind of explanation can be given more specificity and rigour. This will involve some technical material, but I'll simplify and highlight the issues that are relevant to Rey's concerns.

I think it's necessary to go through the details so as to be clear about how a non-intentional explanation will work.

The lexical items that the physical signal connects to have a number of formal properties that distinguish them from each other and from other lexical items. Lexical items will have the following kinds of shape (this is broadly the analysis given in Adger 2003, a standard textbook):

(6)   a.   *which*[D, wh]

        b.   0[C, $u$wh*, T*]

        c.   did[T, uInfl:past, uAgr:_, uD*]

        d.   *woman*[N, sing]

        e.   *you*[D, 2, sing, nom]

        f.   *wonder*[V, uD, uC*]

        g.   ...

The representations in (6) refer to actual hypothesized items in a speaker's grammar, where the formal differences in the representations (e.g. D vs C; $u$wh vs wh; T vs T*, D vs uD*, etc.) correspond to hypothesized differences in the cognitive structures referred to (I'll explain some of the details below). Lexical items generally consist of some phonological structure (I use orthography here in the representations for ease of exposition, with 0 to signal a phonology that lacks phonetic expression) and some syntactic structure given as a collection of features. "Feature" is a technical term that just means distinguishing property of the internal structure, much as lego-bricks might be distinguished from each other by their shape, or connective structures, see e.g. Adger and Svenonius 2011. The use/mention collapse that

linguists are prone to means that we talk of features being both in the representations in (6), and being properties of the relevant mental configurations. Contra Rey, the feature (qua distinguishing property of a mental configuration) that a feature (qua symbol of the theory) like [+nasal] or [uD*] represents is, in fact, in the mind, ultimately the brain. It is not in the nose, the acoustic signal, nor is it inexistent.

A speaker of English, encountering (3) will connect various parts of the signal with the lexical items in (6): They will hear an acoustic signal whose linguistically relevant properties we can represent in phonetic notation as [wɪtʃ] but those linguistically relevant properties are an internal structure, a tokening of a particular configuration of the mind. That phonetic configuration will connect to a stored lexical item (in this case the item *which*[D, wh]); similarly, the English speaker will hear a signal whose phonetic properties can be represented as /wumən/ and the mental configuration that that tokens will be related structurally to the configuration previously tokened by [wɪtʃ] (it is ordered after it), and it will connect to the stored lexical item *woman*[N, sing]; similarly for /dɪd/ and *did*[T,Infl:past, uD*]; and so on.

In a real time speech event, this information is part of parsing the signal, and will take place roughly incrementally, with multiple lexical items being considered and rejected as they enter into the parse, in ways that will depend on the parsing strategy, the grammar, and the developing parse. Information from the grammar and from the context is relevant to this process, so the speaker may include lexical items with no corresponding phonetics as part of their parse (for example, the Complementizer 0[C, $u$wh*, T*], whose syntactic properties will

18

create a structure that forces T to be local to it and hence precede the *you*). Lexical items that are more frequently encountered in particular syntactic contexts will also be selected, leading to classical garden path effects such as *The old man the boats*, where the frequent lexical item for *man*, a noun, is selected while the only parse that the grammar will license is one where a less frequent lexical item *man*, a verb, is needed. Syntacticians attempt to abstract away from the vagaries of parsing so that accounting for the patterns of judgments is addressed by answering the question: when we have a particular collection of lexical items in a particular order, what structures built by the grammar will yield those items in that order?

For the case in hand, during the encounter with the physical signal, at some point phonetic processing leads to selection of a lexical item with the phonology (orthography) *wonder*, which syntactically has two 'slots' (think of the lego-brick analogy), represented by what are called 'selectional features' (uD, uC*) in the lexical item in (6): intuitively, these slots are for a subject and a clausal complement; their formal role is to ensure that a particular kind of structure is built involving the verb, a DP and a CP in particular structural relationships. The differentiation between C and D will result in different structural objects being connected to this item, and the presence of the * on C specifies exactly how the structural connection is made (it requires a close structural connection). If the speaker who judges (3) uses this lexical item when encountering (3), then the grammar will build a structure, dependent on the formal properties of the lexical items and the rules of the grammar. This structure is represented in (7):

19

(7) $[_{CP}$ ⟨$[_{DP}$ which woman]⟩ C $[_{TP}$ you T wonder $[_{CP}$ ⟨when⟩ C $[_{TP}$
      ⟨$[_{DP}$ which woman]⟩ T $[_{VP}$ left ⟨when⟩]]]]]

For those not versed in syntactic theory, the only crucial thing to note is that $[_{DP}$ which woman] appears in two places in the structure[4]. $[_{DP}$ which woman] is a single syntactic element but is connected within the structure to two positions (it is local to both the matrix C and to *wonder*). This is also true for the item *when*. These cases, where a single syntactic element appears in multiple positions in structure, usually correlate with a meaning that logicians represent as variable binding, so (7) would have a meaning representable as:

(8)   Who is the x such that you asked when x left

However, there is something awry with the structure which (7) represents. Based on many phenomena, generative syntax makes the claim that, when a single syntactic element is in multiple positions in a structure, those positions must be related to each other in a particular way: specifically that relationship cannot cross a clausal boundary (a CP) unless it is at the edge of that boundary structurally. This is part of the specification of the generative procedure (indeed, it is hypothesized to be part of the procedure that is universal to human beings). The structure in (7) is, then, not generable by the grammar because the positions that $[_{DP}$ which woman] occupies are not related in this way. This means that a speaker who encounters (3) will be unable to assign

---

[4]I have notated this, as is common, by angled braces in the representation of the configuration; the actual configuration is usually thought to involve some kind of looped structure (see Adger 2019a for an exposition of the idea aimed at a popular science audience).

a structure like (8) to the sequence of lexical items they have used (alternatively, they assign the structure but the violation of the condition on syntactic relationships marks it as ill-formed). This is why (3) is unacceptable, and specifically why it can't have the meaning in (8)[5].

The reason that a speaker, on encountering (3), judges it to be ungrammatical is that the speaker's grammar is unable to generate a structure that has the particular lexical items in an appropriately ordered sequence. The theory has explained a core piece of linguistic data (a 'why-not', in Rey's terms): the judgments made by native speakers. It does so without appeal to representations of SLEs; the explanation is one that trades in SLEs themselves.

Of course this explanation requires a connection between the physical properties of the signal, and the phonological properties of the lexical items. This is where transduction comes in. As Pylyshyn (1984), p152, notes a "transducer performs a rather special conversion: converting computationally arbitrary physical events into computational events."

Our speaker on encountering (3) is exposed to a physical signal (a computationally arbitrary physical event). The signal could be aural in the case of spoken language, or visual in the case of sign language. It could be written, or even, as in the cases of the Tadoma method, tactile. Physical properties of the signal have to be mapped into some structure that the grammar can interact with. This appears, at least

---

[5]Some languages can have what are called resumptive pronouns, where a resumptive pronoun is an independent element of syntax that has the interpretation of the bound variable x in (8). Such languages will still not have grammars that generate (7), but they would provide a different route to the meaning in (8). English is not such a language.

for phonetic properties, to be innate.

Mahmoudzadeh et al. (2013), in a study of pre-term babies, investigated cortical responses to the sounds [b] and [g] and compared these with cortical responses to male vs female voices. There are large acoustic differences in the latter case, while the acoustic differences between [b] and [g] are extremely subtle. Mahmoudzadeh's team showed that there was a major response in the language-associated areas of the brain to the [b]/[g] distinction, but not to the male/female distinction. Given these babies are pre-linguistic, it appears their brains are innately attuned to possible phonological differences and not to other acoustically larger effects. Rats showed the opposite effect, suggesting that human brains, but not rat brains, are wired to extract phonologically relevant information from the physical signal.

Acoustic differences in perception between [b] and [g] correlate to a featural difference between them in their phonological feature specification. [b] is [+labial] and [g] is [+velar]. Mahmoudzadeh's team's work suggests that these phonological features are differentiated brain states, so there is a transduction between a computationally arbitrary physical event (signal) and a computationally relevant event (a tokening of a phonological structure). Part of the human linguistic endowment, then, is a transducer that connects signal to structure (that is, to the phonological features of the lexical items)[6].

---

[6]Given the phonology of sign languages is similar in many respects to the grammar of spoken languages, a similar transduction must be available for visual signals, where the phonological structure of signs is accessed through a function that maps from aspects of hand-shape, motion, direction, etc. It may be that the transduction in general is underpinned by the human imitative capacity: the child uses its internal structures to

We have now completed the connection. A transducer maps physical properties of the signal to physical properties of a brain state, which are, at a level of abstraction from the physical, the phonological features linguists posit (and represent via symbols like [+nasal]). Those phonological features are stored in long term memory with certain syntactic features as lexical items. The grammar specifies the unbounded relationship between sets of lexical items and structures. The speaker's parser uses the lexical items the speaker has selected together with the grammar to create a structure for the incoming signal. Partial structures can be built, but it may not be possible to build a complete structure, depending on the lexical items available, or the (perhaps universal) constraints on the structures that the grammar allows. What Rey calls abstruse properties, like being a noun, are simply not abstruse: they are given by the syntactic features the lexical items have and the possible ways that the grammar can combine them. No appeal to representations or intentionality is required.

More needs to be said about the way that syntactic features like [D], [uC*] etc. are acquired. Rey again suggests that intentionality is needed, but, given that generative grammar takes many features to be innate and specific to syntax, what is needed is an internal learning mechanism that connects statistical or distributional properties of the encountered data to those features, which then serve to organise further data as it is encountered, bootstrapping the grammar. Intentionality doesn't do the work, computation does.

---

imitate the external signal thereby making those structures an active part of language learning.

5   Conclusion

I have, of course, not provided a full account of acquisition, parsing, production and other aspects of language use here, but I hope to have shown how such accounts do not presuppose the notion of intentionality. This is just as well, as generative grammar has never understood linguistic structures to be representations of anything.

References

Adger, David. 2003. *Core Syntax*. Oxford: Oxford University Press.

Adger, David. 2019a. *Language Unlimited: The Science Behind Our Most Creative Power*. Oxford: Oxford University Press.

Adger, David. 2019b. Linguistic representation: a note on terminology vs. ontology, https://ling.auf.net/lingbuzz/004616.

Adger, David and Svenonius, Peter. 2011. Features in minimalist syntax. In Cedric Boeckx, ed., *Oxford Handbook of Linguistic Minimalism*, 27–51, Oxford, UK: OUP.

Bloomfield, Leonard. 1933. *Language*. New York: Holt.

Chomsky, Noam. 1955. The logical structure of linguistic theory, revised 1956 version published in part by Plenum Press, 1975; University of Chicago Press, 1985.

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

Chomsky, Noam. 1980. *Rules and Representations*. Oxford: Basil Blackwell.

Chomsky, Noam. 1987. Reply to alexander george and michael brody. *Mind and Language* 2:178–197.

Chomsky, Noam. 1995. *The Minimalist Program*. Cambridge, MA: MIT Press.

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

Chomsky, Noam. 2003. Reply to rey. In Louise M. Antony and Norbert Hornstein, eds., *Chomsky and his Critics*, 274–286, Oxford, UK: Blackwell Publishing.

Collins, John. 2004. Faculty disputes. *Mind and Language* 19:503–535.

Collins, John. 2014. Representations without representa: Content and illusion in linguistic theory. In Piotr Stalmaszczyk, ed., *Semantics and Beyond: Philosophical and Linguistic Inquiries*, 27–64, Berlin: De Gruyter.

Collins, John and Rey, Georges. 2021. Chomsky and intentionality. In *A Companion to Chomsky*, 488–502, Hoboken, NJ: Wiley-Blackwell.

Ding, Nai, Melloni, Lucia, Zhang, Hang, Tian, Xing, and Poeppel, David. 2016. Cortical tracking of hierarchical linguistic structures in connected speech. *Nature Neuroscience* 19:158–164.

Gallistel, Charles R and King, Adam Philip. 2011. *Memory and the computational brain: Why cognitive science will transform neuroscience*, volume 6. John Wiley & Sons.

Goodman, Nelson. 1949. On likeness of meaning. *Analysis* 10:1–7.

Harris, Zellig. 1951. *Methods in Structural Linguistics*. Chicago: University of Chicago Press.

Mahmoudzadeh, Mahdi, Dehaene-Lambertz, Ghislaine, Fournier, Marc, Kongolo, Guy, Goudjil, Sabrina, Dubois, Jessica, Grebe, Reinhard, and Wallois, Fabrice. 2013. Syllabic discrimination in premature human infants prior to complete formation of cortical layers. *Proceedings of the National Academy of Sciences* 110:4846–4851.

Pallier, Christophe, Devauchelle, Anne-Dominique, and Dehaene, Stanislas. 2011. Cortical representation of the constituent structure of sentences. *Proceedings of the National Academy of Sciences* 108:2522–2527.

Pylyshyn, Zenon Walter. 1984. *Computation and cognition*. MIT press Cambridge, MA.

Rey, Georges. 2003. Intentional content and a chomskyan linguistics. In A Barber, ed., *Epistemology of Linguistics*, Oxford: Oxford University Press.

Rey, Georges. 2020. *Representation of Language*. Oxford, UK: Oxford University Press.

Soare, Robert I. 1996. Computability and recursion. *Bulletin of symbolic Logic* 2:284–321.