Linguistic Representations:

a note on terminology vs. ontology

DRAFT: comments welcome [*]

David Adger

Queen Mary University of London

May, 2019

1 Introduction

Generative linguists have a bad habit of using non-technical terms technically. Worse, we often use the same term with what Chomsky calls 'systematic ambiguity' (Chomsky 1965), such as the term 'grammar', for both the linguist's theory and the natural (mental) object that it is a theory of. An example somewhat similar to this is the term 'representation', which is what I'll be concerned with here. The purpose of this note is to clarify how linguists use

[*]This is a very belated write-up of a talk I gave in Trondheim 2015 at the workshop on *Linguistics, Representation and Cognitive Science* entitled 'Syntactic Structure as Mental Gesture'. Many thanks to all the participants, and especially to Georges Rey for robust discussion.

that term, and to argue that, despite all the talk of syntactic representations in generative linguistics, the ontology of the theory requires nothing that is represented.

The physical configuration of lines and letters on the page below in (1) is a representation of a syntactic structure (assuming for the moment, probably counterfactually, that the theory that underlies this discussion is true).

(1)

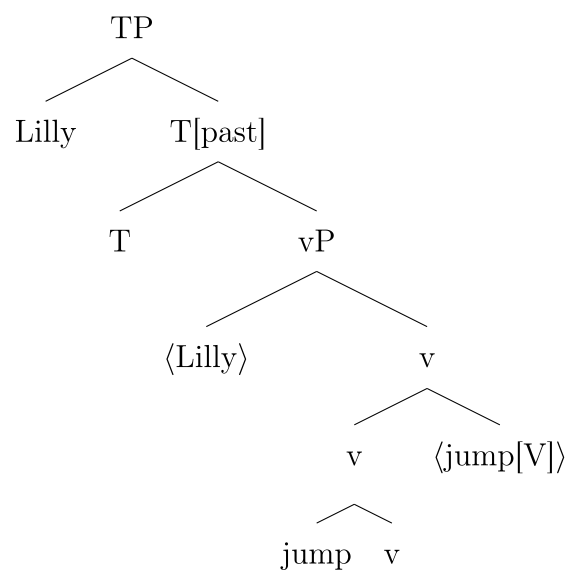

The syntactic structure which (1) is a representation of, however, is something which is not on this page. The structure is a mental event of some kind, a momentary configuration of brain matter, what Adger (2019) calls "a gesture of the mind". When I make a gesture with my hand, it is made possible by the nature of my anatomy and made actual by my intentions. When I utter or comprehend *Lilly jumped*, I make a gesture of my mind, made possible by the anatomy of my brain, and made actual by what I intend to express (or by the linguistic expression I reflexively perceive). (1) represents this gesture of the mind.

Further, (1) is a representation of a *type* of mental event, as opposed to a particular token of it. Linguistic theory abstracts away from particularities of actual mental events so as to come to generalizations about the consistent and recurrent patterns in linguistic data: the relations between meaning and form that are the bedrock of evidence as to the nature of the human linguistic capacity. This level of abstraction to types of events is crucial to motivating, evaluating and accepting or rejecting theoretical proposals.

The typical term used to refer to this abstraction from particular tokens of mental events is 'syntactic structure'. (1) is a representation of a syntactic structure.The generativity of human syntax means that our explanation of language is not simply a list of such structures: it is a generative theory (a computational procedure). Further, as has been clear since at least Chomsky (1957a), generative linguistic theory is not a theory of linguistic patterns, much less of linguistic behaviour. It is a theory of the cognitive system that is (partly) responsible for those patterns or indirectly contributes to the behaviour. Behaviour is evidence for the system, not a primary object of study.

Turning back to (1), it is a representation of a syntactic structure that is legitimized by a particular computational procedure (the mental grammar). Individuals in possession of that computational procedure have the ability to run a mental computation which results in an event that can be seen as a token of the syntactic structure.

This three level distinction (representation; structure; tokened events) is not usually relevant for the daily work of a syntactician, who is concerned to provide accounts of ergativity, long-distance dependencies, passives, or whatever. This leads to slippage in usage: we talk of syntactic representations meaning syntactic structures. But what our theories are about is (abstractions over) configurations of the mind, not representations on the page.

Whereas a representation has intentionality, a structure does not. A representation like (1) is about something in the mind of a human, as just discussed. But the structure (1) represents is not itself a representation of anything.

It is in this sense that the use of the word ‘representation’ is similarly ambiguous to the the use of the word ‘grammar’. The linguist’s grammar is about the mental grammar. It is the latter that is responsible for the fact that humans use structures like (1) when they speak and comprehend sentences, and not structures like (2) or (2).

(2)

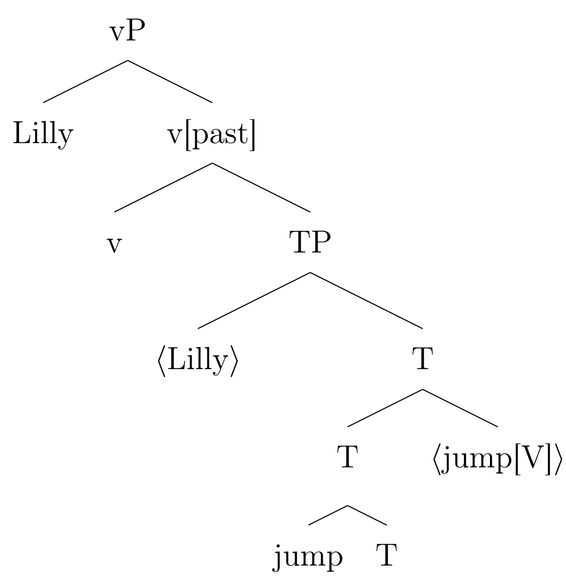

(3)

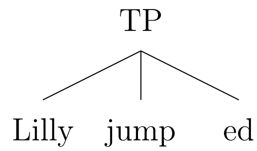

Linguists' representations and theories do have intentionality, in a trivial sense, because they are human constructions. The mental grammar and its effects in time and space do not. Just as the word 'grammar' is used to mean both the linguist's theory, and the mental architecture it characterises, the word 'representation' is sloppily used to mean both the theoretical characterization of a mental structure via a particular diagram and the mental structure itself.

The word 'representation' in this sense is not a relational noun: it is what Grimshaw (1990) calls a result nominal. Think, for example of the word examination in the two following examples:

(4) a. The constant examination of the undergraduate students causes them a lot of stress.
b. The examination that the first year students were set was too easy.

One use of 'examination' here is relational, and eventive. It involves an event of examining and that event is related to the examinees. The second is non-eventive; what is being referred to here by the speaker of (4-b) is a (potentially abstract) object.

Now compare two uses of the word 'representation':

(5) a. The representation of syntactic dependencies via the use of sub-

scripted indices declined after the advent of the copy theory of movement traces.

b. It follows that (30b)-(33b) can be assigned the syntactic representations in (34a)-(34d), explaining the absence of island effects (Kennedy 2003)

The first of these is constructed (it is rather challenging to find natural examples of this use of 'representation' in generative syntax), and 'representation' is clearly used here in a relational way. In (5-b), 'representation is used to mean structure, since it is the structures' properties that are responsible for the absence of island effects. One might say that the language in (5-b) should be reformulated as something like (6):

(6) It follows that the absence of island effects in (30b)-(33b) can be explained if the structures that underlie them have the properties represented in the tree diagrams in (34a)-(34d).

However, Kennedy's intent is obvious, the explanation of the empirical patterns is effective, and the claim about the mental grammar is clear. Practitioners of syntax will typically abstract away from both particular mental events, and issues of representations and what is represented (those mental events qua types).

Georges Rey has, over recent years, made an argument that this terminological slippage is underpinned not by sloppy word use, but rather by a different ontology (e.g. Rey 2003; Rey 2006). Rather than structures simply being abstractions over mental events licensed by computing the procedure that is the mental grammar, he suggests that syntactic structures are relational, representational, and intentional. It is this idea I briefly challenge here.

For Rey, syntactic structures and their components do in fact represent objects. Those objects are what he calls intentional inexistents—they are illusory objects, and syntactic structures represent these (they are about them), in much the same way that words like 'unicorn' are about imaginary horned horses. So syntactic structures and the primitives they are built up out of are, for Rey, intentional: the symbol N in a tree diagram is a (representation of a) representation of an intentional inexistent, the imaginary object noun.

On the alternative view I take, which I outlined above, the issue is about muddy terminology: syntactic 'representations' are not representations with intentionality (construed as an aboutness relation to the external world cf. Chomsky (1995)). They are simply structures/configurations, aspects of which are 'transduced' to/from other cognitive systems (cf. Pylyshyn 1984). This is how syntactic structures connect to the rest of the world, as opposed to via a representation relationship. I sketch the transduction idea below in section 3.

Rey's idea is certainly a more interesting claim than my suggestion that there's only an issue of sloppy terminology at stake. In this very short note I'll only make the argument that Rey's is not a necessary claim: generative linguistics can do without it as there's an interpretation of linguistic theory, which, I think, is the standard one, whereby linguistic structures are not representations of anything (Chomsky 2000). They get their connection to the rest of the world through interfacing with other cognitive systems.

2 Structures and computations

As sketched above, what linguists often (loosely) call a syntactic representation, say a tree, labelled bracketing, set-theoretic structure, set of lines in a derivation etc. is just a configuration, a structure, which is a shape consisting of basic units bearing particular relations to each other, like a gesture of the hand. Such a structure is legitimated (that is, has the shape it has rather than some other kind of shape) by a generative procedure (the generative procedure determines what structures are available). That generative procedure is central to explaining an individual's linguistic abilities: it is a component of the individual's mind, ultimately brain. So if we think of some person, J, then J's knowledge of language is a brain state of J, characterised as a generative procedure g. Possession of g allows J to speak and understand her variety of English (Chomsky 1987).

What is a generative procedure?

A procedure is just a compression of an infinite number of computations (Gallistel and King 2011). The procedure for addition, for example, can be given in finitary form, using, say, a Turing Machine, but will apply across an unbounded domain of numbers. Symbols in a procedure are manipulated, when the procedure is computed, without regard to what those symbols mean, or represent.

A generative procedure does the same: it is written in a finitary form, but compresses an unbounded number of computations. For human syntax, each computation generates a structure. The syntactic representation in (1) is a representation of the kind of structure that is generated. Qua representation, it captures many of the relevant properties of the computation, in a useful static fashion. What (1) represents, however, as stressed above, is an abstraction over a mental event (a computation).

Generative linguistics gets large part of the explanation of the core phenomena that fall within the purview of the theory (hierarchy, lack of numerical bound to hierarchy, discreteness, legitimacy of certain dependencies and not others, etc.) from how it specifies the generative procedure g. That specification has changed over the history of the theory, as one might expect. Early versions of g took it to involve a phrase structure component, which created labelled hierarchies, augmented with a transformational component, which introduced recursion (Chomsky 1957b); later versions incorporated recursion into the phrase structure component, reserving transformations as a means of capturing mismatches in form-meaning correlations (movement type phenomena, Chomsky 1967); the 1980s saw an understanding of g as a 'generate and test' type system where a very general generative procedure built structures that were then filtered out by a separate component; recent versions have generalized the transformational component to allow it to build structure, eliminating phrase structure, and unifying the explanation of local phrase structure building with that of movement phenomena in a single finitary generative procedure, Merge. All of these are hypotheses about the nature of the generative procedure aspects of which are universal to our species, but which develops in different ways in different individuals depending on their linguistic experience.

A generative procedure like Merge specifies the range of possible syntactic computations (usually called derivations) over a finite set of inputs (the lexical items of a particular language): a syntactic derivation is just a theoretically abstract way of characterising a physical configuration of matter and energy in space and time, where that configuration may be used in the mental life of some individual. A linguist writing down such a derivation in an article is no different to a chemist writing H_2O to characterise a physical configuration of matter and energy in space and time which may be used as a mixer for whiskey. No intentionality is necessary in either case (though abstraction to the type level is necessary in both cases for generality of explanation).

Note that I'm not dismissing intentionality. I think intentionality is a deep puzzle about human beings. I'm simply saying that it is not an issue within the domain of syntax. Intentionality is not part of the theory that explains why grammatical structures are hierarchical, why they are unbounded, why they involve certain kinds of dependencies (like c-command) and not others (like linear adjacency), why ergative subjects are resistant to extraction, why strong cross-over effects are universal while weak cross-over effects are not, why unpronounced NPs behave differently when they are A vs A-bar bound, etc. Intentionality is also, contra Rey, not required as part of a theory of syntax for conceptual or philosophical reasons. If we don't need syntactic structures to be intentional to satisfy the demands of syntactic theory, and we don't need them to be intentional to satisfy philosophical scruples, then they probably aren't intentional.

3 Transductions

If syntactic structures are not intentional, how do they connect with the world? Intentionality is, after all, at its core about the connection between mental structures and the external world within which the mind exists. I've characterised a syntactic structure as something legitimized by a computational procedure. But one might say that typical computational procedures do, in fact, typically involve a representation relation. The procedure for addition, though it operates mindlessly on symbols, is interpreted as meaning addition of numbers. The symbols are intended as representations of the numbers.

The computational procedure Merge operates on symbols too: lexical items with particular structural organization (subcategorization features, ϕ-features, categorial features, etc.). If these symbols don't represent anything, contra Rey, how do they have the effect that they do?

I mentioned above that I think that this involves, at some point, transduction. Here is what Pylyshyn (1984) says about transduction:

> By mapping certain classes of physical states of the environment into computationally relevant states of a device, the transducer performs a rather special conversion: converting computationally arbitrary physical events into computational events. A description of a transducer function shows how certain nonsymbolic physical events are mapped into certain symbolic systems. Pylyshyn (1984) 152

Consider first phonetic interpretation. The human being encounters a physical signal, typically sound waves, but the signal could be visual, or even tactile. In all cases the physical state of the signal must be mapped into computationally relevant states. If this is to be done via transduction, the transducer will need to have the relevant relation built in.

There is interesting evidence that this is the case. Mahmoudzadeh, Dehaene-Lambertz, Fournier, Kongolo, Goudjil, Dubois, Grebe, and Wallois (2013) shows a fascinating differential cortical response among pre-term babies to discrimination between the sounds [b] and [g] versus to discrimination between male and female voices. The acoustic distinction between [b] and [g] is extremely subtle compared to the major acoustic differences between male and female voices. This strongly suggests that there are prelinguistic, but language ready, cortical response effects keyed to extracting phonologically relevant information from the acoustic signal, as opposed to non-phonologically relevant information (such as speaker gender). This is exactly Pylyshyn's notion of transduction. Mahmoudzadeh's team also showed that rats showed the opposite effect. Human brains are wired to extract language relevant information from speech signals.

The acoustic signal involves energy transfer of particular physical magnitudes to some brain activity, providing differential categories for the brain to work with. This is the transducer function. These differentiated categories are then accessible to computation, in this case a computation that manipulates phonetic, phonological, and morphophonological symbols. The computation is sensitive to properties of the symbols, and those properties are, ultimately, grounded in acoustic properties via transduction.

The relation between the phonetic/phonological computation and the syntactic one is fairly distant: the phonological computation involves segmental effects, intonational contours, liaison, fine phonetic articulation signalling social meaning, allomorphy, metathesis, reduplication, syncretism, cliticization, etc. Clearly, this is a complex and abstract computation which has a whole subfield fo linguistics to itself (though typically, syntacticians will abbreviate all this complex computation by referring to 'spellout'). There is also a mapping from the phonological computation to the syntactic one, with symbols of the former inaccessible to the latter (which is why we don't find syntactic rules referring to segmental phonological properties, like voicing, or vowel height).

Pylyshyn in his 1985 discussion is focussed on how cognitive systems encounter external events, and he constrains his notion of transducer to this: a function from physical events onto symbolic ones (this is Pylyshyn's constraint that a transducer is an *interrupt* function, not a *test* function). However, for human language, as well as encountering acoustic or visual signals, the system must also output physical events (as in articulation). I therefore propose extending Pylyshyn's proposal so that transduction is a mapping between physical events and computation more generally. It involves both mapping from physical events, and mapping to physical events (where a physical event has temporal and spatial extent).

For articulation, as opposed to perception, we can think of the phonetic/phonological computation as terminating in instructions to the articulators:

> "From this point of view the features that figure in discussions of phonetics and phonology are instructions for specific actions of

an articulator." (Halle, Vaux, and Wolfe 2000, p434)

Carrying out those instructions, mapping from the phonetic/phonological symbols to actions of the articulators, is a function from the output of the phonetic/phonological computation to physical events (which have magnitude, length, temporal status, etc., properties absent from the output of the computation). Rather than transduction from an acoustic signal, it is transduction to a physical instruction.

This same view can be applied to the other side of the computation, the mapping to the systems of the mind that are concerned with planning, thoughts, concepts, etc. Virtually nothing is known about this, but I suggest that the phonetic side of the equation provides a good model. The 'semantic' computation maps to what we might rather clumsily call 'thoughticulators'.

This takes semantic interpretation to map from the output of the syntactic derivation to information that accessible to systems of the brain that are used in various kinds of thinking, providing structure to unstructured thought. The phonetic analogy of instructions to the motor system is apt: syntactic configurations provide instructions to the systems of thought (planning, conceiving, imagining, social thinking, etc.). These systems are thereby boosted by the computational structure that the syntax gives them. Of course, so little is known about this aspect of the use of syntactic derivations by other systems of the mind, that this claim is quite speculative. However, for my purposes here, I want only to show how the generative model of syntax can work without attributing intentionality to the structures is hypothesizes.

Syntactic structures from this point of view connect to syntax-external aspects of the world through further computation that transforms them into structures that can be mapped to acoustic and motor systems on the one hand, and to conceptual and planning systems internal to the brain on the other. Intentionality doesn't figure into the explanation of how syntactic structures are causally efficacious in comprehension or language production.

References

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

Chomsky, Noam. 1957a. *Syntactic Structures*. The Hague: Mouton.

Chomsky, Noam. 1957b. *Syntactic structures*. Mouton.

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

Chomsky, Noam. 1967. The formal nature of language. In Eric Lenneberg, ed., *Biological Foundations of Language*, 397–442, New York, N.Y.: Wiley and Sons.

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

Chomsky, Noam. 1995. Language and nature. *Mind* 104:1–61.

Chomsky, Noam. 2000. Minimalist inquiries: the framework. In R. Martin, D. Michaels, and Juan Uriagereka, eds., *Step by Step: Essays on Minimalist syntax in Honour of Howard Lasnik*, 89–115, Cambridge, MA: MIT Press.

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

Grimshaw, Jane. 1990. *Argument Structure*. Cambridge, MA: MIT Press.

Halle, Morris, Vaux, Bert, and Wolfe, Andrew. 2000. On feature spreading and the representation of place of articulation. *Linguistic inquiry* 31:387–444.

Kennedy, Christopher. 2003. Ellipsis and syntactic representation. In Susanne Winkler and Kerstin Schwabe, eds., *The interfaces: Deriving and interpreting omitted structures,*, 29–53, Amsterdam: John Benjamins.

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.

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. 2006. Conventions, intuitions and linguistic inexistents: A reply to devitt. *Croatian Journal of Philosophy* 6:549–569.