From "Naturalizing Phenomenology" to Formalizing Cognitive Linguistics (I): Cognitive Transform Grammar

Nathaniel Christen

2021/05/24

Abstract

Start with the question of conditions of possibility for language to exist: how spoken, written, or inscribed signs differ from other kinds of things, enough to establish language as something that exists in the world. Language is a regime of multiple conversation partners, an ambient surrounding where they share actions, perceptions, and foci of attention, and a common posit of structured rules. If communication is successful, each sign's referent is isolated, collectively/phenomenologically, from a (canonically spatiotemporal) reality that extends beyond it. My goal here is to sketch this process with the help of formal and Computational linguistic theories, but keeping in sight a cognitive nuance and intersubjectivity that is (thankfully) an essential, unavoidable part of human language. I will incorporate technical models mostly from Link Grammar and Type-Theoretic Semantics, understanding the latter as possibly layering on the former to complete a syntax/semantic pairing. I propose both cases for and limits on formal theories' applicability for Cognitive Linguistics/Humanities/Phenomenology, through the lens of the Philosophy of Science — arguing how Phenomenology and Cognitive Linguistics, while distinct traditions, can work powerfully in consort.

> On connaît la célèbre affirmation de Claude Lévi-Strauss: "les sciences humaines seront structurales ou ne seront pas". Nous aimerions lui en adjoindre une autre: "les sciences humaines seront des sciences naturelles ou ne seront pas". Evidemment, sauf à en revenir à un réductionnisme dogmatique, une telle affirmation n'est soutenable que si l'on peut suffisamment généraliser le concept classique de "naturalité", le généraliser jusqu'à pouvoir y faire droit, comme à des phénomènes naturels, aux phénomènes d'organisation structurale.
> — Jean Petitot, [55, p. 1]

> The nature of any entity, I propose, divides into three aspects or facets, which we may call its form, appearance, and substrate. In an act of consciousness, accordingly, we must distinguish three fundamentally different aspects: its form or intentional structure, its appearance or subjective "feel", and its substrate or origin. In terms of this three-facet distinction, we can define the place of consciousness in the world.
> — David Woodruff Smith, [81, p. 11]

© 2021/05/24 Nathaniel Christen

The question of whether computers can be programmed to understand language may be philosophical, but it overlaps with broad methodological bifurcations: after all, linguists *are* programming computers to "understand" language, at least to some approximation. Given that computational linguistics is now a well-established practice, we can consider how this program for investigating the nature of language orients into linguistics as a whole: to what degree are the computers really "understanding" their linguistic input? How much does *behavior* consistent with language-understanding suggest actual understanding? Is linguistic competence mostly a behavioral phenomenon, or something more holistic and (inter-) subjective? Are the imperfections of automated Natural Language Processing inevitable, and if so, does that foreclose the possibility of NLP engines being considered truly linguistic? That is, should we treat flawed and oversimplistic (but practically useful) NLP software — or

"personas" driven by this software, like "digital assistants" — as bonafide (if rather primitive) participants in the world of human language? Or are they merely machines that simulate linguistic behavior without manifesting real linguistic behavior, as a computer simulation of a celestial galaxy is not a real galaxy?

These are methodological as well as thematic questions. There is a wide swath of formal and computational linguistics, for instance, for which the measure of a theory is its chance of being operationalized on NLP terms and within NLP tools, yielding automated systems whose accuracy and/or computational efficiency competes favorably with other systems. Faithfulness to how *humans* process language is at most a secondary concern. Conversely, there is a broad literature in Cognitive Linguistics and the Philosophy of Language for which uncovering the cognitive and interpretive registers through which *we* understand, produce, and are affected by language is the main goal. For scholars chasing that telos, failure to encode theoretical models in mathematical or software systems is not *prima facie* an explanatory limitation — conversely, we might take this as evidence that cognitive models are addressing the deep, subtle realities of language that are opaque to computer simulation.

Then there is hybrid work, like attempts to formalize Cognitive Grammar (Matt Selway [67], Kenneth Holmqvist [29], [30]), or other branches of Cognitive Linguistics (cf. Terry Regier's influential [62]), or Conceptual Space Theory as initiated by Peter Gärdenfors (which has seen several attempts at mathematical-computational formalization, such as Frank Zenker, Martin Raubal, and Benjamin Adams's metascientific perspectives [1], [2], [27], and more recent Category-Theoretic structures linked to mathematicians such as David Spivak and Bob Coecke [18]). To this list we could add research that extends beyond language alone to broader cognitive-perceptual and conceptual themes, like formal descriptions rooted in Husserlian analyses by phenomenologists whose methods encompass some computer-scientific techniques, such as Barry Smith (as in [12]) and Jean Petitot (see [56]); we can see these accounts as generalizing cognitive-linguistic theories by noting the phenomenological basis of linguistic phenomena, as articulated by (say) Olav K. Wiegand ([78], [79]) and Jordan Zlatev ([84]). In each of the works just cited (prior anyhow to the last three) we can find formal/computational models whose rationale is, in large part, to shed light on human cognitive processes (albeit not necessarily translating to practical NLP components in any straightforward way).

This kind of "intermediate" research is perhaps underappreciated, because it neither accepts the dismissive attitude that formal models are a distraction from the analysis of "real" language, nor the reductionistic faith that language is *intrinsically* computational — that progress toward ideal NLP avatars is just a matter of time. To be sure, layering formal systems on a cognitive/phenomenological foundation adds a complexity of theoretical structure, which could prompt questions about the efficacy of the theoretical dilation: is extra complexity desirable as an end in itself, if the new formalizations have only limited explanatory or practical pay-offs? If there is a human kernel in language that is intrinsically non-computable and non-formalizable, does analysis of language truly benefit from complex but only partly applicable structural overlays? On the other hand, if language *is* computationally tractable, shouldn't NLP implementation be a factor in assessing which formal models are worthy of attention?

Perhaps for these kinds of considerations, linguistics seems to bifurcate between a camp that essentially ignores computational methodology and resources and a camp that centers its whole attention on building better automated NLP tools. Left out of this division is research that invokes formal models as explanatory vehicles while not enmeshing them in an ecosystem oriented toward automation — the difference between deploying formal representations to model (some aspects) of language processing, syntax, and semantics, and trying to program software to *automate* the construction, translation, and pipeline between and among formal models. When situating research relative to computational linguistics, we should bear in mind the metatheoretical point that *incorporating* formal schema into linguistic models does not *necessarily* mean committing ourselves to a task of programming computers to discover the target representations on their own, given raw linguistic input.

But this perspective is not only metatheoretical: I believe that the nature of language is *intrinsically* "hybrid" in a manner that warrants neither blind faith in automation nor *a priori* disengagement with formalizing projects. This is first of all because language is neither wholly isolated from other cognitive phenomena nor without some

© 2021/05/24 Nathaniel Christen

structural autonomy. It is reasonable to suppose that there are distinct intellectual faculties internal to our understanding of language, while other reasonings intrinsic to parsing the form and intentions of a linguistic unit are drawn from the wider inventory of situational, conceptual, social, and practical/enactive cognition. Sentences can vary in terms of their context-sensitivity and the degree to which extralinguistic rationality is implicit in grasping intended meanings. So neither a theory which ignores extralinguistic cognition, nor one which treats *all* linguistic processing as inseparable from the totality of our cognitive processes from moment to moment, are complete. A fully-forged theory needs to place sentences on a spectrum, where extralinguistic and (I'll say) "intralinguistic" theoretical machinery is available to analyze different sentences as their form and context demands.

Extending this point, I believe a comparable spectrum matches the duality of language seen as intrinsically formalizable and computable or as too subtle, social/cultural, embodied, and context-dependent to be tractable to any computer or any idealized logicomathematical abstraction. Some sentence are more logically straightforward; others are more elusive, requiring holistic and context-sensitive interpretation on conversants' parts to be understood. Combined with my claims last paragraph, I argue accordingly that we can (at least as a suggestive picture) view linguistic artifacts (canonically, sentences) along a two-axis spectrum defined both by extralinguistic integration (or lack thereof) and by formal tractability (or lack thereof), like so: (this is intended as an intuitive sketch, not a formal model).

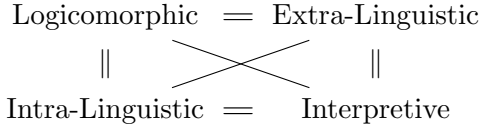

I will elucidate the terms on that picture later. Summarily, though, I claim that *some* sentences evince logically straightforward compositional patterns that can be analyzed *either* at the language-specific (syntactic or semantic) level *or* within cognitive registers outside of language proper (e.g., situational schemata); conversely, some sentences have nuances that call for interpretive judgments which appear to transcend formal simulacra outside the full range of human intelligence, emotion, and embodiment, *either* in terms of parsing complex syntactic or semantic structures *or* in grounding linguistic phenomena in ambient contexts. My overall point is then that sentences take a spectrum of models spanned by these axes; no one paradigm is self-contained as a metalinguistic commitment.

In effect, the choice between paradigms wherein language is or is not formally/computationally tractable, and between paradigms wherein language is or is not intellectually autonomous vis-à-vis our total cognitive faculties, should not be seen as a metaphysical alternative anterior to language as a totality. Instead, these spectra are threaded into language internally, competing polarities which rise or fall from sentence to sentence. Language is not *intrinsically* either formal or non-formal; autonomous or non-autonomous.

But at the same time, sentences are clearly phenomena of the same ilk; the distinctions I have made are not so sharp as to disrupt the ontological similitude among sentences, so that two sentences (however much they differ on the spectra of my diagram) are still manifestations of the same ontological place; they are still roughly the same *sort* of existents. Accordingly, we should conclude philosophically that there are certain aspects of sentences that lie beneath the formal/interpretive and intralinguistic/extralinguistic dualities. There are, in short, paleostructures in language that manifest *either* with formal specificity or with contextual nuance; *either* internal to syntax or semantics or external to intrinsically linguistic cognition; varying from sentence to sentence. This paper will present a theory of one such paleostructure, drawing inspiration both from formal theories (Dependency Grammar, Type-Theoretic Semantics, Generative Lexicon) and from more philosophical approaches (Cognitive Grammar, Semiotics, Phenomenology).

The central element in my analysis is the *conceptual modification* implied or effected by one word in the presence of another word. More precisely, some words' linguistic roles can be analyzed as adding cognitive detail (conceptual and/or perceptual and/or pragmatic) to the ideas or referents signified via other words. The underlying scheme, at this level where the model is undeveloped and thus fairly simple, is close to Dependency Grammar: in lieu of a head/dependent relation we can treat one word (or similar lexical unit) as a *modifier* and the second as a *ground*. I believe this two-pronged picture is not completely vacuous, but is general and underspecified enough to spread over both syntax and semantics, and

© 2021/05/24 Nathaniel Christen

over competing paradigms. I will call the modifier's effect on its ground a *transform*, and the two words together (taking "word" as a convenient designation for lexemes in general) as a *transform pair*.

I will argue that the underlying transform-pair concept is amenable to both more formal and more philosophical/interpretive development. One the one hand, transform-pairs can sometimes be mapped explicitly to dependency or link-grammar pairs, so the theory can be treated as a philosophical preliminary or motivation for Dependency or Link Grammar.[1] Relatedly, transforms themselves can be integrated into type systems (e.g., the kinds of transforms associated with adjectives depend on their ground being typed as a noun), so a theory of transform-pairs can motivate elaboration in a Type-Theoretic context. On the other hand, we can focus on the interpretive and situational nuance often evident in cognitive transformations to find evidence for linguistic phenomena which do not fit neatly into Dependency Grammar or Type-Theoretic (or any other) formalization. These various continuations, which I will present further over the next several sections, try not to foreclose either formal or philosophical/interpretive paradigms. The goal is to trace both formal models of language and alternative models — for which excessive formalization is reductionistic and depends on ad-hoc avoidance of many real-world significations — to a common structural kernel from which both perspectives can be deployed on a case-by-case (e.g., sentence-by-sentence) basis.

In keeping with the perspective that formal models have *some* merit, I have tried to orient the presentation around certain computational-linguistic techniques. Some of my examples are drawn from popular annotated corpora, and I provide models of other examples as processed by representative NLP technologies (e.g., Malt Parser trained against the most recent Universal Dependency training data at the time of writing). I also select multiple examples from the Blackwell *Handbook of Pragmatics* [32] (apart from these examples, a full sample-collection from that volume, and code I used to extract it, are available on request). From these sources (as well as my own novel demonstrative and hypothetical example-cases addressing various cognitive-linguistic themes) I have packaged the samples and supporting code into an open-source data set for purposes of demonstration.

Although some of these examples emerge from an NLP-inspired paradigm, I do not dwell on the accuracy of the NLP components, in part because I am motivated in this paper to examine how computational methods can be employed as explanatory tools separate and apart from their feasibility in automated pipelines. That is, the computational resources I present here are designed more as technological supplements to philosophical, interpretive, and speculative examinations of interesting linguistic examples, rather than as algorithms ultimately targeted at fully automated NLP frameworks. I incorporate code and data as a supplement to my argumentation in the hope that this can serve as an example of computational linguistics adopted outside the priorities of NLP automation and Artificial Intelligence in its more reductive, science-fictional sense.

In short, I am not aspiring to develop code or theoretical claims that could advance the hypothetical project of implementing artificial agents that can mimic and understand human language and behavior in all its subjectivity and complexity. However, I *do* want to leverage certain computational techniques as offering their own explanatory perspectives on structures in human language.

The remainder of this paper will draw and expand on the outline of terms and structures sketched thus far, specifically the contrast of intra/extra-linguistic and "logicomorphic"/interpretive aspects from my "diagram" of sentences' paradigmatic spectrum, and the basic modifier/ground transform-pair account. I use the phrase Cognitive Transform Grammar for the core notion of "transform-pairs" as, at core, cognitive phenomena, which nonetheless allow for further exposition via different methods. I will explain some of this variation in the first section.

[1] Dependency and Link Grammar both build syntactic structures from inter-word pairs (rather than larger phrasal units), so the latter is at least in part a species of the former. In lieu of head/dependent relations, however, Link Grammar identifies each linguistically meaningful connection between words (and lexical or constituent sentence units in general) as motivated by a specific connection rule, which depends on compatible connecting potentiality present on both words in a connection, which are part of their lexical profile. Formally then Link Grammar removes one theoretical structuring posit (the "ranking" of pair-elements as head or dependent) but adds a different posit related to the two "potentialities" which get fused into a single link; linkage capacities become rooted in the lexicon arguably more so than in (conventional) Dependency Grammar [66], [68], [28], [50].

© 2021/05/24 Nathaniel Christen

1 Cognitive Transform Grammar and Transform-Pairs

The idea that inter-word pairs are a foundational linguistic unit — from which larger aggregates can be built up recursively — is an central tenet of Dependency Grammar. Here I will generalize this perspective outside (but not excluding) grammar, to overall semantic, pragmatic, and even extralinguistic relations indicated via interword relations.

In some cases word-relations can still be theorized mostly via syntax. Consider hypothetical, example sentences like:

- (1) His having lied in the past damages his credibilty in the present.
- (2) Voters question whether he is truthful this time around.

In (1), *having* is necessary to syntactically transform its ground *lied* from a verb-form to a noun (something which can be inserted into a possessive clause). Analogously, in (2) *whether* modifies *is* (since this is the head of a subordinate clause), wrapping a propositional clause into a noun so that it furnishes a direct object to the verb *question*. The essential transformation in these cases is motivated by grammatic considerations, particularly part-of-speech: a verb and a subordinate, propositionally complete clause (in (1) and (2), respectively) need (for syntactic propriety) to be modified so as to play a role in a site where a noun is expected (in effect, they need to be bundled into a noun-phrase).

The relevant transforms here — signified by *having* and *whether* — have a semantic dimension also, and we can speculate that the syntactic rules (requiring a verb or propositional-clause to be transformed into a noun) are actually driven by semantic considerations. Conceptually, for example, *his having lied* packages a verb into a possessive context because the sentence is not foregrounding a specific lying-event but rather the fact of the existence of such occasions. We cannot perhaps "possess" an event, but we possess (as part of our nature or history) the fact of past occurrences, viz., events in the form of things we have done. In this sense *his having lied* marks a conceptual transformation, from events qua occurrents to events (as factical givens) qua states or possessions, and the grammatical norm — how we cannot just say "his lied" — is epiphenomenal to the conceptual logic here; the erroneous "his lied" sounds flawed because it does not match a coherent conceptual pattern in how events and states fit together. But, still, the syntactic requirement — the expectation that a noun or noun-phrase serve as the ground of a possessive adjective, or the direct object of a verb — manifests these underlying conceptual patterns in the order of everyday language.[2] Syntactic patterns become entrenched *because* they are comfortable translations of conceptual schema, but *as* entrenched we hear these patterns as grammatically correct, not just as conceptually well-formed. Likewise, we hear errata like "his lied" as *ungrammatical*, not as conceptually incongruous.

I contend, therefore, that many conceptually-motivated word-pairing patterns become syntactically entrenched and thus engender a class of transform-pairs where the crucial, surface-level transformation is syntactic, often in the form of translations between parts of speech, or between morphological classes (singular/plural, object/location, etc.). Consider locative constructions like:

- (3) Let's go to Grandma.
- (4) Let's go to the lawyers.
- (5) Let's go to the press.

Here nouns like *Grandma*, *the lawyers*, and *the press* are used at sites in the surrounding sentence-forms that call for a designation of place — this compels us to read the nouns as describing a place, even while they are not intrinsically spatial or geographical (e.g., Grandma is associated with the place where she lives).

In (5) and perhaps (4), this locative figuring may be metaphorical: going *to the press* does not necessarily mean going to the newspaper's offices. Indeed, each of these usages are to some degree conventionalized: going *to Grandma* is subtler than going *to Grandmas house*, because the former construction implies that you are

[2] Using terminology generally associated with *structural* linguistics, we can also say that expectations provide the conceptual dynamics underlying *syntagmatic* relations between linguistic units — syntagmatic in the sense that semantics are associated with aggregate structures of units (not single units in isolation) but also that, in contrast e.g. to phrases which have fixed constructions (like *determiner* plus *noun* being a general pattern which, when matched, signifies a concretized noun-concept), syntagmatic structure is built up *dynamically*. We do not passively match constructions to patterns, on this account, but rather construction-patterns *emerge* from the dynamics, tension, and expectations latent in cross-word or (in general) cross-unit syntagmatic relations (see [71] for a related "structural" reading of Cognitive Grammar).

© 2021/05/24 Nathaniel Christen

going to a *place* ("Grandma" is proxy for her house, say), but also that Grandma is actually there, and that seeing her is the purpose of your visit. In other words, the specific *go to Grandma* formation carries a supply of situational expectations. There are analogous implications in (4) and (5) — going *to the press* means trying to get some news story or information published. But the underlying manipulation of concepts, which structures the canonical situations implicated in (3)-(5), is organized around the locative grammatical form as binding noun-concepts to a locative interpretation. However metaphorical or imbued with additional situational implications, a person-to-location or institution-to-location mapping is the kernel conceptual operation around which the further expectations are organized. Accordingly, the locative case qua grammatic phenomenon signals the operation of these situational conventions, and the syntactic norms in turn are manifest via word-pairings, such as *to Grandma*.

In short, a transform-pair like *to Grandma* can be analyzed in several registers; we can see it as the straightforward syntactic rendering of a locative construction (via inter-word morphology, insofar as English has no locative case-markers) or explore further situational implications. In these examples, though, there is an obvious grammatic account of pairs' transformations, notwithstanding that there are also more semantic and conceptual accounts. Part-of-speech transforms (like *having lied*) and case transforms (like *to Grandma*) are mandated by syntactic norms and therefore can be absorbed into conventional grammatic models, such as Dependency Grammar: the head/dependent pairings in *having lied* and *to Grandma* are each covered by specific relations within the theory's inventory of possible inter-word connections. So a subset of transform-pairs overlaps with (or can be associated with) corresponding Dependency Grammar pairings.

Another potential embedding of transform-pairs into formal models can be motivated by Type Theory. Such analysis may proceed on several levels, but in general terms we can assume that parts of speech form a functional type system (as elucidated, say, in Combinatory Categorial Grammar; e.g., [11]). For instance, we can recognize nouns and propositions (sentences or sentence-parts forming logically complete clauses) as primitive types, and treat other parts of speech as akin to "functions" between other types. A verb, say, combines with a noun to form a proposition, or complete idea: *go* acts on *We* to yield the proposition *We go*. Schematically, then, verbs are akin to functions that map nouns to propositions. Similarly, adjectives map nouns to nouns, and adverbs map verbs to other verbs (here I use "noun" or "verb" to mean a linguistic unit which functions (conceptually and/or vis-à-vis syntactic propriety) as a noun, or verb; in this sense a noun-phrase is a kind of noun — i.e., a linguistic unit whose *type* is nominal). This provides a type-theoretic architecture through which transform-pairs can be analyzed. An adverb modifies a verb; so an adverb in a transform pair must have a verb as a ground. Moreover, the "product" of that transform is also a verb, in the sense that the adverb-verb pair, parsed as a phrase, can only be situated in grammatic contexts where a verb is expected.

In effect, we can apply type-theoretic models to both parts of a transform-pair and to the pair as a whole, producing structural requirements on how words link up into transform-pairs. We can then see an entire sentence as built up from a chain of such pairs, with the rules of this construction expressed type-theoretically. Given, say, *his having lied flagrantly* we can identify a chain of pairs *flagrantly-lied*, *having-flagrantly*, and *his-having*, where the "outcome" of one transform becomes subject to a subsequent transform. So *flagrantly* modifies *lied* by expressing measure and emphasis, adding conceptual detail; grammatically the outcome is still a verb. Then *having*, as I argued earlier, applies a transform that maps this verb-outcome to a noun, which is then transformed by the possessive *his*. Each step in the chain is governed by type-related requirements: the output of one transform must be type-compatible with the modifier for the next transform. This induces a notion of *type-checking* transform-chains, which is analogous to how type-checking works in formal settings like computer programming languages, Typed Lambda Calculus, and Dimensional Analysis.

This gloss actually understates the explanatory power of type-theoretic models for linguistics, since I have mentioned only very coarse-grained type classifications (noun, proposition, verb, adjective, adverb); more complex type-theoretic constructions come into play when this framework is refined to consider plural/singular, classes of nouns, and so forth, establishing a basis for more sophisticated structures adapted to language from formal type theory, like type-coercions and dependent types (I will revisit these theories in a later section). Here, though, I

© 2021/05/24 Nathaniel Christen

will just point out that Dependency Grammar and Type-Theoretic Semantics can often overlap in their analysis of word-pairs (inter-word relations is not centralized in type-oriented methodology as much as in Dependency Grammar, but type concepts can certainly be marshaled toward word-pair analysis).

Even though Dependency and Type-Theoretic analyses will often reinforce one another, they can offer distinct perspectives on how pairs aggregate to form complete phrases and sentences. In the transform-pair *having lied*, *lied* is clearly the more significant word semantically. This is reflected in *having* being annotated (at least according to the Universal Dependency framework) as auxiliary, and the dependent element of the pair, while *lied* is the "head". Then *lied* is also connected to *his*, establishing a verb-subject relation. So *lied* becomes the nexus around which other, supporting sentence elements are connected. This is a typical pattern in Dependency Grammar parses, where the most semantically significant sentence elements also tend to be the most densely connected (if we treat the parse-diagram as a graph, these nodes tend to have the highest "degree", a measure of nodes' importance at least as this is reflected in how many other nodes connect to it). Indeed, by counting word connections we can get a rough estimation of semantic importance, distinguishing "central" and "peripheral" elements. These are not standard terms, but they suggest a norm in Dependency Grammar that the structure of parse-graphs generally reflects semantic priority: the central "spine" of a graph, so to speak, captures the primary signifying intentions of the original sentence, while the more peripheral areas capture finer details or syntactic auxiliaries whose role is for grammatical propriety more than meaningful content.

Conversely, a type-theoretic analysis might incline us to question this sense of semantic core versus periphery: in the case of *his having lied*, the transform *having* supplies the outcome which is content for the possessive *has*. If we see the sentence as a cognitive unfolding, a series of mental adjustments toward an ever-more-precise reading of speaker intent, then each step in the transformation contributes consequential details to the final understanding. Moreover, *lied* is only present in the transformation signified by *his* insofar as it has in turn been transformed by *having*: each modifier in a transform-pair has a degree of temporal priority because *its* effects are directly present in the context of the following transformation. This motivates a flavor of Dependency Grammar where the head/dependent ordering is inverted: a seemingly auxiliary component (like the function-word to a content-word) can be notated as the head because its output serves as "input" to a subsequent transform. In the analogy to Lambda Calculus, *his having lied* would be graphed with *having* being the head for *lied*, and *his* the head for *having*, reflecting the relation of functions to their arguments. In lisp-like code, this could be written functionally as (his (having lied)), showing *having* as one function, and *his* as a second one, the former's output being the latter's input (later I will include diagrams contrasting these different styles of parse-representation).

Implicitly, then, Type-Theoretic Semantics and Dependency Grammar can connote different perspectives on semantic importance and the unfolding of linguistic understanding. I will explore this distinction further below, with explicit juxtaposition of parse graphs using the two methods. I contend, however, that the distinction reflects a manifest duality in linguistic meaning: we can treat a linguistic artifact as an unfolding process or as a static signification with more central and more peripheral parts. Both of these aspects coexist: on the one hand, we understanding sentences via an unfolding cognitive process; on the other hand, this cognition includes forming a mental review of the essential points of the sentence, a collation of key ideas such as (for 1) *his*, *lied*, *damages*, and *credibility*. Given this two-toned cognitive status — part dynamic process, part static outline — it is perhaps understandable that different methodologies for destructuring a sentence into word-pair aggregates would converge on different organizational norms for how the pairs are interrelated, internally and to one another.

This analysis, which I will extend later, has considered transform-pairs from a syntactic angle — in the sense that I have highlighted pairs which obviously come to the fore via grammatic principles. As I indicated, I believe the notion of transform-pairs cuts across both syntax and semantics, so I will pivot to some analyses which attend more to the semantic dimension.

1.1 Semantic Analyses of Transform-Pairs

In the simplest cases, a transform-pair represents a

© 2021/05/24 Nathaniel Christen

modifier adding conceptual detail to a ground, like *black dogs* from *dogs*. But the nature of this added detail — and its evident relation to surface language — can be highly varied, even among similar sentences at the surface level. Compare between examples like:

- (6) I saw my neighbor's two black dogs.
- (7) I saw my neighbor's two rescued dogs.
- (8) I saw my neighbor's two latest dogs.

Whereas (6) presents a fairly straightforward conceptual transformation, the detailing in (7) is a lot subtler; mentioning *rescued* dogs makes no reference to perceptual qualities, but rather implies intricate situational background. The term *rescued dogs* strongly suggests that the dogs were adopted by their current owner, probably after an animal-welfare organization found them abandoned, or removed them from a prior abusive owner. This kind of backstory is packaged up, as a kind of situational prototype, in the conventionalized phrase *rescued dogs*, implying a level of specificity more precise than the adjective *rescued* alone implies. Correspondingly, the verb *to rescue* when applied to dogs suggest more information than in more generic contexts.

The phrase *latest dogs* carries implications in its own right; we assume the neighbor had owned other dogs before. Of course "latest" implies some temporal order, but the understood time-scale depends on context. If we hear talk about a *vet*'s two latest dogs, we would presumably interpret this in terms of patients the vet has seen over the course of a day:

- (9) We have to wait until after the vet's two latest dogs.
- (10) I'm concerned for the rescue organization's two latest dogs.

Understanding the relevant time-frame depends on understanding the relation between the dogs and the possessive antecedent. In (8) the neighbor (in a typical case) actually owns the dogs, so the situational context grounding the modifier *latest* would be understood against the normal time-scale for dog ownership (at least several years). In (9), the vet only "possesses" the dogs in the sense of endeavoring to examine them, a process of minutes or hours. In (10), the implication of the *organization's* possessive vis-à-vis rescued dogs is that the group endeavors to rehabilitate and find permanent homes for the rescuees. So in each case *latest* implies a succession of dogs, leading over time to two most recent ones, but the implied time-frame for our conceptualizing this sequence can be minutes-to-hours, or days-to-months, or years.

We should also observe that the implied time-frames and backstories in (7)-(10) are not directly signified via morphosyntax or lexical resources alone. The word *rescued* only carries the *rescued dog* backstory when used in a context involving the dogs' eventual owners; in some context the more generic meaning of *rescue* could supersede:

- (11) Boatmen rescued dogs from the flooded streets.
- (12) Firemen rescued dogs from the burning building.

Neither (11) nor (12) imply that the dogs were abandoned, or will have new owners, or be sent to a shelter, or that their rescuers are members of an animal-welfare organization — in short, no element of the conventionalized backstory usually invoked by *rescued dogs* is present. Analogously, there is no lexical subdivision for *latest* which regulates the variance in time-frames among (8)-(10). It is only by inferring a likely situational background that conversants will make time-scale assumptions based on one situation involving dog ownership, another involving veterinary exams, and a third involving animal-welfare rehabilitation.

That is to say, the time-scale inference I have analyzed is essentially *extralinguistic*: there is no specific *linguistic* knowledge (lexical or grammatical, or even pragmatic inferences in the sense of deictic or anaphora resolution) which warrants the situational classification of (8)-(10) into different time scales. Instead, the inference is driven by (to some degree socially or culturally specific) background-knowledge about phenomena like veterinary clinics or animal rescue groups. Whether or not the nuances in *rescued dogs* are similarly extralinguistic is an interesting question — we can argue that the phrase is now entrenched as a *de facto* lexical entrant in its own right, so the role of *rescued* is not only to lend adjectival detail but to construct a recurring phrase with a distinct meaning, like *red card* (in football) or *stolen base* (in baseball). Lexical entrenchment is, I would argue, an intra-linguistic phenomenon, in the sense that understanding entrenched phrases is akin to familiarity with specific word-senses, which is a properly linguistic kind of knowledge. But even in that case, entrenchment is only possible because the phrase has

© 2021/05/24 Nathaniel Christen

a signifying precision more rigorous than its purely linguistic composition would imply. There are, in short, extralinguistic considerations governing *when* phrases are candidates for entrenchment, and a language-user's ability to learn the conventionalized meaning (which I believe is an intra-linguistic cognitive development) depends on their having the relevant (extra-linguistic) background knowledge.

If we consider then the contrast between transform-pairs like *black dogs*, *rescued dogs*, and *latest dogs*, the similar grammatic constructions — indeed similar semantic constructions, in that each pair has an adjective modifying a straightforward plural noun (*dogs* designates a similar concept in each case; this is not a case of surface grammar hiding semantic diversity, like *strong wine* vs. *strong opinion* vs. *strong leash* or *long afternoon* vs. *long history* vs. *long leash*) — package transforms whose cognitive resolution spans a range of linguistic and extralinguistic considerations. Straightforward adjectival modification in *black dogs* gives way to lexical entrenchment in *rescued dogs* which, as I argued, carries significant extra-linguistic background knowledge even though possession of this knowledge is packaged into basic linguistic familiarity with *rescued dogs* as a signifying unit; and in the case of *latest dogs* the morphosyntactic evocation of temporal precedence and two different multiplicities (the latest dogs and earlier ones) is fleshed out by extra-linguistic estimations of time scale. The same surface-level linguistic structures, in short, can (or so such examples argue for) lead conversants on a cognitive trajectory in which linguistic and extra-linguistic factors interoperate in many different ways.

This diversity should call into question the ability of conventional syntactic and semantic analysis to elucidate sentence-meanings with any precision or granularity. Lexical and morphosyntactic observations may certainly reflect details which *contribute* to sentence-meanings, but the overall understanding of each sentence in context depends on holistic, interpretive acts by competent language users in light of extra-linguistic, socially mediated background knowledge and situational understanding. Contextuality applies here not only in the pragmatic sense that pronoun resolution, say, depends on discursive context (who is *her* in *her dogs*); more broadly, transcending even pragmatics, context describes presumptive familiarity with conceptual structures like veterinary clinics, animal shelters, and any other real-world domain which provides an overall system wherein particular lexical significations can be standardized. Without the requisite conceptual background it is hard to analyze how speakers can make sense even of well-established variations in word-sense, like *treat* as in a veterinarian treating a dog, a doctor treating a wound, a carpenter treating a piece of wood, or how an actor treats a part. These senses have lexical specificity only in the domain-specific contexts of medicine, carpentry, theater, and so forth.

The problem of holistic cognitive interpretation (as requisite for sentence-meanings) can be seen even more baldly in examples where semantic readings bifurcate in ways wholly dependent on extra-linguitic conceptualization. Consider for instance:

- (13) All New Yorkers live in one of five boroughs.
- (14) All New Yorkers complain about how long it takes to commute to New York City.
- (15) The south side of Cambridge voted Conservative.
- (16) Lower Manhattan voted Republican.
- (17) Staten Island voted Republican.

Sentence (15) is taken from the *Handbook of Pragmantics* (example 40, page 346/379, chapter 15 — note that when referencing the *Handbook* I will cite the volume "print" page number followed by the PDF page number) which borrows in turn from Ann Copestake and Ted Briscoe. In the *Handbook* analysis (chapter by Geoffrey Nunberg), (15) is seen as ambiguous between reading "The south side of Cambridge" as an oblique description of the *voters* in that territory or as topicalizing the territory itself as a civic entity:

> On the face of things, we might analyze ([their example] 40) in either of two ways: either the description within the subject NP has a transferred meaning that describes a group of people, or the VP has a transferred meaning in which it conveys the property that jurisdictions acquire in virtue of the voting behavior of their residents.

The duality is significant because the designation in "south side of Cambridge" would be more informal in the prior reading, both geographically and in the how the implied collective of people is figured. The prior reading accommodates a hearing wherein the speaker

© 2021/05/24 Nathaniel Christen

construes "south side" not as a precise electoral district (or districts) but as a vaguely defined part of the city. That imprecision also allows the claim "voted Conservative" to be only loosely committal, implying that some majority of voters appeared to vote Conservative but not that this tendency is directly manifest in election results. In short, we can interpret (15) as making epistemically more rigorous or more noncommittal claims depending on how we read the geographical reference "south side of Cambridge" (as crisp or fuzzy), the group of people selected via that reference (mapping the region to its inhabitants, a kind of "type coercion" since places do not vote), and the assertive force of the speech-act: how precisely the speaker intends her claim to be understood. Each of these "axes" contribute to the sentence's meaning insofar as they constrain what would be dialogically appropriate responses.

Meanwhile, in (13), *New Yorkers* refers specifically to everyone who lives in the City of New York, since the five boroughs collectively span the whole of city. In (14), by contrast, we should understand *New Yorkers* as referring to residents of the metropolitan area *outside* the city itself (who commute *to* the city); and moreover *All* should be read less then literally: we do not hear the speaker in (14) committing to the proposition that *every single* New Yorker complains. So both *All* and *New Yorkers* have noticeably different meanings in the two sentences.

And yet, I cannot find any purely linguistic mechanism (lexical, semantic, syntactic, morphological) which would account for these difference as linguistic signifieds *per se*: the actual differences depend on conversants knowing some details about New York (or, respectively, Cambridge) geography, and also general cultural background. It does not make too much sense to commute to a place where you already live, so our conventional picture of the word *commute* constrains our interpretation of (14) — but this depends on *commute* having a specific meaning, of traveling in to a city, usually from a suburban home, on a regular basis; a meaning in turn indebted to the norms of the modern urban lifestyle (it would be hard derive an analogous word-sense in the language spoken by a nomadic tribe, or a pre-industrial agrarian community). Likewise, reading *All* in (13) as *literally* "all" depends on knowing that the five boroughs are in fact the whole of the city's territory. I am from New York, not Cambridge; perhaps residents of the latter city would clearly read "south side" as referencing a fixed civic/electoral area (like *Staten Island*) or as only vaguely defined (like *Lower Manhattan*). For New Yorkers, (16) would be read as fuzzy and (17) as fixed: the latter sentence has a clearly prescribed fact-check (since Staten Island is a distinct electoral district) which the former lacks.

Given that in everyday speech quantifiers like *all* or *every* are often only approximate — and that designations like *New Yorker* are often used imprecisely, with not-identical alternative meanings intended on a case-by-case basis — these kind of examples point to signifying ambiguities that can easily arise as a consequence. Often extra-linguistic considerations resolve the ambiguity by rejecting one or another (otherwise linguistically plausible) reading as non-sensical. Consider:

- (18) The Leafs failed to beat the Habs for the first time this year.
- (19) The Leafs failed to win two consecutive games for the first time this year.
- (20) The Leafs failed to score a goal for the first time this year.

Sentence (18) has two competing readings: either the Toronto Maple Leafs won *all* or *none* of their previous games, in the relevant year, against the Montreal Canadiens. The difference is whether *for the first time this year* attaches to *beat* or to *fail*. In (19), on the other hand, the only sensible interpretation is that the Leafs had not yet won two games: while it is logically accurate to describe a team on a long winning streak as repeatedly winning two consecutive games, it would be very unexpected for (19) to be used in a case where the Leafs lost for the first time, after a three-plus-game winning streak. And in (20) any hockey fan would hear that the Leafs had scored at least one goal in all prior games; even though there is no linguistic rule foreclosing the reading such that the Leafs have not scored a goal in *any* game.

These variations — the degree to which superficial ambiguity is actually perceived by competent language-users as presenting competing plausible meanings — depend on background factors; the contingencies of hockey fix how potential ambiguities resolve out because one or another alternative is extralinguistically incoherent. But these cases point to how linguistic criteria alone, no matter how broadly understood, cannot necessarily predict in what sense linguistic structurations have empirically

© 2021/05/24 Nathaniel Christen

plausible meanings — or whether they have sensible meanings at all.

Notice however that all these examples have alternate versions which are less subtle or ambiguous, which shows that the complications are not localized in the communicated ideas themselves, but in their typical linguistic encoding:

- (21) All residents of the city of New York live in one of five boroughs.
- (22) Many residents of the New York metropolitan area complain about how long it takes to commute to New York City.
- (23) All districts on the south side of Cambridge voted Conservative.
- (24) For the first time this year, the Leafs failed to beat the Habs.

These versions are more logically transparent, in that their propositional content is more directly modeled by the structure of the sentences. Indeed, hearers unfamiliar with New York (respectively Cambridge) or with hockey might find these versions easier to understand; more context-neutral and journalistic. But perhaps for this reason the "journalistic" versions actually sound stilted or non-idiomatic for everyday discourse.

Geoffrey Nunberg's chapter (15) in the *Handbook* also has interesting examples that point to extralinguistic factors determining proxies and substitutions in referential contexts:

- (25) (example 3, page 346/369)I am parked out back.
- (26) (example 5, 347/370)I am parked out back and have been waiting for 15 minutes.
- (27) (example 6)I am parked out back and may not start.
- (28) (page 348/371)I'm in the Whitney Museum.
- (29) (352/375)I drank two beers.
- (30) (352/375)I drank two Michelobs.
- (31) (352/375)I drank two Sauternes last night.

Someone says "I am parked" as proxy for their car, or "in the Whitney" as proxy for their art work. Nunberg questions the acceptability of (27) on the premise that once we establish talking of people as an oblique reference to their cars, it sounds awkward to switch to content that can *only* apply to the car, like *may not start*. He also argues that we interpret (29) and (30) as references, quite possibly, two glasses or bottles of *the same* beverage, while (31) implies two samples of *different* wines. The analysis turns on conventions for creating referring expressions; in particular, he argues, extralinguistic factors (albeit not using this term) regulate when referential substitutions as exemplified in (25)-(31) are comfortable. We accept an artist proxying her work in a museum, but (28) arguably does not generalize to a case like (32), and (26) does not generalize to (33):

- (32) (said by an author whose books are there)I'm in the library.
- (33) I've been in the auto shop for the last 15 days.

Nunberg attributes these discrepancies to cultural conventions, like the difference in prestige between having a painting in a museum compared to having one copy of one's book in a library; or chugging a cheap beer versus savoring a celebrated wine. Here too there are logically transparent alternatives:

- (34) My car is parked out back and have been waiting for 15 minutes.
- (35) I drank two pints of beer last night.
- (36) I tried two different Sauternes last night.

Again, though, the seemingly simpler versions — whose forms generalize more readily to variant cases — seem *less* natural or fluent, as if their very generalizability makes them sound awkward compared to logically more opaque, but idiomatically more sociable, renderings. Impersonal, journalistic language can sound rather cold or unfriendly in contexts where speakers expect a dialect register characteristic of conversations among peers.

In short, even if sentences have a basically transparent logical content, *how* sentences holistically signify this content does not always emerge straightforwardly from semantic or syntactic structures on their own. I think this weakens the case for semantic paradigms that concentrate on logically-structured content which appears to be signified through sentences — even if we grant that this propositional ground of meaning is real, it does not follow that propositional contents are designated by purely linguistic means, rather than by a cohort of cognitive processes many of which are extra-linguistic. This is the basis of my proposing "logicomorphic" qualities as one axis for evaluating sentences, which I will now discuss further.

© 2021/05/24 Nathaniel Christen

1.2 Gaps in Logical Phrase-Models

Assume we have a baseline lambda-calculus-like functional summary of sentences and derived types. That is, any sentence can be rewritten as if a sequence of "function calls", assuming an underlying representational vocabulary of a typed Lambda Calculus, with sentences having overall "proposition" types (I will present this model in detail next section).

My overall goal is to embrace a hybrid methodology — accepting formal analyses when they shed light on linguistic processes, but not going so far as to treat logical, mathematical, or computational models as full explanations for linguistic rationality qua scientific phenomenon. Cognitive Grammar, in particular, challenges our assumption that grammar and semantics are methodologically separate. Received wisdom suggests that grammar concerns the "form" of sentences whereas semantics considers the meaning of words — implicitly assuming that *word combinations* produce new meanings, and that the *order* by which words are combines determines how new meanings are produced. This notion, in turn, is allied with the essentially logical or propositional picture of signifying via doxa: the idea that inter-word relations cue up different logically salient transformations of an underlying predicate model. Thus — to initiate a case-study I will return to several times — *many students* as a phrase is more significatorily precise than *students* as a word, because the phrase (with intimations of quantitative comparison) has more logical detail. Similarly *many students complained* is more logically complete because, provisioning both a verb-idea and a noun-idea, it represents a whole proposition.

In general, then, phrases are more complete than words because they pack together more elements which have some logical role, establishing individuals, sets, spatiotemporal setting, and predicates which collectively establish sufficiently completed propositional attitudes. On this account the key role of phrase-structure is to establish phrases as signifying units on a logical level analogous to how lexemes are signifying units on a referential or conceptual level. Moreover, phrases' internal structure are understood to be governed by rule defining *how* word-combinations draw in extra logical detail. A link between words is not a random synthesis of concepts, but rather implies a certain logical connective which acts as a de facto "third party" in a double-word link, proscribing with orientation to predicate structures *how* the words' semantic concepts are to be joined. In *many students* the implied connector is the propositional act of conceiving a certain quantitative scale to a conceptualized set; in *students complained* the implied connector is a subject-plus-verb-equals-proposition assertiveness. Phrases acquire logical specificity by building up word-to-word connections into more complex aggregates.

One implication of this model is that phrases are semantically substitutable with individual lexemes that carry similar meanings, having been entrenched by convention to capture a multipart concept which would otherwise be conveyed with the aggregation of a phrase: consider "MP" for *member of parliament*, or "primaried" for *subject without your own party to a primary challenge*. Conversely, phrases can be repeatedly used in a specific context until they function as quasi lexical units in their own right. These patterns of entrenchment imply that we hear language in term of phrases bearing semantic content; and insofar as we are comfortable with how we parse a sentence, each word sited in its specific phrasal hierarchy, we do not tend to consider individual words semantically outside of their constituent phrases.

This theory of the syntax-to-semantic relationship is a paradigmatic partner, at the grammatic level, to the belief that meanings are fixed by "truth makers": that is a sentence asserts a true fact, this fact is its first and foremost meaning, and if not, the meaning is somehow the description of a possible world where it *is* true. One of my projects here is to criticize (to some extent) these "truth-theoretic" paradigms. At present I want to point out how this philosophical paradigm can influence our appraisal of phrase-structure. Simplistic "logic-based" conceptions of phrases can be based on two concerns: first, the idea that lexemes retain some syntactic and semantic autonomy even within clearly defined phrases where they are included; and, second, that the shape of phrases insofar as they are perceived as holistic signifying units is often driven by figurative or "gestalt" principles rather than neat logical structuration. I'll call the former the issue of "phrasal isolation" (or lack thereof): syntactic and semantic effect often cross phrasal boundaries, even outside the overarching hierarchy whose apex is the whole sentence. Both of these lines of reasoning — arguably especially the second — are developed in

© 2021/05/24 Nathaniel Christen

Cognitive Grammar literature.

In Ronald Langacker's *Foundations of Cognitive Grammar*, the sentence

- (37) Three times, students asked an interesting question.

is used to demonstrate how grammatical principles follow from cognitive "construals" of the relevant situations, those which language seeks to describe or takes as presupposed context.[3] In particular, Langacker argues that "students" and "question" can both be either singular or plural: syntax is open-ended here, with neither form being more evidently correct. Langacker uses this example to make the Cognitive-Linguistic point that we assess syntactic propriety relative to cognitive frames and conversational context. In this specific case, we are actually working with two different cognitive frames which are interlinked — on the one hand, we recognize distinct events consisting of a student asking a question, but the speaker calls attention, too, to their recurrence, so the events can also be understood as part of a single, larger pattern. There are therefore two different cognitive foci, at two different scales of time and attention, a "split focus" which makes both singular and plural invocations of "student" and "question" acceptable.

Supplementing this analysis, however, we can additionally focus attention directly on grammatical relations. The words *student* and *question* are clearly linked as the subject and object of the verb *asked*; yet, contrary to any simple presentation of rules, no agreement of singular or plural is required between them (they can be singular and/or plural in any combination). Moreover, this anomaly is only in force due to the context established by an initial phrase like *three times*; absent some such framing, the singular/plural relation would be more rigid. For example, "A student asked interesting questions" would (in isolation) strongly imply *one* student asking *several* questions. So the initial "Three times" phrase alters how the subsequent phrase-structure is understood, while remaining structurally isolated from the rest of the sentence. Semantically, it suggests a "space builder" in the manner of Gilles Fauconnier or Per Aage Brandt [24]; [15], but we need to supplement Mental Space analysis with a theory of how these spaces influence syntactic acceptability, which would seem to be logically prior to the stage where Mental Spaces would come in play.

The mapping of (37) to a logical substratum would be more transparent with a case like:

- (38) Three students asked interesting questions.

Here (38) is a more direct translation of the facts which the original sentence conveys. But this "more logical" example has different connotations than the sentence Langacker cites; (37) places the emphasis elsewhere, calling attention more to the idea of something temporally drawn-out, of a recurrence of events and a sense of timescale. The "more logical" sentence lacks this direct invocation of time scale and temporal progression.

We can say that the "Three students" version is a more direct statement of fact, whereas Langacker's version is more speaker-relative, in the sense that it elaborates more on the speaker's own acknowledgment of belief. The speaker retraces the steps of her coming to appreciate the fact — of coming to realize that the "interesting questions" were a recurrent phenomenon and therefore worthy of mention. By situating expressions relative to cognitive processes rather than to the facts themselves, the sentence takes on a structure which models the cognition rather than the states of affairs. But this shift of semantic grounding from the factual to the cognitive also apparently breaks down the logical orderliness of the phrase structure. "Three times", compared to "three students", leads to a morphosyntactic choice-space which is "underdetermined" and leaves room for speakers' shades of emphasis.

This is not an isolated example. Many sentences can be provided with similar phrase-structure complications, particularly with respect to singular/plural agreement.

- (39) Time after time, tourists (a tourist) walk(s) by this building with no idea of its history.
- (40) The streets around here are confusing; often people (someone) will ask me for directions.
- (41) Student after student came with their (his/her) paper to complain about my grade(s).
- (42) Student after student — and their (his/her) parents — complained about the tuition increase.

On a straightforward phrase-structure reading, *student after student* reduces to an elegant equivalent of *many students*, with the rhetorical flourish abstracted away to

[3] For example, [40, pp. 119 and 128], discussed by [14, p. 189], and [53, p. 9].

© 2021/05/24 Nathaniel Christen

a logical form. But our willingness to accept both singular and plural agreements (his/her/their parents, grades, papers) shows that clearly we don't simply substitute *many students*; we recognize the plural as a logical gloss on the situation but engage the sentence in a more cognitively complex way, recognizing connotations of temporal unfolding and juxtapositions of cognitive frames. The singular/plural underdeterminism is actually a signification in its own right, a signal to the listener that the sentence in question demands a layered cognitive attitude. Here again, syntactic structure (morphosyntactic, in that syntactic allowances are linked with variations in the morphology of individual words, such as singular or plural form) serves to corroborate conversants' cognitive frames rather than to model logical form.

The contrast between the phrases *Student after student* and *Many students* cannot be based on "abstract" semantics alone — how the evident temporal implications of the first form, for example, are concretely understood, depends on conversants' mutual recognition of a relevant time frame. The dialog may concern a single day, a school year, many years. We assume that the speakers share a similar choice of time "scale" (or can converge on one through subsequent conversation). *Some* time-frame is therefore presupposed in the discursive context, and the first phrase invokes this presumed but unstated framing. The semantics of the phrase are therefore somewhat open-ended: the phrase "hooks into" shared understanding of a temporal cognitive framing without referring to it directly. By contrast, the second phrase is less open-ended — it is consistent with both a more and less temporally protracted understanding of *many*, but leaves such details (whatever they may be) unsignified. The factual circumstance is designated with a level of abstraction that sets temporal considerations outside the focus of concern. The second ("*Many students*") phrase is therefore both less open-ended and also less expressive: it carries less detail but accordingly also relies less on speaker's contextual understanding to fill in detail.[4]

One consequence of these analyses is that grammar needs to be approached holistically: the grammatic structure of phrases cannot, except when deliberate oversimplification is warranted, be isolated from surrounding sentences and still larger discourse units. Phrases' semantic roles have some effect on their syntax, but phrases are nonetheless chosen from sets of options, whose variations reflect subtle semantic and syntactic maneuvers manifest at super-phrasal scales. The constituent words of phrases retain some autonomy, and can enter into inter-word and phrasal structures with other words outside their immediate phrase-context. We can still apply formal models to phrase structure — for example, Applicative and Cognitive Grammar (ACG) considers phrases as "applications" of (something like) linguistic or cognitive functions, with (say) an adjective modeled as a function applied to a noun, to yield a different noun (viz., something playing a noun's conceptual role) [22] — but we should not read these transformations too hastily as a purely semantic correlation within a space of denotable concepts *such that* the new concept wholly replaces the contained parts, which then cease to have further linguistic role and effect. Instead, applicative structures represent shifts or evolutions in mental construal, which proceed in stages as conversants form cognitive models of each others' discourse. Even if phrase structure sets landmarks in this unfolding, phrases do not wholly subsume their constituents; the parts within phrases do not "vanish" on the higher scale, but remain latent and may be "hooked" by other, overlapping phrases.

Consider the effect of "Many students complained". Propositionally, this appears to say essentially that *students* complained; but, on hermeneutic charity, the speaker had *some* reason to say "many". The familiar analysis is that "many" suggests relative size; but this is only half the story. If the speaker chose merely *students complained*, we would hear an assertion that more than one student did, but we would also understand that there were several occasions when complaints happened. Adding "many" does not just imply "more" students, but

[4] The examples I have used so far may also imply that a choice of phrase structure is always driven by semantic connotations of one structure or another; but seemingly the reverse can happen as well — speakers choose a semantic variant because its grammatic realization lends a useful organization to the larger expression. There are many ways to say "many", for example: *a lot of*, *quite a few*, not to mention "time after time" style constructions. Whatever their subtle semantic variations, these phrases also have different syntactic properties: *Quite a few* is legitimate as standalone (like an answer to a question); *A lot of* is not, and *A lot* on its own is awkward. On the other hand the "of" in *A lot of* can "float" to be replicated further on: "A lot of students, of citizens, believe education must be our top priority" sounds more decorous than the equivalent sentence with the second "of" replaced by "and". If the cadence of that sentence appeals to the speaker, then such stylistic preference will influence taking "A lot of" as the "many" variant of choice. So speakers have leeway in choosing grammatic forms that highlight one or another aspect of situations; but they also have leeway in choosing rhetorical and stylistic pitch. Both cognitive framings and stylistic performance can be factored when reconstructing what compels the choice of one sentence over alternatives.

© 2021/05/24 Nathaniel Christen

suggests a mental shift away from the particular episodes. In the other direction, saying *a student complained* is not just asserting how at least one student did so, but apparently reports one specific occasion (which perhaps the speaker wishes to elaborate on). In other words, we cannot really capture the singular/plural semantics, or different varieties of plural, just by looking at the relative size of implied sets; we need to track how representations of singleness or multitude imply temporal and event-situational details.

Against this backdrop, *Student after student complained* captures both dimensions, implying both a widespread unrest among the student body and also temporal recurrence of complainings. The ——`after`—— idiom is a special case vis-à-vis "after", where the "argument" to *after* is repeated in both positions, suggesting an unusual degree of repetition, something frustratingly recurrent: *He went on and on*; *Car after car passed us by*; *Time after time I got turned down.* Although I have no problem treating these constructions as idiomatic plurals, I also contend (on the premise of phrase-overlap) that the dependent constituents in the ——`after`—— construction can be hooked to other phrases as well (which is why "and [their/his/her] parents" can also be singular, in this case). I dwell on this example because it shows how type-theoretic accounts of phrase structure, which I will explore next section, can be useful even if we treat phrases more as frames which overlay linguistic structure, not as rigid compositional isolates. Each "students" variation uses morphology to nudge cognitive attention in one direction or another, toward events or the degree to which events are representative of some global property (here of a student body), or both. The pluralizing transformation — from one student to many students — is not *the* morphosyntactic meaning, but instead the skeleton on which the full meaning (via cognitive schema) is designed.

If this analysis has merit, it suggests that a Combinatory Grammar or type-logical approach to phrases like *many students* or *student after student* (singular-to-plural or plural-to-plural mappings) should be understood not just as functions among Part of Speech (POS) types but as adding cognitive shading, foregrounding or backgrounding cognitive elements like events or typicality in some context. In other words, *many students* is type-theoretically a pluralizing action, but, in more detail, it adds a kind of cognitive rider attached to the mapping which focuses cognition in the subsequent discourse onto events (their recurrence and temporal distribution); similarly "student after student" has a "rider" suggesting more of a temporal unfolding. The second form implies not only that many students complained, but that the events of these complainings were spread out over some stretch of time. Each such functional application (mappings between POS understood as linguistic types) produces not only a resulting POS "type", but also a reconfiguration of cognitive attitudes toward the relevant situation and context. Language users have many ways to craft a sentence with similar meanings, and arguably one task for linguistic analysis is to model the space of choices which are available in a given situation and represent what specific ideas and effects are invoked by one choice over others.

Logical semantic models build up propositional representations from the word level to phrase and sentence levels. Even when we find that a logical *telos* is in effect, however, at each of these levels we can find interpretive and situational details which may complicate, flesh out, or cognitively supersede a more transparent logical meaning. So at the word-pair level, there is more going on in a transform like *rescued dog* than simply applying the predicate "rescue" to the ground "dog". The construction drafts a complex backstory that influences our cognizing both component words. Scaling up, as I argued in the last several paragraphs, a phrase like *student after student* carries conceptual effects beyond just effectuating a logical operation of pluralization. I will argue that similar effects can be identified at the sentence level — so there are certain lacunae in narrowly logically-based reconstructions of language components at the distinct levels of word (transform) pairs, phrases, and whole sentences. I'll sketch my arguments here and give a more thorough analysis later in Section 3.

1.3 Logical Structure versus Sentence Structure

Let us grant in general that particular sentences can be mapped to distinct, relatively transparent propositional contents. In some cases sentences express propositional attitudes to such content (requests, commands, questioning) rather than unadorned locutionary assertions.

© 2021/05/24 Nathaniel Christen

To properly respond to speech-acts, however (even ones with illocutionary force) conversants need to derive the content which is logically conjured via the discourse, either as the speaker's primary intent or as a condition for that intent. In effect, a proposition like *the window is closed* furnishes logical content to assertions like *The window is closed now* but also statements of belief (*I think the window is closed*) or requests or opinions (*The window should be closed*; *Please close the window*).

Philosophical "truth-theoretic" paradigms imply that such propositional contents are the *essential* meanings within language; that analyzing semantic forms via logical structure is the core of a rigorous theory of semantics. It is certainly true that many elements of language can be translated, or deemed as conventionalized encodings for, structures in predicate logic — invocations of multiplicity and quantification; logical connectives between propositions; negations, modalities, and possibilia. This provides an analytic matrix wherein *some* sentences' structures may be analyzed. I will argue, however, that in typical cases logical forms are invoked only indirectly — which calls into question the applicability of logical analyses as explanatory vehicles for *linguistic* analysis in itself (as opposed to more general cognitive/extralinguistic processing).

There are several cognitive operations requisite for grasping sentence-meanings as a logical gestalt: figuring individuals or multiplicities as conceptual foci (verb subjects or objects); establishing relationships between individuals and multiplicities or among multiplicities (member/part of, larger/smaller, overlap/disjoint); predicating properties to individuals or multiplicities; quantification; logical conjunction or disjunction, between predicates (also negation). In some cases we can find these operations fairly directly encoded in explicit language form — sentences which are precise in figuring multiplicities numerically, or through unambiguous use of determiners like *all* and *every*; which are structured to avoid scope ambiguities; which use transparent semantic resources to describe verb subjects and objects; and so forth. In the most recent Universal Dependencies Shared Task corpus we can find examples like:

- (43) It is the most common tumour found in babies, occurring in one of every 35,000 births. `(en_pud)`
- (44) Dengue fever is a leading cause of illness and death in the tropics and subtropics, with as many as 100 million people infected each year. `(en_partut-ud-train)`
- (45) Many Taliban living in Afghanistan voted for President Karzai. `(en_ewt-ud-train)`
- (46) Most of the girls I was meeting had grown up in Mujahedeen schools in Ashraf, where they lived separated from their parents. `(en_ewt-ud-train)`
- (47) Most experts believe China intends to develop a small space station of its own over the next several years. `(en_ewt-ud-train)`
- (48) Check out their wine tastings every Friday night! `(en_ewt-ud-train)`
- (49) For each start tag, there is a corresponding end tag. `(en_lines-ud-train)`
- (50) Each collection donated by the Andy Warhol Photographic Legacy Program holds Polaroids of well-known celebrities. `(en_gum-ud-train)`

These sentences have straightforward logical structure, in terms of how they establish topical foci (*one of every 35,000 births*; *as many as 100 million people*; *Many Taliban*; *every Friday night*; *For each start tag*), and how predicates or references are bound together to create more precise significations (*the tropics and subtropics*; *in Mujahedeen schools in Ashraf*; *a corresponding end tag*). Properties ascribed to subject foci are neatly drawn, both in conveying the property intended and its bearer, according to the sentence's terms: *the most common tumour found in babies*; *a leading cause of illness and death*; *China intends to develop a small space station*; *holds Polaroids of well-known celebrities.* With aggregate foci and/or quantification, there is an unambiguous framing of predication and quantifier scope — Each collection has its set of Polaroids; the set of Karzai voters, Dengue infections, birth tumours, etc., are crisply figured.

For many philosophers of language, identifying similar logical structuration is an intrinsic aspect of coming to terms with human language in general. This paradigm also reinforces the goal of Artificial Intelligent Natural Language Processing, because computers can certainly engage in the kind of symbolic-logical reasoning outlining signified meanings in cases where language reciprocates propositional morphology very clearly. The problem is that language artifacts often cloak their logical core, such that examples like (43)-(50) are not representative of language as a whole. Logical patterns may certainly be present, but they are not necessarily structurally reproduced in surface-level formations; rather a sentences'

© 2021/05/24 Nathaniel Christen

propositional content can depend on a subtle interpretive trajectory. I will present examples throughout this paper, but a few further corpus items are reasonable case-studies:

- (51) A furry black band of ants led up a cupboard door to some scrap that had flicked from a plate. `(en_lines-ud-train)`
- (52) The current waiting period is eight weeks. `(en_pud)`
- (53) I think that's why they immersed themselves in pattern and colour. `(en_pud)`
- (54) With her appearance finalized, Jasmine became Disney's first non-white princess as opposed to being of European heritage. `(en_pud)`

It requires a certain cognitive flexibility to understand a band of ants as "flurry", or to parse the disjoint time-frames in *current waiting period*. In (53), the presumed sense of "immerse" transcends any immediate, perceptual immersion, instead involving scholarship or engagement with artistic form; and (54) depends on us understanding the meaning of temporality in Jasmine's appearance being *finalized*, and also her *becoming* non-white. As a fictional character, discourse about Jasmine can be evaluated in the time-frame of her artistic creation, distinct from the fictional time of her narrated world.

I think the intended propositional content in (51)-(54) is no less evident than in (43)-(50); however, interpreting the topical foci and predicate attributions constituting such propositional content requires a holistic reading whose compositional structure is not recapitulated in the sentence-forms themselves. In the latter examples, then, merely notating propositional content in logical fashion does not yield a very informative *linguistic* analysis, since it does not address the key question of *how* the sentences signify those propositions.

I propose to use the term "logicomorphic" for sentence in the former vein; in such cases, pointing out propositional content is linguistically useful because we can treat that content as a prototype for sentence organization. That is, propositional content is not only *holistically* signified but, in its logical structure, sheds light on patterns in the language. The purpose of phraseology like *most common tumour*, *Many Taliban living in Afghanistan*, *Most of the girls I was meeting*, etc., is to circumscribe a focus or a property suitable for predication, and we can logically model the tools used to do so: logical superlative (*most common*), assertions of magnitude (*Many*, *Most of*), refining an multiplicity with some further criteria (*the girls I was meeting*, *Taliban living in Afghanistan*), and so forth. These are "logicomorphic" constructions in that we can read the logical structure of signified propositional content as a direct cause of the given phrasal morphology.

On the other hand, I call examples like (51)-(54) "interpretive" because the sentences' propositional content, with its logical structure, does not explain the compositional rationale for the explicit linguistic form — we cannot read any pattern in the logic as a direct motivation for how the sentence is pieced together. The spectrum between *logicomorphic* and *interpretive* represents different strategies by which language is composed in anticipation of its cognitive reception, with the eventual goal of establishing a signified propositional content, but in different ways. On the logicomorphic side, logical form informs language directly; on the interpretive side, the actual rationale for compositional structures transcends exact predicative structure — a more perceptual or indirect figuring of topical focus, for instance, or a more elliptical construal of predicate attributes, leaving the hearer to piece together the final propositional content via some pragmatic or extralinguistic calculation.

I emphasize that an "interpretive" mode of constructing sentences (and sentence-like units) is not at all equivalent to logical ambiguity, imprecision, underdetermination, or any other effect of a sentence *failing* (perhaps by intent) to convey a logically straightforward meaning. The interpretive "register" is perhaps more conducive to the emergence of logical ambiguity in one sense or another, but may also yield productions which are very clear in propositional terms. Moreover, when ambiguity arises, such opacity cannot necessarily be explained by the employment of a more interpretive style to begin with. Instead, the point of contrasting *interpretive* with *logicomorphic* is to highlight how propositionally clear content does not necessarily get conveyed in language via discursive assemblies that *structurally resemble*, or represent or recapitulate, the propositional complexes they seek to invoke. The gap between *interpretive* and *logicomorphic* correlates with how far logical structure lies from the "surface" of language.

In sum, the *logicomorphic*/*interpretive* distinction — along with the overall contrast between linguistic and ex-

© 2021/05/24 Nathaniel Christen

tralinguistic aspects of meaning — are contrasts between sentences that become manifest in the compositional maxims evident at subsentence (phrasal and inter-word) scales. We can apply all four criteria to estimate the cognitive as well as syntactic and semantic paradigms in effect for given inter-word pairs and phrasal structure; identifying sentences as logicomorphic or interpretive propagates down to how phrasal and interword patterns should be analyzed. With this in mind, having presented certain claims as to the holistic nature of sentences vis-à-vis propositional content, I will now switch attention to the composition of sentences from the interword level upward.

2 Functional Type Theory in Cognitive and Dependency Grammars

My discussion so far has focused on characterizing sentences' holistic meaning. On the face of it, such holistic analysis is more semantic than syntactic. However, syntactic paradigms can be grounded in theories of how language elements aggregate *toward* holistic meaning.

Here I propose a notion of "cognitive transforms" — that holistic meanings emerge from a series of interpretive and situational modeling modifications which progressively refine our understanding of a speaker's construal of our environing context and her propositional attitudes. While elucidation of these transforms as cognitive phenomena may belong to semantics, syntactic structure dictates the *sequence* of transforms. Many transforms are expressed by individual word-pairs. Taking the temporal or logical order of transforms into consideration, we can derive a syntactic model of sentences by introducing an order among word-pairs — a methodology akin to using Dependency Grammar parse-graphs as an intermediate stage, then ordering the graph-edges around an estimation of cognitive aggregation. One transform is a successor to a predecessor if the modifications induced by the predecessor are consequential for the cognitive reorientation pertinent to the successor, and/or to the morphosyntactic features which trigger it.

In this spirit I talk of Cognitive Transform *Grammar*, because while in the general case transforms are semantic and interpretive — not the purview of grammar per se — we can theorize grammar as governing the *order of precedence* among transforms. More precisely, there is a particular order of precedence germane to sentence meaning; sentences have their precise syntax in order to compel recipients' reception of the linguistic performance according to that same ordering.

From this perspective, an essential aspect of grammar theory is that whatever units are understood as syntactic constituents — like phrase structure or word-pairs — an order of precedence should "fall out" of grammatic reconstructions. We should be able to supplement parse-representations with a listing of salient syntactic features in order, retracing the *cognitive* steps by which localized sense-alterations synthesize into holistic meaning. The details of this precedence-establishment will vary across grammatic paradigms, so one way to assess grammar theories is to consider how they engender corresponding cognitive-transform models.

Type theory can be adopted in this context because most versions of type theory include a notion of "function-like" types: types whose instances modify instances of some other type (or types). This establishes an order of precedence: anything modified by a function is in some sense logically prior to that function. In formal (e.g., programming) languages, the procedure whose output becomes input to a different procedure must be evaluated before the latter procedure begins (or at least before the output-to-input value is used by that latter procedure).

A common paradigm is to consider natural-language types as generated by just two bases — a noun type N and a proposition type $\mathcal{P}rop$, the type of sentences and of sentence-parts which are complete ideas — having in themselves a propositional content (see e.g. [6] or [38]). Different models derive new types on this basis in different ways. One approach, inspired by mathematical "pregroups", establishes derivative types in terms of word pairs — an adjective followed by a noun yields another noun (a noun-phrase, but N is the phrase's *type*) — e.g., *rescued dogs*, like *dogs*, is conceptually a noun. Adjectives themselves then have the type of words which form nouns when paired with a following noun, often written as N/N. Pregroup grammars distinguish left-hand and right-hand adjacency — *bark loudly*, say, demonstrates an adverb *after* a verb, yielding a verb phrase: so "loudly" here has the type of a word producing a verb in combination with a verb to its *left* (sometimes written $V\backslash N$); by contrast adjectives combine with nouns to their *right*.

© 2021/05/24 Nathaniel Christen

A related formalization, whose formal analogs lie in Typed Lambda Calculus, abstracts from left-or-right word order to models derived types as equivalent (at least for purposes of type attribution) to "functions". This engenders a fundamental notion of functional *application* and operator/operand distinctions:

> *Categorial Grammars* make the connection between the first and the second level of the ACG. These models are typed applicative systems that verify if a linguistic expression is syntactically well-formed, and then construct its functional semantic interpretation. They use the *fundamental operation of application* of an operator to an operand. The result is a new operand or a new operator. This operation can be repeated as many times as necessary to construct the linguistic expressions' functional semantic representation. [64, p. 1]

So, e.g., an adverb becomes a function which takes a verb and produces another verb; an adjective takes a noun and produces another noun; and a verb takes a noun and produces a proposition. By "function" we can consider some kind of conceptual transformation: *loudly* transforms the concept *bark* into the concept *loud bark*. Assuming all lexemes are assigned a Part of Speech drawn from such a type system, the definition of functional types directly yields a precedence order: instances of functional types are functionally dependent on their inputs, which are therefore precedent to them. On this basis, any well-typed functional expression has a unique precedence ordering on its terminal elements (i.e., its "leaves" when the expression is viewed as a tree, or its nodes when viewed as a graph), which can be uncovered via a straightforward algorithm (one implementation is part of this paper's data set; see the "parse_sxp" method in file **ntxh-udp-sentence.cpp**).

Functional parts of speech that can be formally modeled with one "argument" (the most common case), having a single input and output type, conveniently lend themselves to cognitive transforms defined through word pairs — an adjective modifies a noun to another noun, an adverb maps a verb to a verb, an auxiliary like *that* or *having* can map verbs or propositions to nouns, and so forth. The only main complication to this picture is that verbs, which typically have subjects as well as objects, can take two or three "inputs" instead of just one. Instead of a transform *pair* we can then consider a three- or four-part transform *structure* (e.g., verb, subject, direct object, indirect object). We can still assign a precedence ordering to verb-headed phrases, however, perhaps by stipulating that the subject takes precedence before the direct object, and the direct object before the indirect. This ordering seems cognitively motivated: our construal of the significance of a direct object appears to intellectually depend on the verb's subject; likewise the indirect object depends on the direct object to the degree that it is rationally consequential.

A secondary complication involves copulae like *and*, which can connect more than two words or clauses. Here, though, a natural ordering seems to derive from linear position in the sentence: given *x, y, and z* we can treat x as precedent to y, and y to z, respectively.

In total, sentences as a whole can thus be seen as structurally akin to nested expressions in lambda calculi (and notated via "S-Expressions", like code in the Lisp programming language). S-Expressions are occasionally recommended as representations for some level of linguistic analysis (cf. [61], [69], [17]), and if this form by itself may add little extra data, it does offer a succinct way to capture the functional sequencing attributed to a sentence during analysis. Given, say,

▼ (55) The city's ambience is colonial and the climate is tropical.
`(en_gum-ud-train)`

the gloss `(and (is ((The ('s city)) ambience) colonial) (is (the climate) tropical))` summarizes analytic commitments with regard to the root structure of the sentence (in my treatment the copula is the overall root word) and to precedence between words (which words are seen as modifiers and which are their ground, for instance). So even without extra annotations (without, say, the kind of tagging data included by treebanks using S-Expression serializations), rewriting sentences as nested expressions captures primitive but significant syntactic details.

Nested-expression models also give rise directly to two other representations: a precedence ordering among lexemes automatically follows by taking function inputs as precedent to function (words) themselves[5]; moreover,

[5]But note that using "function-words" as terminology here generalizes this term beyond its conventional meaning in grammar.

© 2021/05/24 Nathaniel Christen

S-Expression formats can be rewritten as sets of word-pairs, borrowing the representational paradigms (if not identical structures) of Dependency Graphs. This allows Dependency Graphs and S-Expressions to be juxtaposed, which I will discuss in the remainder of this section.

2.1 Double de Bruijn Indices

Assume then that all non-trivial sentences are nested expressions, and that all lexemes other than nouns are notionally *functions*, which take typed "inputs" and produce typed "outcomes". Expression "nesting" means that function inputs are often outcomes from other functions (which establishes a precedence order among them). Since there is an obvious notion of "parent" — instances of functional types are parents of the words or phrase-heads which are their inputs — nestable expressions are formally trees. Via tree-to-graph embedding, they can also be treated as graphs, with edges linking parents to children; since parse-graphs are canonical in some grammar theories (like Link and Dependency grammar), it is useful to consider the graph-style representation as the intrinsic structure of linguistic glosses based on S-Expressions. That is, we want to define a Category of labeled graphs each of whose objects is isomorphic to an S-Expression (using this terminology in the sense of mathematical Category Theory); equivalently, a bijective encoding of S-Expressions within labeled graphs given a suitable class of edge-labels.

Indeed, labels comprised of two numbers suffice, generalizing the lambda-calculus "de Bruijn Indices". The de Bruijn notation is an alternative presentation of lambda terms using numeric indices in lieu of lambda-abstracted symbols (cf., for a linguistic or discourse-representation context, [75] and [76]). The *double* indices accommodate the fact that, in the general case, the functional component of an expression may be itself a nested expression, meaning that "evaluation" has to proceed in several stages: a function (potentially with one or more inputs) is evaluated, yielding another function, which is then applied to inputs, perhaps again yielding a function applied to still more inputs, and so forth. I use the term "evaluate" which is proper to the computer-science context, but in linguistics we can take this as (only) a suggestive metaphor. More correctly, we can say that a function/input structure represents a cognitive transform which produces a new function (i.e., a phrase with a function-like part of speech), that is then the modifier to a new transform, and so forth. In general, the result of a transform can either be the *ground* of a subsequent transform, which is akin to passing a function-result to another function; or it can be the *modifier* of a subsequent transform, which is akin to evaluating a nested expression to produce a new function, then applied to other inputs in turn.

For a concrete example, consider

▾ (56) The most popular lodging is actually camping on the beaches. `(en_gum-ud-train)`

with gloss, as I take it, `((actually is) ((The (most popular)) lodging) ((on (the beaches)) camping))`. Here the adverb *actually* is taken as a modifier to *is*, so we imagine that interpreting the sentence involves first refining *is* into *is actually*, yielding a new verb (or "verb-idea") that then participates in a verb-subject-object pattern. Hence, the parse opens with the evaluation `(actually is)` in the head-position of the sentence's top-level expression. Similarly, I read *on the beaches* as functionally a kind of adverb, like "outside" in *camping outside*. In the generic pattern, a verb can be paired with a designation of location to construct the idea of the verb happening at such location; the designation-of-location is then a modifier to the verb's ground. When this designation is a locative construction, the whole expression becomes a modifier, while it also has its own internal structure. In *on the beaches*, *on* serves as a modifier which maps or reinterprets *the beaches* to a designation of place.

So here is the unfolding of the phrase: in *the beaches* the determinant (*the*) is a modifier to the ground *beaches*, signifying that *beaches* are to be circumscribed as an aggregate focus. Then *on* modifies the outcome of that first transform, re-inscribing the focus as a place-designation. Then *that* transform's output becomes a modifier for *camping*, wherein the locative construction becomes a defacto adverb, adding detail to the verb *camping* (camping on the beach as a kind of camping, in effect).[6]

[6] If it seems better to read camping as a *noun* — the act or phenomenon of camping — then we could treat the locative as an *adjective*, with the stipulation that the operation converting verbs to nouns (from X to the phenomenon, act, or state of X-ing) propagates to any modifiers on the verb: modifying constructions that refine X as verb are implicitly mapped to be adjectives likewise modifying X in the nominal sense of

© 2021/05/24 Nathaniel Christen

Notice in this review that *the* as modifier in *the beaches* yield a pair whose outcome is the *ground* for *on*. If we take the modifier as representative for a modifier-ground pair, *the* is the *modifier* in its own transform pair but then the *ground* in the subsequent pair; the pattern is **modifier**-then-**ground**. However, *on* is the modifier vis-à-vis *the* and then *also* modifier vis-à-vis *camping*; the pattern is **modifier**-then-**modifier**. This latter case is the scenario where a lexeme will be a modifier on two or more different "levels", giving rise to the "doubling" of de Bruijn indices. The first index, that is, represents the "level" tieing a modifier to a ground, while the second index is the *normal* notation of lambda-position. In *camping on the beaches*, the indices for the pair *on*/*the* would be `1,1` (meaning *the* is the first argument to *on* on the first transform level); the indices for *on*/*camping* would be `2,1` (*camping* is the first argument to *on* on the *second* transform level).

By combining an index for "transform levels" — capturing cases where a modifier produces an outcome which is a modifier again, not a ground — with an index for lambda position (e.g. the direct object has index 2 relative to the verb, and the indirect object has index 3), we can transform any expression-tree into a labeled graph. Parse graphs can then be annotated with these double-indices via the same presentations employed for Link or Dependency Grammar labels. Sentence (56) could be visualized as in Figure 1, with the double-indices juxtaposed alongside conventional Dependency labels (the indices below the sentence, and relation labels above it); the upper parse is drawn from the Universal Dependency corpus (other annotated examples are included in this paper's downloadable data set).

As a natural corollary to this notation, parts of speech can have "type signatures" notionally similar to the signatures of function types in programming languages: a verb needing a direct object, for example, "transforms" two nouns (Subject and Object) to a proposition, which could be notated with something like $N \dot{\rightarrow} N \dot{\twoheadrightarrow} \mathcal{P}rop$.[7] The notation is consistent so long as each constituent of a verb phrase has a fixed index number.[8] Transforms (potentially with two or more arguments) then combine lexemes — having the right signatures — with one or more words or phrases. Type analysis recognizes criteria on these combinations, insofar as the phrases or lexemes at a given index have a type consistent with (potentially a subtype of) the corresponding position in the head-word's signature. Ideally, this representation can explain *ungrammatical* constructions via *failure* for these types to align properly. If the combination "type-checks", then we can assign the phrase as a whole the type indicated in the signature's output type — $\mathcal{P}rop$ in $N \dot{\rightarrow} N \dot{\twoheadrightarrow} \mathcal{P}rop$, say; in this case the signature points out how $\mathcal{P}rop$ derives from the phrase with two Ns as its components (along with the verb), a derivation we could in turn notate like $N \dot{\rightarrow} N \dot{\twoheadrightarrow} \mathcal{P}rop \gg \mathcal{P}rop$. A resulting phrase can then be included, like a nested expression, in other phrases; for instance a $\mathcal{P}rop$ joined to *that* so as to create a noun-phrase (recall my analysis of *question whether* last section); notationally $\mathcal{P}rop \dot{\twoheadrightarrow} N \gg N$.[9]

Type "signatures" like $N \dot{\rightarrow} N \dot{\twoheadrightarrow} \mathcal{P}rop$ may seem little more than notational variants of conventional linguistic wisdom, such as how sentences requiring a noun (-phrase) and a verb (`S = NP + VP`). Even at this level, however, type-theoretic intuitions offer techniques for making sense of more complex, layered sentences, where integrating Dependency Graphs and phrase structures can be complex. One complication is the problem of applying Dependency Grammar where phrases do not seem to have an obviously "most significant" word for linkage with other phrases.

Often phrases are refinements of one single crucial component — a phrase like *many students* becomes in some sense collapsible to its semantic core, *students*. In real-world examples, however, lexemes tend to be neither wholly subsumed by their surrounding phrase nor wholly autonomous:

▼ (57) Many students and their parents came to complain about the tuition hikes.

"the phenomenon of X-ing".

[7] A note on notation: I adopt the Haskell convention (referring to the Haskell programming language and other functional languages) of using arrows both between parameters and before output notation, but for visual cue I add one dot above the arrow in the former case, and two dots in the latter: $Arg_0 \dot{\rightarrow} Arg_1 \dot{\rightarrow} \cdots \dot{\twoheadrightarrow} Result$. I will use N and $\mathcal{P}rop$ for the broadest designation of nouns and propositions/sentences (the broadest noun type, respectively type of sentences, assuming we are using type-theoretic principles). I will use some extra markings (in diagrams below) for more specific versions of nouns.

[8] The subject at position one, for instance; direct object at position two; and indirect object at position three.

[9] Numeric indices take the place of left and right adjacency — "looking forward" or backward — in Combinatory Categorial Grammar; the type-theoretic perspective abstracts from word order. The theory of result types "falling out" from a type-checked phrase structure, however, carries over to this more abstract analysis.

© 2021/05/24 Nathaniel Christen

Figure 1: Parse Graph with Dependency Relations as well as Double (Functional-Type) Indices

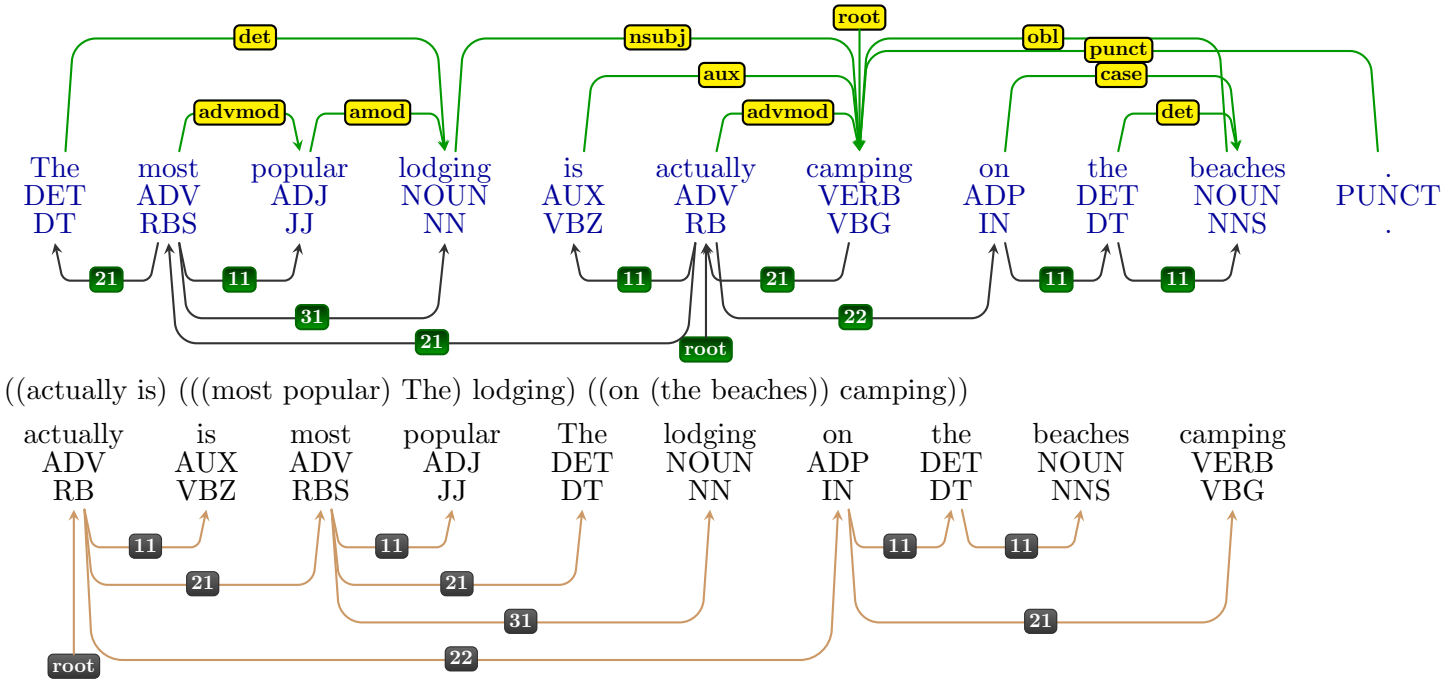

- (58) Many students came by my office to complain about their grades.
- (59) Student after student complained about the tuition hikes.
- (60) Student after student came with their parents to complain about the tuition hikes.

In (57) and (58), we read *Many students* as topicalizing a multitude, but we recognize that each student has their own parents, grade, and we assume they came to the office at different times (rather than all at once). So *students* links conceptually with other sentence elements, in a way that pulls it partly outside the *Many students* phrase; the phrase itself is a space-builder which leaves open the possibility of multiple derived spaces. This kind of space-building duality is reflected in how the singular/plural alternative is underdetermined in a multi-space context; consider Langacker's example:

- (61) (repeating 37) Three times, students asked an interesting question.
- (62) Three times, a student asked an interesting question.

Meanwhile, in (59) and (60) the phrase *Student after student* invokes a multiplicity akin to *Many students*, but the former phrase has distinct syntactic properties; in particular we can replace *their parents* (which is ambiguous between a plural and a gender-neutral singular reading) with, say, *his parents* (at a boy's school), a valid substitution in (59)-(60) but not (57)-(58).

Cases like *Student after student* (or consider *time after time*, *year after year*, and so forth; this is a common idiomatic pattern in English) present a further difficulty for Dependency Grammar, since it is hard to identify which word of the three is the more significant, or the "head". Arguably Constituency Grammar is more intuitive here because then phrases as a whole can get linked to other phrases, without needing to nominate one word to proxy the enclosing phrase. As I feel the "students" examples illustrated, however, it is too simplistic to treat phrases as full-scale replacements for semantic units, as if any phrase is an ad-hoc single lexeme (of course some phrases *do* get entrenched as de facto lexemes, like *Member of Parliament*, or my earlier examples *red card* and *stolen base*). In the general case though component words retain some syntactic and semantic autonomy (entrenchment diminishes but does not entirely eliminate such autonomy). There is, then, a potential dilemma: phrases link to other phrases (sometimes via subsumption and sometimes more indirectly, as in anaphora resolution), but phrases are not undifferentiated units; lexemes, which on

© 2021/05/24 Nathaniel Christen

this sort of analysis *are* units, can be designated as proxies for their phrase; but then we can have controversy over which word in a phrase is the most useful stand-in for the whole ([51] has an interesting review of a similar controversy in Dependency Grammar).

Incorporating type theory, we can skirt these issues by modeling phrases through the perspective of type signatures: given Part of Speech annotations for phrasal units and then for some of their parts, the signatures of other parts, like verbs or adjectives linked to nouns, or adverbs linked to verbs, tend to follow automatically. A successful analysis yields a formal tree, where if (in an act of semantic abstraction) words are replaced by their types, the "root" type is something like $\mathcal{P}rop$ and the rest of a tree is formally a reducible structure in Typed Lambda Calculus: $N \dashrightarrow N \dashrightarrow \mathcal{P}rop$ "collapses" to $\mathcal{P}rop$, $\mathcal{P}rop \dashrightarrow N$ collapses to N, and so forth, with the tree "folding inward" like a fan until only the root remains — though a more subtle analysis would replace the single $\mathcal{P}rop$ type with variants that recognize different forms of speech acts, like questions and commands. In Figure 2, this can be seen via the type annotations: from right to left $N \dashrightarrow N$ yields the N as second argument for *is*, which in turn yields a $\mathcal{P}rop$ that is mapped (by *that*) to N, finally becoming the second argument to *know*. This calculation only considers the most coarse-grained classification (noun, verb, proposition) — as I have emphasized, a purely formal reduction can introduce finer-grained grammatical or lexico-semantic classes (like *at* needing an "argument" which is somehow an expression of place — or time, as in *at noon*). Just as useful, however, may be analyses which leave the formal type scaffolding at a very basic level and introduce finer type or type-instance qualifications at a separate stage.

In either case, Parts of Speech are modeled as (somehow analogous to) functions, but the important analogy is that they have *type signatures* which formally resemble functions'. Words with function-like types proxy their corresponding phrase, not because they are necessarily more important or are Dependency "heads", but because they supply the pivot in the type resolutions which, collectively/sequentially, progress to a propositional culmination. This epistemological telos induces a sequencing on the type resolutions — there is a fixed way that trees collapse — which motivates the selection of function-words to proxy phrases; they are not semantically more consequential, necessarily, but are landmarks in a dynamic figured syntactically as "folding inward" and semantically as a progressive signifying refinement. Phrases are modeled via a "function-like" Parts of Speech along with one or more additional words whose own types match its signature; the type calculations "collapsing" these phrases can mimic semantic simplifications like *many students* to *students*, but here the theory is explicit that the simplification is grammatic and not semantic: the collapse is acknowledged at the level of *types*, not *meanings*. In addition, tree structures can be modeled purely in terms of inter-word relations — as I have proposed here with double-indices — so a type-summary of a sentence's phrase structure can be notated and analyzed without leaving the Link or Dependency Grammar paradigm.

In sum, then, function-like words can always be represented as the "head" of corresponding phrases, but this implies neither greater semantic importance nor that phrases are conceptual units that fully subsume their parts. Instead, a "head-function" notation captures the idea that sentence-synthesis is bounded by considerations of conceptual integrity that we can model, to some coarse approximation, via type theory. Designating *after* as the "head" in *student after student* means that we have a type machinery that models (coarsely but formally) how successive cognitive refinement converges onto something with the conceptual profile of a proposition, and that we can leverage this formality by identifying certain words as functional pivots in the synthesis-toward-proposition: these words are not necessarily pivotal in term of meaning, but they are the core skeleton of the schematic type-based model which we attach to a sentence while modeling the constraints on its synthesizing process (there is a further dimension in the *student after student* phrase — the repetition of "student" — which I will discuss later).

Or at least, this is one area of analysis where type theory is relevant for linguistics. I will argue that there are several different methodologies where linguists have turned to type theory, and moreover that they can be integrated into a unified picture organized around the granularity of types themselves, according to the relevant theories.

© 2021/05/24 Nathaniel Christen

Figure 2: Dependency-style graph with type annotations

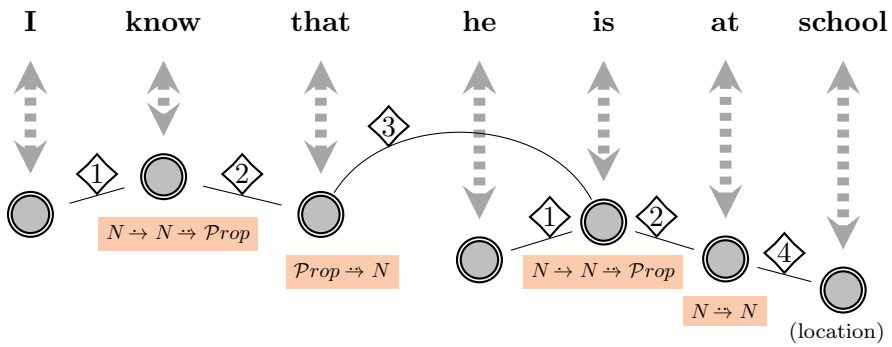

- 1 Verb's subject argument
- 2 Verb's direct object argument
- 3 Propositional "packaging" ("typed" as $\mathcal{P}rop \dashrightarrow N$)
- 4 Locative auxiliary link
  (may be typed as converting nouns to place-designations)

© 2021/05/24 Nathaniel Christen

2.2 *Three tiers of linguistic type theory*

When explaining grammaticality as type-checking — the concordance between function-like signatures and word or phrase "arguments" — types are essentially structural artifacts; their significance lies in the compositional patterns guiding phrases to merge into larger phrases in a well-ordered way — specifically, that the "outermost" expression, canonically the whole sentence, is type-theoretically a proposition. I proposed earlier that sentence-understanding be read as an accretion of detail culminating in a complete idea; type-checking then imposes regulatory guidelines on this accretion, with each constituent phrase being an intermediate stage. Assigning types to phrases presents a formal means of checking that the accretion stays on track to an epistemological telos — that the accumulated detail will eventually cohere into a propositional whole, a trajectory formally captured by the progressive folding-inward of phrase types to a propositional root.

Types themselves are therefore partly structural fiats — they are marks on intermediate processing stages embodying the paradigm that type-checking *captures* the orderliness of how successive cognitive transforms accrue detail toward a propositionally free-standing end-point. At the same time, types also have semantic interpretations; the N vs. $\mathcal{P}rop$ distinction, for example, is motivated by the cognitive difference between nominals and states of affairs as units of reason. Type-theoretic semantics allows the structural paradigm of type-checked resolutions, the tree "folding inward" onto its root, to be merged with a more semantic or conceptual analysis of types qua categories or classifications of meanings (or of units comprising meanings). I have described this merger at a coarse level of classification, taking broad parts of speech as individual types, but similar methods apply to more fine-grained analysis. By three "tiers" of linguistic organization, I am thinking of different levels of granularity, distinguished by relative scales of resolution, amongst the semantic implications of putative type representations for linguistic phenomena.

From one perspective, grammar is just a most top-level semantics, the primordial Ontological division of language into designations of things or substances (nouns), events or processes (verbs), qualities and attributes (adjectives), and so forth. Further distinctions like count, mass, and plural nouns add semantic precision but arguably remain in the orbit of grammar (singular/plural agreement rules, for example); the question is whether semantic detail gets increasingly fine-grained and somewhere therein lies a "boundary" between syntax and semantics. The mass/count distinction is perhaps a topic in grammar more so than semantics, because its primary manifestation in language is via agreement (*some* wine

in a glass; *a* wine that won a prize; *many* wines from Bordeaux). But what about distinctions between natural and constructed objects, or animate and inanimate kinds, or social institutions and natural systems, matters more of grammar or of lexicon?

For example, the template *I believed X* generally requires that *X* be a noun (?*I believed run*), but more narrowly a certain *type* of noun, something that can be interpreted as an idea or proposition of some kind (?*I believed flower*). Nicholas Asher and James Pustejovsky point out the anomaly in a sentence like "Bob's idea weighs five pounds" [5, example 2, p. 5], which possesses a flavor of unacceptability that feels akin to Part of Speech errors but are not in fact syntactic errors. The object of *weigh* is "five pounds" and its subject is "Bob's idea", which is admissible *syntactically* but fails to honor our semantic convention that the verb "to weigh" should be applied to things with physical mass (at least if the direct object denotes a quantity; contrast with *Let's all weigh Bob's idea*, where the *idea* is object rather than subject). These conventions are analogous to Part of Speech rules but more fine-grained: there is a meaning of *weigh* which has (like any transitive verb) to be paired with a subject and object noun, but beyond just being nouns the subject must be a physical body (in effect a sub-type of nouns) and the object a quantitative expression (another sub-type of nouns). Potentially, type restrictions on a coarse scale (e.g. that the subject of a verb must be a noun) and those on a finer scale (as in this sense of *to weigh*) can be unified into an overarching theory, which spans both grammar and semantics — for instance, both Part of Speech rules and usage conventions of the kind often subtly or cleverly subverted in metaphor and idioms (see *flowers want sunshine*, *my computer died*, *neutrinos are sneaky*, as rather elegantly compactified by assigning sentient states to inert things). This is one way of reading the type-theoretic semantic project.

Certainly "Ontological" qualities of signifieds engender agreements and propriety which appear similar to grammatic rules. *The tree wants to run away from the dog* sounds wrong — because the verb *want*, suggestive of propositional attitudes, seems incompatible with the nonsentient *tree*. Structurally, the problem with this sentence seems analogous to the flawed *The trees wants to run away*: the latter has incorrect singular/plural linkage, the former has incorrect sentient/nonsentient linkage, so to speak. But does this structural resemblance imply that singular/plural is as much part of semantics as grammar, or sentient/nonsentient as much part of grammar as semantics? It is true that there are no morphological markers for "sentience" or its absence, at least in English — except perhaps for "it" vs. "him/her" — but is this an accident of English or revealing something deeper?

In effect, assessments of propriety seem to operate on several levels. First (not to imply an actual temporal priority, though) we may consider fine-grained word-sense: which lexical entrant for such-and-such word is plausible in the current context? Then we can consider larger-scale Ontological criteria: is the subject of this sentence figured as material or immaterial, sentient or nonsentient, natural or sociocultural, spatially and/or temporally extended or pointwise; and so forth? And finally, what Part of Speech is consistent for various words given syntactic principles and morphological cues — distinguishing noun/verb/adjective, and etc., along with singular/plural, mass/count, verb tense and case, and other criteria of morphosyntactic fit?

So type-related observations can be grouped (not necessarily exclusively or exhaustively) into those I will call *macrotypes* — relating mostly to Parts of Speech and the functional treatment of phrases as applicative structures; *mesotypes* — engaged with existential/experiential qualities and "Ontological" classifications like sentient/nonsentient, rigid/nonrigid, and others I have discussed; and *microtypes* — related to lexemes and word-senses. This last lexical level can include "microclassification", or gathering nouns and verbs by the auxiliary prepositions they allow and constructions they participate in (such as, different cases), and especially how through this they compel various spatial and force-dynamic readings; their morphosyntactic resources for describing states of affairs; and, within semantics, when we look toward even more fine-grained classifications of particular word-senses, to reason through contrasts in usage.[10] Microclasses can point out similarities in mental "pictures" that explain words' similar behaviors, or study why different senses of one word succeed or fail to be acceptable in particular phrases. There are *stains*

[10] So, conceiving microclasses similar in spirit to Steven Pinker in Chapter 2 of [57], though I'm not committing to using the term only in the way Pinker uses it. Cf. also [77], which combines a microclass theory I find reminiscent of *The Stuff of Thought* with formal strategies like Unification Grammar.

© 2021/05/24 Nathaniel Christen

all over the tablecloth and *paint splattered all over the tablecloth*, but not (or not as readily) *dishes all over the tablecloth*. While "stains" is count-plural and "paint" is mass-aggregate, they work in similar phrase-structures because both imply extended but not rigid spatial presence; whereas "dishes" can work for this schema only by mentally adjusting to that perspective, spatial construal shifting from visual/perceptual to practical/operational (we might think of dishes "all over" the tablecloth if we have the chore of clearing them). Such observations support microclassification of nouns (and verbs, etc.) via Ontological and spatial/dynamic/configuration criteria.

Type-theoretic semantics can also apply Ontological tropes to unpack the overlapping mesh of word-senses, like *material object* or *place* or *institution*. This mode of analysis is especially well illustrated when competing senses collide in the same sentence. Slightly modifying two examples:[11]

- (63) The newspaper you are reading is being sued.
- (64) Liverpool, an important harbor, built new docks.

Both have a mid-sentence shift between senses, which is analyzed in terms of "type coercions" (see also [45] and [46]). The interesting detail of this treatment is how it correctly predicts that such coercions are not guaranteed to be accepted:

- (65) ?The newspaper fired a reporter and fell off the table.
- (66) ?Liverpool beat Tottenham and built new docks.

(again, slightly modifying the counter-examples). Type coercions are *possible* but not *inevitable*. Some word-senses "block" certain coercions — that is, certain sense combinations, or juxtapositions, are disallowed. These preliminary, motivating analyses carry to more complex and higher-scale types, like plurals (the plural of a type-coercion works as a type-coercion of the plural, so to speak). As it becomes structurally established that type rules at the simpler levels have correspondents at more complex levels, the use of type notions *per se* (rather than just "word senses" or other classifications) becomes more well-motivated.

Clearly, for example, only certain kinds of agents may have beliefs or desires, so attributing mental states forces us to conceive of their referents in those terms:

- (67) Liverpool wants to sign a left-footed striker.
- (68) That newspaper plans to fire its editorial staff.

This *can* be analyzed as "type coercions"; but the type-theoretic machinery should contribute more than just obliquely stating linguistic wisdom, such as maintaining consistent conceptual frames or joining only suitably related word senses. The sense of *sign* as in "employ to play on a sports team" can only be linked to a sense of Liverpool as the Football Club; or *fire* as in "relieve from duty" is only compatible with newspapers as institutions. These dicta can be expressed in multiple ways. But the propagation of classifications (like "inanimate objects" compared to "mental agents") through complex type structures lends credence to the notion that type-theoretic perspectives are more than just an expository tool; they provide an analytic framework which integrates grammar and semantics, and various scales of linguistic structuration. For instance, we are prepared to accept some examples of dual-framing or frame-switching, like thinking of a newspaper as a physical object and a social institution (but we reject other cases, like *Liverpool voted in a new city government and signed a new striker* — purporting to switch from the city to the Football Club). The rules for such juxtapositions appear to reveal a system of types with some parallels to those in formal settings, like computer languages.

In short, "Ontological" types like *institution* or *place* serve in some examples to partition senses of one multi-faceted word. Here they reveal similar cognitive dynamics to reframing-examples like *to the press*, where Ontological criteria (like reading something as a place) are triggered by phrase-scale structure. But there are also interesting contrasts: the *newspaper* and *Liverpool* examples imply that some words have multiple framings which are well-conventionalized; newspaper-as-institution feels less idiomatic and metaphorical than press-as-place. So these examples suggest two "axes" of variation. First, whether the proper Ontological framing follows from other word-choices (like "fire" in *the newspaper fired the reporter*, which has its own semantic needs), or from morphosyntax (like the locative in *to the press*); and, second, whether triggered framings work by selecting from established word senses or by something more metaphorical. Metaphors like *to the press* do have an element of stan-

[11] [16, p. 40] (former) and [48, p. 4] (latter).

© 2021/05/24 Nathaniel Christen

dardization; but apparently not so much so to be distinct senses: note how *the press* as metaphorical place does not work in general: ?*at the press*, ?*near the press* (but *at the newspaper*, *near the newspaper* — imagine two journalists meeting outside the paper's offices — sound quite reasonable).

The "type coercion" analysis works for mid-sentence frame-shifts; but other examples suggest a more gradual conceptual "blending". For example, the place/institution dynamic is particularly significant for *restaurant* (whose spatial location is, more so, an intrinsic part of its identity). Being a *place* implies both location and extension; most places are not single points but have an inside where particular kinds of things happen. I am not convinced that restaurant as place and as institution are separate word senses; perhaps, instead, conversations can emphasize one aspect or another, non-exclusively. We need not incorporate all framing effects via "subtypes" (restaurant as either subtype of hypothetical "types of all" places or institutions, respectively). But "placehood", the Ontological quality of being a place — or analogously being a social institution — identify associations that factor into cognitive frames; types can then be augmented with criteria of tolerating or requiring one association or another. So if "restaurant" is a type, one of its properties is an institutionality that *may* be associated with its instances. In conversation, a restaurant may be talked about as a business or community, foregrounding this dimension. Or (like in asking for directions) its spatial dimension may be foregrounded. The availability of these foregroundings is a feature of a hypothetical restaurant type, whether or not these phenomena are modeled by subtyping or something more sophisticated. The "newspaper" examples suggest how Ontological considerations clearly partition distinct senses marked by properties like objecthood or institutionality (respectively). For "newspaper" the dimensions are less available for foregrounding from a blended construal, than "unblended" by conventional usage; that is why reframings evince a type *coercion* and not a gentler shift of emphasis. The example of *restaurant*, in contrast, shows that competing routes for cognitive framing need not solidify into competing senses, though they trace various paths which dialogs may follow. But both kinds of examples put into evidence an underlying cognitive-Ontological dynamic which has potential type-oriented models.

At the most general level — what I called *macrotype* modeling — a type system recognizes initially only the grammatical backbone of expressions, and then further type nuances can be seen as shadings and interpretations which add substance to the syntactic form. So in type-theoretical analysis at this more grammatic level, we can still keep the more fine-grained theory in mind: the relation of syntax to semantics is like the relation of a spine to its flesh, which is a somewhat different paradigm than treating syntax as a logical or temporal stage of processing. Instead of a step-by-step algorithm where grammatical parsing is followed by semantic interpretation, the syntax/semantics interface can be seen as more analogous to stimulus-and-response: observation that a certain grammatic configuration appears to hold, in the present language artifact, triggers a marshaling of conceptual and cognitive resources so that the syntactic backbone can be filled in. Perhaps a useful metaphor is grammar as gravitation, or the structure of a gravitational field, and semantics is like the accretion of matter through the interplay of multiple gravitational centers and orbits. For this analogy, imagine typed lambda reductions like $\mathcal{P}rop \nrightarrow N \nrightarrow N$ taking the place of gravitational equations; and sentences' grammatic spine taking the place of curvature pulling mass into a planetary center.

As I have argued, sentences' progression toward complete ideas can be assessed more semantically — accretion of conceptual detail — or more syntactically, in terms of regulated type resolutions pulling in from a tree's leaves to its root. The latter model is a kind of schematic outline of the former, marking signposts in the accretion process rather like a meetings' agenda. Type theory allows points in conceptual accretion to be selected — corresponding to nested phrases — where type-checking signals that the accretion is progressing in an orderly fashion. Or, more precisely, type-checking acts as a window on a cognitive process; phrasal units are like periodic gaps in a construction wall allowing us to reconstruct interpretive processes, and the possibility of certain linguistic elements being assigned types marks the points where such windows are possible. So type theory can impose a formal paradigm on our assessment of sentence structure, but at the cost of sampling only discrete steps of an unfolding completion toward understanding. In practice, this discrete analysis should be supplemented with a more holistic and interpretive paradigm, which explores — perhaps speculatively, without demanding

© 2021/05/24 Nathaniel Christen

thorough formalization — the gaps between the formalizable windows. I will transition toward this style of analysis in the next section.

At the same time, I feel that the foundations of "cognitive" linguistics deserve a little more attention. I have used *cognitive* rather informally, depending on the intuitive picture of "cognitive" linguistics, grammar, or indeed "cognitive phenomenology", which emerges from the speculative project we associate with linguists/philosophers like George Lakoff, Mark Johnson, Leonard Talmy, Ronald Langacker, and Peter Gärdenfors — along with, as Gärdenfors points out, phenomenologists like Jean Petitot.[12] At the same time, Petitot also links with a tradition that combines elements of phenomenology and Analytic Philosophy, represented by such philosophers as Barry Smith and David Woodruff Smith and by "Analytic Phenomenology" or "Naturalizing Phenomenology" projects, the latter also being a large volumes Petitot co-edited.[13] "Cognitive" in these contexts tends to imply a desire to ground analyses on holistic human experience, in its experiential, embodied, pragmatically-oriented, first-personal, and intersubjective dimensions.[14]

Cognitive linguistics can be called "speculative" because its methodology generally relies on linguists' assessments of acceptability rather than empirical data from surveys, computational or statistical analyses of copora, or psycholinguistic studies of language processing or acquisition. Analytic Phenomenology is speculative in similar ways; although in some cases phenomenological structures are related to formal/mathematical theories like Mereotopology or Differential Geometry (cf. Barry Smith and Petitot, respectively, or Kit Fine's *Part-whole* [23]), phenomenologists' assessments of common perceptual patterns (and how they are situationally or conceptually interpreted and engaged with) is principally introspective. Both phenomenologists and cognitive linguists, in short, introspect on conscious and linguistic experience to identify patterns which they believe are not eccentric to their own cognition, but have some public disputability and theoretical merit. Of course, the overall dialog wherein philosophers debate and compare their own introspective reports allows this speculative method to have some rigor, because descriptions of cognitive processes — in their first-personal facticity — which seem both subjectively faithful and structurally revealing will emerge as analyses of general patterns so long as (and insofar as) multiple philosophical treatments agree on their fidelity to experience. So analyses get theoretically favored if they meet these different criteria of structural productivity — in the sense of producing new insight into cognitive processes, rather than just describing cognition as an experienced givenness — plus both faithfulness to each person's conscious experience and also generality to many people's experience.[15]

At their best, then, both phenomenology and cognitive linguistics combine introspective analysis and public disputation to develop theories of cognitive-experiential structures — of how the immediate structuration of perceptual experience as primordial conscious content unfolds into the schematic and rational models of our surrounding environment and situations, for purpose of goal-directed activity and inter-personal, collective reasonableness. The underlying assumption is that raw

[12] Gärdenfors mentions Lakoff, Langacker, Talmy, Fauconnier, and others alongside "a French semiotic tradition, exemplified by [Jean-Pierre] Desclés ... and [Jean] Petitot-Cocorda ... which shares many features with the American (mainly Californian) group" [26, p. 4].

[13] Petitot's and Barry Smith's formalizing projects were parallel and collaborative to some extent. Maxwell James Ramstead in a 2015 master's thesis reviews the history elegantly:

> Now, the "science of salience" proposed by Petitot and Smith (1997) illustrates the kind of formalized analysis made possible through the direct mathematization of phenomenological descriptions. Its aim is to account for the invariant descriptive structures of lived experience (what Husserl called "essences") through formalization, providing a descriptive geometry of macroscopic phenomena, a "morphological eidetics" of the disclosure of objects in conscious experience (in Husserl's words, the "constitution" of objects). Petitot employs differential geometry and morphodynamics to model phenomenal experience, and Smith uses formal structures from mereotopology (the theory of parts, wholes, and their boundaries) to a similar effect. [60, p. 38]

[14] Note that in this sense "Cognitive" connotes a very different perspective than this same term in AI research, for example; on the one side we have a philosophical commitment to the irreducibility of human reason to computable "symbol processing", whereas on the other there is a paradigm where, in effect (perhaps simply because mental activity is presumably reducible to low-level biological processes), there exists *some* computable core of cognition, which scientists can unlock to build powerful "Artificial General Intelligence".

[15] So for instance, Husserl's examination of protention and retention in sensing spatial form during perceptual episodes focused on discrete, extended objects (e.g. in *Thing and Space*), or his investigation of how intersubjectivity contributed to consolidating our conceptual integration of perceptual givens (e.g. in the *Cartesian Meditations*), can be considered classic phenomenological analyses because they have been deemed both experientially accurate and theoretically insightful by subsequent generations' worth of public review.

© 2021/05/24 Nathaniel Christen

structures, below the threshold of conscious deliberation, are enmeshed in immediate experience, and that from there we can identify ambient and situational patterns; mental representations of the structural and material properties of surrounding objects and places and the social/pragmatic rules governing interpersonal situations. We can then posit situational prototypes and morphological principles apparently operating in these representations, which become the basis of systematic theorizing of cognitive activity in general, including language.

In this paradigm, situational and organizational prototypes and patterns lend their structure to language, so — in many cases, i.e., with respect to many language artifacts — the scenarios influencing linguistic structure are these extra- or pre-linguistic gestalts rather than semantic to syntactic rules *per se*. But for this belief about the origin of (at least some) surface-level linguistic form to be leveraged as a diligent semantic or syntactic method, we need a systematic account of how phenomenological patterns evolve into (by grounding) linguistic structures. A thorough treatment of this problem is far beyond the scope of one paper, but I will offer a few ideas in the remainder of this section.

2.3 *Types and Phenomenology*

In their proper context, understanding linguistic expressions requires binding language to objects or "phenomena" in speakers' collective perceptual (or conceptual) horizon. Referents are not always material things; they can even be the *absence* of objects, or of substance:

- (69) There are footprints on the beach.
- (70) There's a hole in the bucket.
- (71) There are footprints leading up the hill.
- (72) There are train tracks leading up the hill.

In (69) our attention is directed not to any *object*, but to a certain perceptual pattern which has some factual significance, enough to warrant a distinct concept. Insofar as *one* footprint is a focus of attention, we notice some pattern of discontinuity which allows a foreground to emerge from a background — but there is a perceptual blend of discontinuity and continuity, since a footprint is literally situated in a material expanse with its surroundings. To focus on the footprint — a cognitive act which is both perceptual and conceptual — we have to retain awareness both of the footprint materially continuous with the surrounding sand (say) and also distinct from it via a somewhat different color or composition (e.g., the sand in the footprint may be darker or more compact than the sand around it). When talking about *footprints*, plural, we have to direct attention to a perceptual totality, something extending over our visual field and possessing inner parts, but also suggesting a conceptual totality, a worthiness of being cognized in aggregate.

In (71), moreover, we conceive the totality in conjunction with a sense of direction, and protention. The implication is that we do not see *all* of the footprints, but we can discern from their pattern a direction which, we anticipate, will reveal more footprints. Meanwhile the footprints are presumably disjoint, unlike train tracks. So the perceptual foreground in (71) is phenomenologically complex, including a totality we perceive that implies a greater totality, part of which is experienced anticipatorily rather than explicitly, along with a fragmentation which nonetheless permits a perceptual unification into a coherent whole. These kinds of perceptual/conceptual complexities in arrangements — blending continuity and discontinuity, wholeness and fragmentation, sensation and protention — are canonical to the phenomenology of attentional foci in any cognition engaged with ambient situations, which certainly includes language. We are not robots experiencing the world as a tableau of simply discrete, integral object-"things".

Language serves as a guide to negotiating the complexities of perceptual foci in inter-personal environments. We therefore have to understand how language "hooks" into conversants' phenomenological faculties. So we have, at a basic level, a contrast between situationally grounded or conceptually generalized references:

- (73) Salesmen are intelligent.
- (74) Salesmen are knocking on the door.
- (75) Should I let them in?

The effects of (74) (which with (73) are from the *Handbook*, example 44, page 139/162, chapter 6) are manifest in several changes to conversants' collective understanding: (74) establishes both a domain of potential reference (the salesmen could be subsequently identified as *the salesmen* or *those salesmen* or *them*), and the fact of

© 2021/05/24 Nathaniel Christen

their being at the door established as a basis for further dialog (e.g., (75)). The more generic (73) does not have any comparable situational effects, though it permits further dialog on a more generic plane.

Sometimes, of course, people dialog about things that are perceptually evident around them collectively. Probably more common, though, is that the structures of presentational perception translate to general cognitive patterns that are signified through language. That is to say, the rules of perceptual gestalts — the partial continuity and partial discontinuity between background and foreground; and the mixture of singular integrity and divisibility characterizing attentional foregrounds — are taken for granted as patterns typifying perception in general, whether that be perceptions presently occurring in the situational context or those conceptualized indirectly, abstractly, or hypothetically. How foci of attention figure into perceptual continua become part of their extralinguistic background, reflected in lexical conventions:

- (76) There are ink stains on the pages of this book.
- (77) There is a pretty sunset over the river.
- (78) In some neighborhoods, cicada insects make loud sounds from sunset to sunrise. `(en_gum-ud-train)`
- (79) The audience rose to give thundering applause.

The unifying theme in these examples is that their central concepts are experienced in terms of a matrix of phenomenological criteria that can be harnessed into a general theory: how an ink stain is (in part) materially continuous with the stained paper that extends around it; how *a pretty sunset* is an imprecisely bounded but conceptually integral atmospheric event; how (79) figures the audience as discrete individuals who, at that moment, act in such a way as to present a conceptual and perceptual totality; how (78) likewise proposes a totality but in a more complex fashion, because the speaker is describing a typical *kind* of experience rather than the specific sounds of cicadas on one specific occasion.

The matrix of continuity/discontinuity, and individual coherence balanced (within intentional "noemata") against internal structuration and diversity, are *conceptually* intrinsic to notions like *ink stains* (or *footprints*), *applause*, *sunsets*, and even *cicadas* and *noises*. They become part of a language's lexical machinery and therefore embed, in lexical conventions, phenomenological prototypes that imply a situational background, called forth by the lexical potency of the specific words. The point is not that the hearer of (77), say, sees the sunset also, and thus *explicitly* undergoes the disclosure-experience of a sunset (in all its sensory vividness and individuating vagueness). Instead, the patterns of sensation and (dynamic, imprecise, un-fixed) individuation are abstracted to prototypes poised within the lexicon itself to be perceived as applicable to the talked-about situation. We have some idea of the *kind* of experience that is involved in a sunset, not just perceptually but in the nexus of conceptual and rational processes which leads us to identify the sunset as such and estimate its phenomenal specificity — its spatial form, its temporality — so that it has some determinate epistemic content, something that can be discussed with others (*Did you see the sunset*; *Is it too late to see the sunset*; etc.).

Phenomenological patterns in construing conceptual and referential foci, solicited by lexical and idiomatic conventions, give language flexibility and rhetorical flair, often bounding expressive possibilities on phenomenological rather than narrowly semantic grounds. I would dispute some of Nunberg's analyses I mentioned in the last section, for example, on the dubiousness of (27) or the exceptionality of (31). I do not find (27) exceptionally jarring, considering the situational construal common to (27) along with (25) and (26): the formation *I am parked* is not *just* a matter of a person proxy-referring their car. Usually a person saying *I am parked* is *in* their car, so they are describing a larger situation — the conceptual foreground is the totality of the car and themselves in it. Were the addressee to find them, her attention would be directed to the car *and* the speaker, as a complex but integral (in that episode of time) whole. This then carries over to cases where the speaker is *not* in their car. Saying *I am parked* (along with a place-description) outlines the steps needed to return to the car; it is a way of framing the car's location operationally. Using the first person for where the car is parked — even when the speaker is not at that spot — indicates that the speaker is conceiving the car's location in first-personal terms; that is, in terms of the actions the speaker is taking or anticipates taking.

This analysis then implies that (27) is acceptable on similar operational grounds; first-personalizing the car's location, so to speak, foregrounds that location only in the context of an operational totality involving (presum-

© 2021/05/24 Nathaniel Christen

ably) going to the car and then driving it. If the speaker is concerned about the car not starting, this is relevant to how the overall situation is cognized, and therefore the transition from the speaker-centered *I am parked ...* to the discordant *and may not start* is understandable. I would similarly argue that the apparent difference between (31) and (29)-(30) depends on how referential foci are established. The phrase *a beer* can designate a specific glass or bottle of beer — something perceptually and enactively/kinaesthetically individuated — or also a specific preparation of beer, individuated by the unique flavor of the beverage rather than the material identity of the liquid. These differences have phenomenological overtones: if I say *I drank two beers* I could mean two glasses, pints, or bottles; I could also mean two *kinds* of beer (maybe multiple pints of each). The phrasing itself is ambiguous, but the ambiguity is extra-linguistic: it relates to how linguistic content binds to empirical givens with their rules of phenomenological individuation. The "two pints" reading implies one kind of phenomenological architecture governing "word-to-world fit", where elements in language link to discrete, perceptually/operationally integral objects like glasses (though of course a glass of beer is also an integral complex where the beer and the glass are semi-autonomous parts). The "two kinds" reading is more subtle, resting conceptual focus on the belief that a distinct kind of beer has a distinct flavor that is both unique vis-à-vis other beers and also consistent across bottles (or kegs), so it has *a* taste that people can jointly experience and talk about. Nunberg makes the reasonable claim that we are more likely the hear the former interpretation when it comes to beers, and the latter one for Sauternes, but I can imagine plausible conversations where these conventions would be reversed.

To return to type theory, then, this style of formalization is a potential window onto how the lexicalized concepts negotiate the phenomenological options in "binding" words to phenomena. The matrix of continuity/discontinuity and individuation (singularity)/complexity (internal structure or diversity) forms a tableau which different word senses and, in explicit phrase contexts, different usages "hook into" in different ways. So beer *qua* liquid has one conventional pattern in "word-to-world fit", while beer *qua* consumer product, or a brewer's creative endeavor, evinces a different phenomenological pattern. This is modeled, to some approximation, by the "type" (or as I put it "meso-type") distinction of a liquid (more generally a substance) against a consumer product or social good (more generally a socio-cultural artifact). We can refer to beers in both senses, and I have mentioned theories of type "coercions" or juxtapositions (cf. Pustejovsky's "dot product" theories in [59] or [5]) that explore when the different "Ontological" construals or sentences can be combined or alternated. At this point I would add that these Ontological details are not only relevant to semantics; they also govern how phenomenological patterns can be reified and typified in lexical and idiomatic conventions.

For successful conversation, participants need to converge for each sentence on a common conceptual focus; not surprisingly, this process often reciprocates the phenomenology of perceptual focus, or enactive/kinaesthetic attention. Drawing form from phenomenological structure, referential signification often takes shapes that can seem opaque on purely logical considerations. In analysis of referring expressions, then — analogously to semantic analyses as earlier in this section — we need to be sensitive to the possibilities of language form molding to perceptual and situational gestalts rather than predicate structure.

Nunberg, for instance, investigates how the proxying effect in *I am in the Whitney* can be generalized to other cases. I'll point out that this fits a not-uncommon rhetorical pattern, as in:

- (80) I am in the Hall of Fame.
- (81) I am on TV.

Nunberg argues that reference, in these kinds of cases, will "transfer" between conceptually linked designata, so I can refer to my car's location as my location, or to my painting as myself. If we analyze this logically, the rule appears to be that I can substitute first-person reference for reference to something associated with myself (that I own or have created, etc.). He then finds that the pattern does not generalize to a scenario such as a painter — referring to the location of her work in transit — saying something like (his example 12):

- (82) I'm in the second crate on the right.

© 2021/05/24 Nathaniel Christen

In other words, Nunberg's analysis turns on theorizing various "I am in..." constructions as a kind of referential transfer or (as I would say) proxying, and then seeking the logical rules behind how and when such proxying works in a first-person (morphosyntactic) context.[16]

On the other hand, I would argue that the pattern in (80)-(81) is not reducible to referential proxying as a logical maxim. We can speculate that *I* obliquely references, say, the bust and info about an athlete who has been elected to a Hall of Fame, or to the image of a TV personality which viewers see on screen. So analyzing this construction as a case where reference is transferred from the "subject of enunciation" to some subordinate vehicle — her bust, image, and so forth — is a plausible gloss on the pattern. But looking for a single maxim to accommodate these different cases overlooks the contextual particulars, in particular the backstory which is lexically embedded in expressions like *Hall of Fame* and *on TV*. There is a long process leading to retired athletes being recognized as worthy of a Hall; there is also a long process whereby personalities get a chance to appear on television. Actually, (81) has two interpretations — someone could be on TV just momentarily, e.g. as a bystander during on-location news coverage — or as someone like a policy expert who is interviewed on occasion. But the more interesting reading of (81) involves a person being described as "on TV" not on the occasion of their appearing on-screen in that moment, but as a recurring event. In this case *being on TV* encapsulates a backstory reflecting a measure of personal success and esteem, not unlike an athlete in a Hall of Fame or an artist in a prestigious museum.

The relevant backstories govern interpretations for first-person reference: *I am in the Hall* or *on TV* indirectly reports that the speaker has undergone some process which the addressee, assumed competent with the relevant lexicon, at least minimally understands. So a pattern like (80) packages up that backstory *in the guise* of a referring expression; it is not so much the athlete's bust as a target of referential proxying, but more the act of referring to that bust (i.e., to the referentiable object manifesting the speaker's being elected) entering the Hall of Fame backstory into the discursive ledger. This carries over to the Whitney case: calling oneself "in" a museum does not just proxy the self for one's art, but uses patterns of reference to the art work as a vehicle for introducing the backstory of one's rise through the art world (or whatever is the relevant autobiographical context) into the conversation.

I would argue, in short, that referential proxying is not a primarily logical operation, but relies on backstory context to regulate how referring patterns are received. To establish conceptual foci, addressees have to negotiate the language given to them, in search of the focal element or foreground, often an explicit or hypothetical perceptual nexus. So a painter *in the Whitney* means that a visitor can, in the right room and orientation, perceptually encounter her work; a personality *on TV* means that the addressee can, at times, see her on-screen. Those potential perceptual givens ground the semantics of expressions like (28) or (81). This does not mean that the sentences are only meaningful to an addressee who wants to go find the speaker at the museum or on the television. But it does present the perceptual ground as a foundation for the conceptual implications which lead from that perceptual situation: from the perceptual presence of an art work in a museum to the backstory of its provenance and implications for the artist's career or place in history; or from the on-screen image to the context of how TV shows are produced and the careers and reputations of the figures that show up.

We approach referring expressions in these contexts at two levels: we identify the canonical perceptually-oriented references which supply a conceptual ground (the painting as perceptual object, the on-screen image as perceptual simulacrum); but then we recognize the backstory and situational context which determines how the perceptual ground should be understood. Usually the discourse is "about" that backstory, not about the precise situations where explicit perceptual grounds would be relevant — an artist probably would not say (28) only to someone looking to visit a museum to see her art work. The "I am in..." pattern turns on how these two different levels are played off in interpretation. In particular, the addressee needs to ascertain why the speaker's information is being proffered: is boasting of their art being hung in an elite museum mostly a way of conveying the speaker's status, so the main theme is essentially autobiographical; or is the actual location of the painting

[16] Obviously "first person" in this setting concerns verb tense and other linguistic cues linking to the speaker/enunciator of a piece of language; the same term is also encountered (including elsewhere in this paper) in the phenomenological (and Philosophy-of-Mind) sense of conscious, intentional experience (in contrast to experience, or thoughts and feelings, we attribute to others based on their behavior).

© 2021/05/24 Nathaniel Christen

directly relevant to the present conversation?

A statement like (82) (the speaker being "in the crate on the right") does not have the more "autobiographical" reading, since (absent some bizarre imaginary scene) there is no status or reputation attached to which crate is carrying your paintings. But the second kind of reading, where specific location is operationally relevant, is plausible in some contexts; perhaps several different artists' works are in several different crates and the speaker wants to direct focus onto the one holding her own works. So, in the proper context, I have no objection to (82) or to *I am in that crate* (forms which Nunberg rejects). Sure, we need to imagine peculiar scenarios to make these plausible; but almost all linguistic meaning plays out in the context of *some* scenario — some kinds of situations forming the backdrop to discourse are more typical than others, which can bias our hearing linguistic acceptability, but situational idiosyncracy is not a linguistic phenomenon per se (*in* the scenario, conversants don't experience their discourse as unusual). The gist in the "crate" examples, as in (28), is to establish conceptual focus onto a painting or art work, and to do so via a first-person referential lead-in. So the situation determines first *why* the focus is thus singled out and second *how* the first-person reference can proxy that focus. In (28), we single out a painting insofar as it is a product of the artist's creativity, so appreciation of the work is manifestly appreciation of the artist. The artist is then referentially linked to the work because the backstory of how museums acquire art works overlaps with the story of artists' careers and reputations.

In (82) the situation is different: presumably the focus is on one or several art works because of some concern about transporting or accessing them, so we're attending to the works in their guise as (fragile) physical objects. Moreover the link between the artist and the work, which guides the referential proxying, is presumably that the artist is concerned about accessing and protecting those objects. The situational elements are different, but in both these cases — whether paintings are in the crate or in the Whitney — there is some structural resonance in how situational, backstory, and hypothetical-perceptual gestalts are all integrated into patterns of grounding referential interpretation in ambient conceptual contexts. These points seem like additional datums pointing to the "Whitney" and "crate" examples being more similar than they may appear.

I think that linguists sometimes underestimate the multiplicity of layers — situational, contextual, phenomenological (explicit and hypothetical), referential, contextual — that all converge on meaning and reference. Perhaps this tendency can be empirically examined by testing judgments of acceptability: the more that we incorporate multiple layers of possible context, the wider becomes the circle of sentences that feel reasonable (cf. the reasonableness, in my estimation, of (82) and (31)).[17] Failure to identify these layers' workings can lead analysis toward searches for reproducible logical rules governing which constructions are accepted by a language-community as recurring patterns, and logical explanations for the limits on that generalizability. This is one example of overestimating the logicality of language in general, an issue I have approached in this section from both semantic and reference-theoretic angles. I will focus on this issue of logicality — both its applicability and its limits — in the next section.

3 Gaps in Truth-Theoretic Semantics

I consider as a given that typical sentences have a propositional core, against which they take a performative stance (which can be outright assertion, or else asserting speakers' more complex propositional attitudes). I would further say that *truth-theoretic* semantics, in particular, is organized around this propositional content as the core target of linguistic analysis. I am thinking of truth-theoretic semantics in a broad sense, perhaps the most influential paradigm in the Philosophy of Language and by extension philosophy and linguistics in general (not to mention Computer Science and Artificial Intelli-

[17] Granted that, perhaps, with a lot of imagination almost any construction could be deemed plausible in *some* context. Accordingly, one could argue that it is analytically reasonable to distinguish sentences that exemplify proper language in a wide range of contexts, as compared to sentences which could only be meaningful in very select circumstances. Even on this perspective, however, we have to unpack the distinction between "generic" from "select" circumstances — what qualities of these situations, together with relevant word-meanings, make these examples of broad usage patterns as against unexpected usages that are nonsensical without a very specific background? Rather than being a neutral arbiter of acceptability, such genericity is a phenomenon that needs to be explained. I would argue that assessments of unacceptability will in many cases overrate the distinction between "normal" and "exceptional" circumstances. For example, I dispute that the beer/wine contrast, or painting in a museum or in a crate contrast, are such a divergence between normal and unusual contexts that (82) has essentially different plausibility than (28), or (31) compared to (29) and (30).

© 2021/05/24 Nathaniel Christen

gence research). The most notable counter-paradigm is Cognitive Linguistics; consider George Lakoff and Mark Johnson's extended critique of truth-theoretic paradigms in *Philosophy in the Flesh* [39]. Adherents of the latter perspective need not dispute the logical substance of language artifacts' propositional content, but tend to direct theoretical attention not to the nature of propositional content itself, but to the cognitive processes through which this content is understood.

As I argued earlier, many sentences do not simplistically reproduce the logical structure of their propositional content, so models of that structure are only tangentially relevant to analysis on the language side. This is why we need distinct analyses, beyond a mere logical gloss, covering the interpretive steps leading to holistic sentence-understanding. This section will consider several cognitive and pragmatic themes moving toward a general theory of this phenomenon.

3.1 *Enaction and Illocutionary Force*

I will start by reviewing illocutionary pragmatics, to identify some of the contextual and interpretive transformations that pertain to mapping surface language to propositional contents. My point is to establish what should be a common theory of logicality that can be shared by both critics and defenders of "truth-theoretic" paradigms, on which basis their legitimate disputes can be investigated.

Many linguists (on both sides, I would say, of my central truth-theoretic pro/con) seem to analyze hedges like "could you please" as merely dressing over crude commands: we don't want to come across as giving people orders, but sometimes we do intend to ask people to do specific things. As a result, we feel obliged to couch the request in conversational gestures that signal our awareness of how bald commands may lie outside the conversational norms. These ritualistic "could you please"-like gestures may have metalinguistic content, but — so the theory goes — they do not *semantically* alter the speech-act's directive nature.

The problem with this analysis is that sometimes directive and "inquisitive" dimensions can overlap:

- (83) Do you have almond milk?
- (84) Can you get MsNBC on your TV?
- (85) This isn't a screw-cap bottle: I need a corkscrew.

These *can* be read as bare directives, and would be interpreted as such if the hearer believed the speaker already knew that yes, he has almond milk, and yes, he gets MsNBC. In (85), if both parties know there's one corkscrew in the house, the statement implies a directive to fetch *that* corkscrew. But, equally, (83)-(85) can *also* be read as bare questions with no implicature: say, as fans of almond milk and MsNBC endorsing those selections, or pointing out that opening the bottle will need *some* corkscrew. And, meanwhile, (83)-(85) can *also* be read as a mixture of the two, as if people expressed themselves like this:

- (86) I think the window is open, can you close it?
- (87) I see you have almond milk, can I have some?
- (88) If you get MsNBC, can you turn on Rachel Maddow?
- (89) If there is a corkscrew in the house, can you get it?

I think the mixed case is the most prototypical, and pure directives or inquiries should be treated as degenerate structures where either directive or inquisitive content has dropped out. After all, even a dictatorial command includes the implicit assumption that the order both makes sense and is not impossible. On the other hand, we don't ask questions for no reason: "do you have almond milk" may be a suggestion rather than a request, but it still carries an implicature (e.g., that the addressee *should* get almond milk).

Ordinary requests carry the assumption that addressees can follow through without undue inconvenience, which includes a package of assumptions about both what is currently the case and what is possible. "Close the window" only has literal force if the window is open. So, when making a request, speakers have to signal that they recognize how the request involves certain assumptions and are rational enough to accept modifications of these assumptions in lieu of literal compliance. This is why interrogative forms like "can you" or "could you" are both semantically nontrivial and metadiscursively polite: they leave open the possibility of subsequent discourse framing the original request just as a belief-assertion. Developments like *can you open the window — no, it's closed* are not ruled out. At the same time, interrogative forms connote that the speaker assumes the addressees

© 2021/05/24 Nathaniel Christen

can fulfill the request without great effort: an implicit assumption is that they *can* and also *are willing to* satisfy the directive. This is an assumption, not a presumption: the speaker would seem like a bully if he acted as if he gave no thought to his demands being too much of an imposition — as if he were taking the answer to "can you" questions for granted. This is another reason why requests should be framed as questions. So, in short, "commands" are framed as questions because the speaker literally does not know for sure whether the command is possible; given this uncertainty a command *is* a question, and the interrogative form is not just a non-semantic exercise in politesse.

Sometimes the link between directives and belief assertions is made explicit. A common pattern is to use *I believe* or *I believe that* as an implicature analogous to interrogatives:

- (90) I believe you have a reservation for Jones?
- (91) I believe this is the customer service desk?
- (92) I believe we ordered a second basket of garlic bread?
- (93) I believe you can help me find computer accessories in this section?

These speakers are indirectly signaling what they want someone to do by openly stating the requisite assumptions — *I believe you can* in place of *can you?*. The implication is that such assumptions translate clearly to a subsequent course of action — the guest who *does* have that reservation should be checked in; the cashier who *can* help a customer find accessories should do so. But underlying these performances is recognition that illocutionary force is tied to background assumptions, and conversants are reacting to the propositional content of those assumptions as well as the force itself. If I *do* close the window I am not only fulfilling the request but also confirming that the window *could* be closed (a piece of information that may become relevant in the future).

In sum, when we engage pragmatically with other language-users, we tend to do so cooperatively, sensitive to what they wish to achieve with language as well as to the propositional details of their discourse. But this often means that I have to interpret propositional content in light of contexts and implicatures. Note that both of these are possible:

- (94) Do you have any milk?
- (95) Yes, we have almond milk.
- (96) No, we have almond milk.

A request for milk in a vegan restaurant could plausibly be interpreted as a request for a vegan milk-substitute. So the concept *milk* in that context may actually be interpreted as the concept *vegan milk*. Responding to the force of speech-acts compels me to treat them as not *wholly* illocutionary — they are in part statements of belief (like ordinary assertions). One reason I need to adopt an epistemic (and not just obligatory) attitude to illocutionary acts is that I need to clarify what meanings the speaker intends, which depends on what roles she is assigning to constituent concepts.

Suppose my friend says this, before and after:

- (97) Can you put some almond milk in my coffee?
- (98) Is there milk in this coffee?

Between (97) and (98) I do put almond milk in his coffee and affirm "yes" to (98). I feel it proper to read (98)'s "milk" as really meaning "almond milk", in light of (97). Actually I should be *less* inclined to say "yes" if (maybe as a prank) someone had instead put real (animal) milk in the coffee. In responding to his question I mentally substitute what he almost certainly *meant* for how (taken out of context) (98) would usually be interpreted. In this current dialog, the *milk* concept not only includes vegan milks, apparently, but *excludes* actual milk.

It seems — on the evidence of cases like this one — as if when we are dealing with illocutionary force we are obliged to subject what we hear to extra interpretation, rather than resting only within "literal" meanings of sentences, conventionally understood. This point is worth emphasizing because it complicates our attempts to link illocution with propositional content. Suppose grandma asks us to close the kitchen window. Each of these are plausible and basically polite responses:

- (99) It's not open, but there's still some cold air coming through the cracks.
- (100) It's not open, but I closed the window in the bedroom.
- (101) I can't — it's stuck.

In each case I have not fulfilled her request vis-à-vis its literal meaning, but I *have* acted benevolently in terms

© 2021/05/24 Nathaniel Christen

of conversational maxims. Similarly, the *Handbook* has this case (example 12, page 180/203, chapter 8):

- (102) The window, it's still open.
- (103) A window, it's still open.

Chapter authors Jeanette K. Gundel and Thorstein Fretheim suggest that (103) is dubious, but in context it may make perfect sense — particularly if it follows a discourse where a requester wanted *whatever* window closed (whichever window was causing a draft); even if that wish was *expressed* via a *close the window*.[18]

Part of reading propositional content is syncing our conceptual schemas with our fellow conversants. The illocutionary dimension of a request like *can I have some milk?* makes this interpretation especially important, because the addressee wants to make a good-faith effort to cooperate with the pragmatic intent of the speech-act. But cooperation requires the cooperating parties' conceptual schemas to be properly aligned. I therefore have to suspend the illocutionary force of a directive temporarily and treat it as locutionary statement of belief, interpret its apparent conceptual underpinnings in that mode, and then add the illocutionary force back in: if I brought *real* milk to a vegan customer who asked for "milk" I would be *un*-cooperative.

[18] Consider also the case of (*Handbook*, example 9, page 132/155, chapter 6):

- (104) He scribbled on a living-room wall.

Barbara Abbott (chapter author, commenting on examples/discussion from John W. Dubois) finds the indefinite article in (104) awkward. Here, like in (103), I think the acceptability of *both* definite and indefinite articles points to the flexibility of articles in English: we can use *close the window* even if we are referentially ambiguous about which window is open — both (103) and (102) are almost interchangeable. The reason is apparently that the very act of talking about an open or closed window foregrounds it such that it can be approached definitely. In (104) Abbott's intuition is probably that we usually speak of *the* wall of a room — even in the normal case of a room with four walls — because there is no common reason to single out one wall, as with an indefinite article (saying *a* wall implies that the other walls are differentiated from the one referenced). But I believe we can clearly cognize the walls of a room as a collection of discrete things, and that the formation *the* wall — unifying them into a single — is not so much a matter of discounting the perceptual multiplicity of four walls, but of conceiving "walls" themselves functionally as well as perceptually. In my reading, *the wall* refers not only to something perceptually individuated but to an element of the architectural complex of a building: a wall is something that prohibits movement, muffles sounds, provides some privacy, protects people inside the room (part of a building's structural integrity), and so forth. We can say *the* wall because we cognize walls as occupying that phenomenal niche, so all four walls of a room are collectively "the" wall vis-à-vis such a niche, even while on sensory grounds we can switch to treating them as separate (warranting the (104) version).

The upshot is that conversational implicatures help us contextualize the conceptual negotiations that guarantee our grasping the correct propositional contents, and vice-versa. This means that propositionality is woven throughout both assertive and all other modes of language, but it also means that propositional content can be indecipherable without a detailed picture of the current context (including illocutionary content). The propositional content of, say, *there is milk in this coffee* has to be judged sensitive to contexts such as *milk* meaning *vegan milk* — and this propagates from a direct proposition to any propositional attitudes that may be directed towards it, including requests like *please put milk in this coffee*.

Suppose the grandkids close grandma's bedroom window when she asks them to close the kitchen window. The propositional content at the core of grandma's request is that the kitchen window be closed; the content attached to it is an unstated belief that this window is open. Thus, the truth-conditions satisfying her implicit understanding would be that the kitchen window went from being open to being closed. Suppose, as it happens, that window is already closed. So the truth-conditions that would satisfy grandma's initial belief-state do not obtain — her beliefs are false — but the truth conditions satisfying her desired result *do* obtain. The window *is* closed. Yet the grandkids should not thereby assume that her request has been properly responded to; it is more polite to guess at the motivation behind the request, e.g., that she felt a draft of cold air. In short, they should look outside the truth conditions of her original request taken literally, and *interpret* her request, finding different content with different truth-conditions that are both consistent with fact and address whatever pragmatic goals grandma had when making her request. They might infer her goal is to prevent an uncomfortable draft, and so a reasonable "substitute content" is the proposition that *some* window is open, and they should close *that* one.

So the grandkids should reason as if translating between these two implied meanings:

- (105) I believe the kitchen window is open — please close it!
- (106) I believe some window is open — please close it!

They have to revise the simplest reading of the implicit propositional content of grandma's *actual* request, be-

© 2021/05/24 Nathaniel Christen

cause the actual request is inconsistent with pertinent facts. In short, they feel obliged to explore propositional alternatives so as to find a replacement, implicit request whose propositional content *is* consistent with fact and also meets the original request's illocutionary force cooperatively.

In essence, we need to express a requester's desire as itself, in its totality, a specific propositional content, thinking to ourselves (or even saying to others) things like:

- (107) Grandma wants us to close the window.
- (108) He wants a bottle opener.

But to respond politely we need to modify the parse of their requests to capture the "essential" content:

- (109) Grandma wants us to eliminate the cold draft.
- (110) He wants something to open that bottle.

We have to read outside the literal interpretation of what they are saying. This re-reading is something that may be appropriate to do with respect to other forms of speech also; but our conversational responsibility to infer some unstated content is especially pronounced when we are responding to an explicit request for something.

Certainly, in many cases, meanings are not literal. But how then do we understand what people are saying? Trying to formulate a not-entirely-haphazard account of this process, we can speculate that interpreting what someone is "really" saying involves systematically mapping their apparent concepts and references to some superimposed inventory designed to mitigate false beliefs or conceptual misalignments among language users in some context. That means, we are looking for mappings like *milk* to *almond milk* in (111) from a vegan restaurant, or *kitchen window* to *bedroom window* in (112) if it is the latter that is open:

- (111) Can I have some milk?
- (112) Can you close the kitchen window?

The point of these "mappings" is that they preserve the possibility of modeling the *original* propositional content by identifying truth conditions for that content to be satisfied.

A *literal* truth-condition model doesn't work in cases like (111) and (112): the diner's request is *not* satisfied if it is the case that there is now (real) milk in her coffee; and grandma's request is not necessarily satisfied if it is the case that the kitchen window is closed. The proposition "the kitchen window is closed" only bears on grandma's utterance insofar as she believes that this window is open and causing a draft. So if we want to interpret the underlying locutionary content of (111) and (112) truth-theoreticaly, we need to map the literal concepts appearing in these sentences to an appropriate translation, a kind of "coordinate transformation" that can map concepts onto others, like milk/almond milk and kitchen window/bedroom window.

In sum, a theory of sentences' logical nexus can only be complete with some model of discursive context *structured in such a way* that we can repesent the interpretations and concept-transforms internal to parsing sentences to their propositional core. I will now consider what such a "theory of context" might look like.

3.2 The co-framing system and the doxa system

Illucutionary acts expressly signify our desire for something to change in our environment (with the help of our addressees), but similar implications of pragmatic desire are evident even when sentences are more directly assertorial, or less directly illocutionary. Compare between:

- (113) Remember that wine we tasted on the Niagara Peninsula last summer? Can you find it in our local liquor store?
- (114) Remember that wine we tasted on the Niagara Peninsula last summer? What varietal was that again?

The first sentence in each pair attempts to establish a common frame of reference between addresser and addressee — it does not, in and of itself, request any practical (extramental) action. The second sentence in (113) *can* be read as requesting that the addressee buy a bottle, though an alternate interpretation is to learn for *future reference* whether someone *could* buy that bottle. The second sentence in (114) carries no directive implicature at all, at least with any directness; it asks for more information.

© 2021/05/24 Nathaniel Christen

Despite these variations, it seems reasonable to say that language is always performed in an overarching setting where concrete (extralinguistic) activity will *eventually* take place. If in (114) I intend to recommend that grape variety to a friend, I may not be making a direct request of him, but I *am* proposing an eventual action that he might take. If in (113) I am not issuing a directive, I am however establishing (and reserving the future possibility) that such a directive would be reasonable. As a result, some extralinguistic state change seems to be lurking behind the linguistic content: I want my friend to go from having never tasted that varietal to having tasted it. Or I want to go from not having a bottle of that wine to having one. Or, if I do not want these things at the moment, I want to confirm intellectually that these wishes are plausible. We seem to use language to set up the interpersonal understandings needed to *eventually* engage in (usually collective) practical activity, which means effectuating some (extralinguistic) change.

That is, most expressions are not direct requests or suggestions of the "close the window" or "let's get some wine" variety, but they are stitches in the thread of coordinated human actions. Often however we use language to *prepare*, *negotiate*, and *decide upon* joint actions. We may have a *holistic* sense that meanings orbit around extralinguistic and extramental state-change, but at the level of particular sentences most changes that occur, or are proposed, tend to be changes in our conceptualization of situations. Accordingly, we can pursue a semantic theory based on *change of state* if we accept that such changes run the gambit from changes *internal* to language — to conversants' appraisal of dialogic context — to changes effectuated by human activity inspired by language. Dialogs themselves change: the first sentences in (113) and (114) modify the discursive frame so that, for example, a particular wine becomes available as the anaphoric target for *that* and *that wine* — and also, metonymically, *that varietal*, *that grape*, *that winery*. Conceptual frames can change: if we are discussing a visit to Ontario and I mention a particular winery, one effect is to (insofar as the conversation follows my lead) refigure our joint framing to something narrower and more granular that the prior frame (but still contained in it; I am not changing the subject entirely). We can pull a frame out as well as in — e.g., switch from talking about one winery visit to the whole trip, or one Leafs game to the entire season. Moreover, our beliefs can change/evolve: if you tell me the wine was Cabernet Franc, I have that piece of info in my arsenal that I did not have before.

So I assume in this paper that linguistic meanings are grounded in state-changes, with the stipulation that the "register" where the changes occur can vary over several cognitive and extramental options: actual change in our environment (the window closed, milk in the coffee, the bottle opened); changes to the dialog structure (for anaphoric references, pronoun resolution, metalinguistic cues like *can you say that again*, etc.); changes to conceptual framings (zoom in, zoom out, add detail); changes to beliefs. Each of these kinds of changes deserve their own analysis, but we can imagine the totality of such analyses forming an umbrella theory of meanings.

During the course of a conversation — and indeed of any structured cognitive activity — we maintain conceptual frames representing relevant information; what other people know or believe; what are our goals and plans (individually and collectively); and so forth. We update these frames periodically, and use language to compel others to modify their frames in ways that we can (to some approximation) anticipate and encode in linguistic structure.

In the simplest case, we can effectuate changes in others' frames by making assertions they are likely to believe to be true (assuming they deem us reliable). In general, it is impossible to extricate the explicit content of the relevant speech-acts from the relevant cognitive, linguistic, and real-world situational contexts:

- (115) That wine was a Cabernet Franc.
- (116) Those dogs are my neighbor's. They are very sweet.

Although there is a determinate propositional content being asserted and although there is no propositional attitude other than bald assertion to complicate the pragmatics, still the actual words depend on addressees drawing from the dialogic context in accord with how I expect them to (as manifest in open-ended expressions like *that wine*, *those dogs*, *they*). Moreover, the open-ended components can refer outward in different "registers": in *that wine* I may be referencing a concept previously established in the conversation, while *those dogs* may refer to pets we saw or heard but had not previously talked about. Of course, the scenarios could be reversed: I could introduce *that wine* into the conversation by

© 2021/05/24 Nathaniel Christen

gesturing to a bottle you had not noticed before, and refer via *those dogs* to animals you have never seen or heard but had talked about, or heard talk about, in the recent past. These dialog steps need to be resolved via a mixture of linguistic and extra-linguistic cues: surface-level language is not always clear as to whether referring expressions are to work "deictically" (drawing content from the ambient context, signified by gestures, rather than from any linguistic meaning proper), "discursively" (referring within chains of dialog, e.g. anaphora), or "descriptively" (using purely semantic means to establish a designation, like "my next-door neighbor's dogs" or "Inniskillin Cabernet Franc Icewine 2015").

Let's agree to call the set of entities sufficiently relevant to a discourse or conversation context the *ledger*. By "sufficiently relevant" I mean whatever is already established in a discourse so it can be referenced with something less than full definite description (and without the aid of extralinguistic gestures). I assume that gestures and/or descriptions are communicative acts which "add" to the ledger. The purely linguistic case — let's say, *descriptive additions* — can themselves be distinguished by their level of grounding in the current context. A description can be "definitive" in a specific situation without being a *definite description* in Bertrand Russell's sense (see "that wine we tasted last summer").

So, descriptive additions to the ledger are one kind of semantic side-effect: we can change the ledger via language acts. I will similarly dub another facet of cognitive-linguistic frames as a *lens*: the idea that in conversation we can "zoom" attention in and out and move it around in time. "That wine we tasted last summer in Ontario" both modifies the *ledger* (adding a new referent for convenient designation) and might alter the *lens*: potentially compelling subsequent conversation to focus on that time and/or place. Finally, I will identify a class of frame-modifications which do directly involve propositional content: the capacity for language to promote shared beliefs between people whose cognitive frames are in the proper resonance, by adding details to conceptual pictures already established: *those dogs are Staffordshires*, *that wine is Cabernet Franc*, *we have almond milk*, etc.

For sake of discussion, I will call this latter part of the "active" cognitive frame, for some discussion — the part concerning shared beliefs or asserted facts — the *doxa inventory*. This "database"-like repository stands alongside the "ledger" and "lens" to track propositional content asserted, collectively established, or already considered as background knowledge, vis-à-vis some discourse. Manipulations of the lens and ledger allow speakers to designate (using referential cues that could be ambiguous out-of-context) propositional contents which they wish to add to the "doxa inventory". I'll also say that modifying this inventory *can* be done through language, but participants in a discourse are entitled to assume that everyone formulates certain beliefs which are observationally obvious, and can therefore be linguistically presupposed rather than reported (the likes of that a traffic light is red, or a train has pulled into a station, or that it's raining).

So, I will assume that the machinery of frames is cognitive, not just linguistic. We have analogous faculties for "refocusing" attention and adding conceptual details via interaction with our environment, both alone and with others, and both via language and via other means. Some aspects of *linguistic* cognitive framing — like the "ledger" of referents previously established in a conversation — may be of a purely linguistic character, but these are the exception rather than the rule. In the typical case we have a latent ability to direct attention and form beliefs by direct observation *or* by accepting others' reports as proxies for direct observation.

When we are told that two dogs are male, for instance, we may not perceptually encounter the dogs but we understand what sorts of perceptual disclosures could serve as motivation for someone believing that idea. We therefore assume that such belief was initially warranted by observation and subsequently got passed through a chain of language-acts whose warrants are rooted in the perceived credibility of the speaker. Internal to this process is our prior knowledge of the parameters for judging statements like *this dog is male* observationally.

True, sometimes such observational warrants are less on display. If I had never heard of Staffies (Staffordshire pit bulls), I would be fuzzier about observational warrants and could end up in a conversation like:

- (117) Those dogs are Staffordshires.
- (118) What's a Staffordshire?
- (119) It's a breed of dog.

© 2021/05/24 Nathaniel Christen

Here I still don't really have a picture of what it is like to tell observationally that a dog is a Staffordshire. There may not be any visual cues — at least none I know of — which announce to the world that some dog is a Staffy (compared to those announcing that it is male, say). But insofar as I am acquainted with the concept *dog breed*, I also understand the general pattern of these observations. For instance, I may know breeds like poodles or huskies and be able to identify *these* by distinctive visual cues. I also understand that dogs' parentage is often documented, allowing informed parties to know their breeds via those of their forebearers. That is, I am familiar with how beliefs about breeds are formed based on observation rather than just accepting others' reports, so I know the extralinguistic epistemology anchoring chains of linguistic reports in this area to originating observations — even if I cannot in this case initiate such a chain myself.

My overall point is that language enables us to formulate beliefs based on the beliefs of others, but this is possible because we also realize what it is like to formulate *our own* beliefs, and envision that sort of practice at the origin of reports that later get circulated via language. If we can't sufficiently picture the originating observations, we don't feel like we are grasping the linguistic simulacrum of those reports with enough substance. If I never learn what Stafforshire is, an assertion that some dogs are Staffordshires has no real meaning for me — even if I trust the asserter and do indeed thereby believe that the dogs are Staffordshires. Notice that merely knowing Staffordshire is a breed of dog does not expand my conceptual repertoire very much — it does not tell me how to recognize a Staffordshire or what I can do with the knowledge that a dog is one (it cannot, for instance, help me anticipate his behavior). Nevertheless even (only) knowing that Staffordshire is a breed of dog seems to fundamentally change the status of sentences like *those dogs are Staffies* for me: I do not *have* the conceptual machinery to exploit that knowledge, but I understand what *sort* of machinery is involved.

In short, the *linguistic* meaning of concepts is tightly bound to how concepts factor in perceptual observations anterior to linguistic articulation. As a result, during any episode wherein conversants use language to compel others' beliefs, an intrinsic dimension of the unfolding conversation is that people will form their own (extralinguistic) beliefs — and can also imagine themselves in the role of originating the reports they hear via language, whether or not they can actually test out the reports by their own observations.

This extralinguistic epistemic capacity is clearly exploited by the form of language itself. If a tasting organizer hands me a glass and says "This is Syrah", she clearly expects me to infer that I should take the glass from her and taste the wine (and know that the glass contains wine, etc.). These conventions may be *mediated* by language — we are more likely to understand "unspoken" norms by asking questions, until we gain enough literacy in the relevant practical domain to understand unspoken cues and assumptions. But many situational assumptions are extralinguistic because they are (by convention) not explicitly stated, even if they accompany content that *is* explicitly stated. *This is Syrah* accompanied by the gesture of handing me a glass is an indirect invitation for me to drink it (compare to *Please hold this for a second?* or *Please hand this to the man behind you?*).

I bring to every linguistic situation a capacity to make extralinguistc observations, and to understand every utterance in the context of hypothetical extralinguistc observations from which it originates. My conversation peers can use language to trigger these extralinguistic observations. Sometimes the "gap" — the conceptual slot which extralinguistic reasoning is expected to fill — is directly expressed, as in *See the dog over there?*. But elsewhere the "extralinguistic implicature" is more indirect, as in *This is Syrah* and my expected belief that I should take and taste from the glass. But in any case the phenomenon of triggering these extralinguistic observations is one form of linguistic "side effect", initiating a change in my overall conceptualization of a situation by compelling me to augment beliefs with new observations.

All told, then, the language which is presented to me has the effect of initiating changes in what I believe — partly via signifying propositional content that I could take on faith, but partly also via directing my attention and my interpretive dispositions to guide me towards extralinguistic observations. Here I will argue that side-effects like these are not side-effects *of* linguistic meaning, but are in some sense *constitutive* of meaning.

© 2021/05/24 Nathaniel Christen

3.3 Side-Effects and Logical Incompleteness

In the most common analysis, I would argue, side-effects and propositional content overlap. The effect of an utterance, all else being equal, is that addressees understand the claimed report as believed by the speaker and — depending on their deeming the speaker credible — accept it as a provisional doxic given themselves. There are many other effects, such as how one speaker's assertion changes dialogic context in a way that alters the current topical focus and the interpretation of context-specific expressions. But these generally build off of a statements' coherence and credibility — a non-sensical comment is less likely to veer a discourse in new directions or reinscribe the dialogic "ledger". For these reasons, we may be tempted to focus analysis on the architecture of propositional content rather than the itemization of side-effects. I would argue, however, that linguistic side-effects are more universal than logically pristine propositional signification, and that communicating doxic content is a special case of side effects more than side-effects are the passive consequence of rational conversation.

I will defend this perspective via cases where language users seem to traffic in a relative *absence* of semantic determinism, with no detrimental effects to the *telos* of language in context. These cases buttress an idea that language is not targeted at doxic specificity as a precondition for meaning in general, but rather packages doxa along with other contextualizing constituents in the service of pragmatic ends. Consider:

- (120) My colleague Ms. O'Shea would like to interview Mr. Jones, who's an old friend of mine. Can he take this call?
- (121) I'm sorry, this is his secretary. Mr. Jones is not available at the moment.

It sounds like Ms. O'Shea is trying to use personal connections to score an interview with Mr. Jones. Hence her colleague initiates a process intended to culminate in Ms. O'Shea getting on the telephone with Mr. Jones. But his secretary demurs with a familiar phrase, deliberately formulated to foment ambiguity: (121) could mean that Mr. Jones is not in the office, or that he is in a meeting, or he is unwilling to talk, or even missing (like the ex-governor consummating an affair in Argentina while his aides thought he was hiking in Virginia). Or:

- (122) Mr. Jones, were you present at a meeting where the governor promised your employer a contract in exchange for campaign contributions?
- (123) After consulting with my lawyers, I decline to answer that question on the grounds that it may incriminate me.

Here Mr. Jones neither confirms nor denies his presence at a corrupt meeting.

As these examples intimate, the processes language initiates do not always result in a meaningful logical structure. But this is not necessarily a complete breakdown of language:

- (124) Is Jones there?
- (125) He is not available.

The speaker of (125) does not provide any prima facie logical content: it neither affirms nor denies Jones's presence. Nonetheless that speaker is a cooperative conversational partner (even if they are not being very cooperative in real life): (125) responds to the implicature in (124) that what the first speaker really wants is to interview Jones. So the second speaker conducts what I called a "transform" and maps *Jones is here* to *Jones is willing to be interviewed*. Responding to this "transformed" question allows (125) to be (at least) linguistically cooperative while nonetheless avoiding a response at the *logical* level to (124). (125) obeys conversational maxims but is still rather obtuse.

So one problem for theories that read meanings in terms of logically structured content — something like, the meaning of an (assertorial) sentence is what the world would be like if the sentence were true — is that the actual logical content supplied by some constructions (like *Jones is not available*) can be pretty minimal — but these are still valid and conversationally cooperative segments of discourse. To be sure, this content does not appear to be *completely* empty: "Jones is not available" means the conjunction of several possibilities (he cannot be found or does not want to talk or etc.). So (125) does seem to evoke some disjunctive predicate. But such does not mean that this disjunctive predicate is the *meaning* of (125). It does not seem as if (125) when uttered by a bodyguard is intended first and foremost to convey the disjunctive predicate. Instead, the bodyguard is

© 2021/05/24 Nathaniel Christen

responding to the implicature in the original *Is Jones there?* query — the speaker presumably does not merely want to know Jones's location, but to see Jones. Here people are acting out social roles, and just happen to be using linguistic expressions to negotiate what they are able and allowed to do.

Performing social roles — including through language — often involves incomplete information: possibly the secretary or bodyguard themselves do not know where Jones is or why he's not available. We could argue that there is *enough* information to still ground *some* propositional content. But this is merely saying that we can extract some propositional content from what speakers are supposed to say as social acts, which seems to make the content (in these kinds of cases) logically derivative on the enactive/performative meaning of the speech-acts, whereas a truth-theoretic paradigm would need the derivational dependence to run the other way. By saying *Jones is unavailable* the speaker is informing us that our own prior speech act (asking to see or talk to him) cannot have our desired effect — the process we initiated cannot be completed, and we are being informed of that. The person saying *Jones is unavailable* is likewise initiating a *new* process, one that counters our process and, if we are polite and cooperative, will have its own effect — the effect being that we do not insist on seeing Jones. The goal of "Jones is unavailable" is to create that effect, nudging our behavior in that direction. Any *logic* here seems derivative on the practical initiatives.

And moreover this practicality is explicitly marked by how the chosen verbiage is deliberately vague. The declaration "Jones is unavailable" does not *need* logical precision to achieve its effect. It needs *some* logical content, but it exploits a kind of disconnect between logical and practical/enactive structure, a disconnect which allows "Jones is unavailable" to be at once logically ambiguous and practically clear — in the implication that we should not try to see Jones. I think this example has some structural analogs to the grandma's window case: *there* we play at logical substitutions to respond practically to grandma's request in spirit rather than *de dicto*. *Here* a secretary or bodyguard can engage in logical substitution to formulate a linguistic performance designed to be conversationally decisive while conveying as little information as possible. The logical substitution in grandma's context *added* logical content by trying alternatives for the window being closed; here, the context allows a *diminution* in logical content. We can strip away logical detail from our speech without diminishing the potency of that speech to achieve effects. And while the remaining residue of logical content suggests that some basic logicality is still essential to meaning, the fact that logical content can be freely subtracted without altering practical effects suggests that logic's relation to meaning is something other than fully determinate: effect is partially autonomous from logic, so a theory of effect would seem to be partially autonomous from a theory of logic. I can be logically vague without being conversationally vague. This evidently means that conversational clarity is not identical to logical clarity.

Let us agree that — beneath surface-level co-framing complexity — many language acts have a transparent content as "doxa" that gets conveyed between people with sufficiently resonant cognitive frames. So *in the overall course of communication* we have propositional contents that converge among discourse partners, suspended between the various cognitive and pragmatic units which contextualize a given, unfolding dialog. There is in short a *holistic* mapping between units of discourse and "units" of propositionality, or "doxa". This general observation leaves unstated, however, *how* language elements map to corresponding doxic particulars. I have argued that focusing on the *logical structures* of propositions can lead us astray if we seek to find concordant formations on the language side.

Consider our attempts to close grandma's kitchen window. My analysis related to conceptual "transforms" assumed that we can find, substituting for *literal* propositional content, some *other* (representation of a) proposition that fulfills a speaker's unstated "real" meaning. Sometimes this makes sense: the proposition "that the *bedroom* window is closed" can neatly, if the facts warrant, play the role of the proposition that *the kitchen window is closed*. But we can run the example differently: there may be *no* window open, but instead a draft caused by non-airtight windows (grandma might ask us to put towels by the cracks). Maybe there is no draft at all (if grandma is cold, we can fetch her a sweater). Instead of a single transform, we need a a system of potential transforms that can adapt to the facts as we discover them. Pragmatically, the underlying problem is that *grandma is cold*. We can address this — if we want to faithfully respond to her request, playing the role of cooperative conversation partners (and grandkids) — via

© 2021/05/24 Nathaniel Christen

a matrix of logical possibilities:

- (126) If the kitchen window is closed, we can see if other windows are open.
- (127) If no windows are open, we can see if there is a draft through the window-cracks.
- (128) If there is no draft, we can ask if she wants a sweater.

This is still a logical process: starting from an acknowledged proposition (grandma is cold) we entertain various other propositional possibilities, trying to rationally determine what pragmas we should enact to alter that case (viz., to instead make true the proposition that *grandma is warm*). Here we are not just testing possibilities against fact, but strategically acting to modify some facts in our environment.

But the kind of reasoning involved here is not logical reasoning per se: abstract logic does not tell us to check the bedroom window if the kitchen window is closed, or to check for gaps and cracks if all windows are closed. This all solicits practical, domain-specific knowledge (about windows, air, weather, and houses). Yet we are still deploying our practical knowledge in logical ways — there is a logical structure underpinning grandma's request and our response to it. In sum: we (the grandkids) are equipped with some practical knowledge about houses and a faculty to logically utilize this knowledge to solve the stated problem, reading beyond the *explicit* form of grandma's discourse. We use a combination of logic and background knowledge to reinterpret the discourse as needed. By making a request, grandma is not expressing one attitude to one proposition, so much as *initiating a process*. This is why it would be impolite to simply do no more if the kitchen window is closed: our conversational responsibility is to enact a process trying to redress grandma's discomfort, not to entertain the truth of any one proposition.

For all that, there is still an overarching logical structure here that language clearly marshals. We read past grandma's explicit request to infer what she is "really saying" — e.g., *that she is cold* — but we still regard her speech act in terms of its (now indirect) propositional content. However, notice how our ascertaining this content only one step toward legitimate understanding of the original speech-act (even accounting for its illocutionary dimensions). The doxa are *factors* in understanding but, given these cases, are not straightforward *designata* of linguistic compounds. This implies a critique of truth-theoretic paradigms from a semiotic and compositional perspective: language is not *composed* to convey doxa through semantic reference and grammatic form internally (without the mediation of extralinguistic cognition); propositional content does not *fall out* of syntax and semantics. I will expand on this critique in the next section.

To summarize my current arguments, then, I believe that most sentences have an accompanying propositional content, and that during conversations we interpret this content as a factor in sentence meanings, becoming aware of what our partners believe, desire, or inquire to be the case. We retain this awareness in a cumulative model of conversational context — a "doxa inventory" — alongside other referential and deictic axes establishing each dialogic setting. Essentially, I grant that this doxic layer is central to linguistic performance in general — but given this very centrality I will argue that the logical substratum of language cannot be *separated* from the totality of syntactic, semantic, and pragmatic processing such that models based on formal logic could be curated in isolation from the overarching interconnectedness of language as a cognitive system.

3.4 The Illogic of Syntax

As I understand it, a non-trivial truth-theoretic semantics requires more than a holistic association between sentences and propositional content: it requires that this association be established *by linguistic means* and *on linguistic grounds* (syntax, semantics, pragmatics). I will present several arguments against this possibility, in the general case — that is, against the possibility that for *typical* sentences we can analyze syntactic form through the lens of the logical structure of propositions signified via a sentence; or analyze natural-language semantics through a logically well-structured semantics of propositions. I will emphasize two issues: first, that the architecture of linguistic performances *does not*, in the general case, *recapitulate propositional structure*; and, second, that language-acts work through gaps in logical specificity that complicate how we should theorize the triangular relation between surface language, propositional content, and side-effect meanings.

© 2021/05/24 Nathaniel Christen

Since it is widely understood that the essence of language is compositionality, the clearest path to a truth-theoretic semantics would be via the "syntax of semantics": a theory of how language designates propositional content by emulating or iconifying propositional structure in its own structure (i.e., in grammar). This would be a theory of how linguistic connectives reciprocate logical connectives, phrase hierarchies reconstruct propositional compounds, etc. It would be the kind of theory motivated by cases like:

- (129) This wine is a young Syrah.
- (130) My cousin adopted one of my neighbor's dog's puppies.

where morphosyntactic form — possessives, adjective/noun links — seems to transparently recapitulate predicate relations. Thus the wine is young *and* Syrah, and the puppy is the offspring of a dog who is the pet of someone who is the neighbor of the speaker. These are well-established logical forms: predicate conjunction, here; the chaining of predicate operators to form new operators, there. Such are embedded in language lexically as well as grammatically: the conjunction of husband and "former, of a prior time" yields ex-husband; a parent's sibling's daughter is a cousin.

The interesting question is to what extent "morphosyntax recapitulates predicate structure" holds in general cases. This can be considered by examining the logical structure of reported assertions and then the structures via which they are expressed in language. I'll carry out this exercise vis-à-vis several sentences, such as these (supplementing my earlier, more preliminary discussion of (51)-(54)):

- (131) The majority of students polled were opposed to tuition increases.
- (132) Most of the students expressed disappointment about tuition increases.
- (133) Many students have protested the tuition increases.

There are several logically significant elements here that seem correspondingly expressed in linguistic elements — that is, to have some model in both prelinguistic predicate structure and in, in consort, semantic or syntactic principles. All three of (131)-(133) have similar but not identical meanings, and the differences are manifest both propositionally and linguistically (aside from the specific superficial fact that they are not the same sentence). I will review the propositional differences first, then the linguistic ones.

One obvious predicative contrast is that (131) and (132) ascribes a certain *quality* to students (e.g., disappointment), whereas (132) and (133) indicate *events*. As such the different forms capture the contrast between "bearing quality Q" and "doing or having done action A": the former a predication and the latter an event-report. In the case of (132), both forms are available because we can infer from *expressing* disappointment to *having* disappointment. There may be logics that would map one to the other, but let's assume we can analyze language with a logic expressive enough to distinguish events from quality-instantiations.

Other logical forms evident here involve how the subject noun-phrases are constructed. "A majority" and "many" imply a multiplicity which is within some second multiplicity, and numerically significant there. The sentences differ in terms of how the multiplicities are circumscribed. In the case of *students polled*, an extra determinant is provided, to construct the set of students forming the predicate base: we are not talking about students in general or (necessarily) students at one school, but specifically students who participated in a poll.

Interrelated with these effects are how the *tuition increases* are figured. Using the explicit definite article suggests that there is *some specific* tuition hike policy raising students' ire. This would also favor a reading where "students" refers collectively to those at a particular school, who would be directly affected by the hikes. The *absence* of an article on "tuition increases" in (132) leaves open an interpretation that the students are not opining on some specific policy, but on the idea of hikes in general.

Such full details are not explicitly laid out in the sentences, but it is entirely possible that they are clear in context. Let's take as given that, in at least some cases where they would occur, the sentences have a basically pristine logical structure given the proper contextual framing — context-dependency, in and of itself, does not weaken our sense of language's logicality. In particular, the kind of structures constituting the sentences' precise content — the details that seem context-dependent — have bona fide logical interpretations. For example, we can consider whether students are responding to *specific*

© 2021/05/24 Nathaniel Christen

tuition hikes or to hikes in general. We can consider whether the objectionable hikes have already happened or instead are proposed for the future. Context presumably identifies whether "students" are drawn from one school, one governmental jurisdiction, or some other aggregating criteria (like, all those who took a poll). Context can also determine whether aggregation is more set- or type-based, more extensional or intensional. In (131)-(133) the implication is that we should read "students" more as a set or collection, but variants like *students hate tuition hikes* operates more at the level of students as a *type*. In "students polled" there is a familiar pattern of referencing a set by marrying a type (students in general) with a descriptive designation (e.g., those taking a specific poll). The wording of (131) does not mandate that *only* students took the poll; it does however employ a type as a kind of operator on a set: of those who took the poll, focus on students in particular.

These are all essentially logical structures and can be used to model the propositional content carried by the sentences — their "doxa". We have operators and distinctions like past/future, set/type, single/multiple, subset/superset, and abstract/concrete comparisons like tuition hikes *qua* idea vs. *fait accompli*. A logical system could certainly model these distinctions and accordingly capture the semantic differences between (131)-(133). So such details are all still consistent with a truth-theoretic paradigm, although we have to consider how linguistic form actually conveys the propositional forms carved out via these distinctions.

Ok, then, to the linguistic side. My first observation is that some logically salient structures have fairly clear analogs in the linguistic structure. For instance, the logical operator for deriving a set from criteria of "student" merged with "taking a poll" is brought forth by the verb-as-adjective formulation *students polled*. Subset/superset arrangements are latent as lexical norms in senses like *many* and *majority*. Concrete/abstract and past/future distinctions are alluded to by the presence or absence of a definite article. So "*the* tuition increases" connotes that the hikes have already occurred, or at least have been approved or proposed, in the past relative to the "enunciatory present" (as well as that they are a concrete policy, not just the idea), whereas articleless "tuition increases" can be read as referring to future hikes and the idea of hikes in general: past and concrete tends to contrast with future and abstract.

A wider range of logical structures can be considered by subtly varying the discourse, like:

- (134) Most students oppose the tuition increase.
- (135) Most students oppose a tuition increase.

These show the possibility of *increase* being singular (which would tend to imply it refers to a concrete policy, some *specific* increase), although in (135) the *in*definite article *may* connote a discussion about hikes in general.

But maybe not; cases like these are perfectly plausible:

- (136) Today the state university system announced plans to raise tuition by at least 10%. Most students oppose a tuition increase.
- (137) Colleges all over the country, facing rising costs, have had to raise tuition, but most students oppose a tuition increase.

In (136) the definite article could also be used, but saying " *a* tuition increase" seems to reinforce the idea that while plans were announced, the details are not finalized. And in (137) the plural "increases" could be used, but the indefinite singular connotes the status of tuition hikes as a general phenomenon apart from individual examples — even though the sentence also makes reference to concrete examples. In other words, these morphosyntactic cues are like levers that can fine-tune the logical designation more to abstract or concrete, past or future, as the situation warrants. Again, context should clarify the details. But morphosyntactic forms — e.g., presence or absence of articles (definite or indefinite), and singular/plural — are vehicles for language, through its own forms and rules, to denote propositional-content structures like abstract/concrete and past/future.

So these are my "concession" examples: cases where language structures *do*, in their compound architectonics, signifying propositional contents — and moreover the lexical and morphosyntactic cues (like singular/plural or the choice of articles) drive this language-to-logic mapping in an apparently rule-bound and replicable fashion. These are potential case-studies of how a truth theory of language, without neglecting contextual and semantic subtleties, *could* work: capturing granular semantic constructs via sufficiently nuanced logics, and theorizing word-senses and morphology through the aegis of a structural reduction between surface language and predicate structure. My tactic for critiquing truth-theoretic

© 2021/05/24 Nathaniel Christen

paradigms is to argue that many sentences *fail* to display a mapping between lexico/morphosyntactic details and predicate structure *in this relatively mechanical fashion*. By pointing out examples where morphosytax *does* rather seamlessly recapitulate propositional content (e.g. *the tuition hike* plural/definite), we can appreciate the more circuitous hermeneutics for examples I will present wherein the morphosytax-to-logic translation, while present, is not *sui generis*.

Varying the current examples yields cases — in a sense, intermediate between logical transparency and oblique constructions I will discuss next — where logical implications are more elliptical. For instance, describing students as *disappointed* (132) implies that the disliked hikes have already occurred, whereas phraseology such as "students are gearing for a fight" would imply, conversely, that they are still only planned or proposed. The mapping from propositional-content structure to surface language here is less mechanical than, for instance, merely using the definite article on *the tuition increases*. Arguably "disappointment" — rather than just, say, "opposition" — implies a specific timeline and concreteness, an effect analogous to the definite article. The semantic register of "disappointment" bearing this implication is a more speculative path of conceptual resonances, compared to the brute morphosyntactic "the". There is subtle conceptual calculation behind the scenes in the former case. Nonetheless, it does seem as if via this subtlety linguistic resources are expressing the constituent units of logical forms, like past/future and abstract/concrete.

So, I am arguing (and conceding) that there are units of logical structure that are conveyed by units of linguistic structure, and this is partly how language-expressions can indicate propositional content. The next question is to explore this correspondence compositionally — is there a kind of aggregative, hierarchical order in terms of how "logical modeling elements" fit together, on one side, and linguistic elements fit together, on the other? There is evidence of compositional concordance to a degree, examples of which I have cited. In *students polled*, the compositional structure of the phrase mimics the logical construct — deriving a set (as a predicate base) from a type crossed with some other predicate. Another example is the phraseology *a/the majority of*, which directly nominates a subset/superset relation and so reciprocates a logical quantification (together with a summary of relative size; the same logical structure, but with different ordinal implications, is seen in cases like *a minority of* or *only a few of*). Here there is a relatively mechanical translation between propositional structuring elements and linguistic structuring elements.

However, varying the examples — for instance, varying how the subject noun-phrases are conceptualized — points to how the synchrony between propositional and linguistic composition can break down:

- (138) Student after student came out against the tuition hikes.
- (139) A substantial number of students have come out against the tuition hikes.
- (140) The number of students protesting the tuition hikes may soon reach a critical mass.
- (141) Protests against the tuition hikes may have reached a tipping point.

Each of these sentences says something about a large number of students opposing the hikes. But in each case they bring new conceptual details to the fore, and I will also argue that they do so in a way that deviates from how propositional structures are composed.

First, consider *student after student* as a way of designating *many students* (I analyzed the very similar case (59) in terms of "space building"). There is a little more rhetorical flourish here than in, say, *a majority of students*, but this is not just a matter of eloquence (as if the difference were stylistic, not semantic). "Student after student" creates a certain rhetorical effect, suggesting via how it invokes its multiplicity a certain recurring or unfolding phenomenon. One imagines the speaker, time and again, hearing or encountering an angry student. To be sure, there are different kinds of contexts that are consistent with (138): the events could unfold over the course of a single hearing or an entire semester. Context would foreclose some interpretations — but it would do so in any case, even with simpler designations like *majority of students*. What we *can* say is that the speaker's chosen phraseology cognitively highlight a dimension in the events that carries a certain subjective content, invoking their temporality and repetition. The phrasing carries an effect of cognitive "zooming in", each distinct event figured as if temporally inside it; the sense of being tangibly present in the midst of the event is stronger here than in less temporalized language, like

© 2021/05/24 Nathaniel Christen

"many students". And then at the same time the temporalized event is situated in the context of many such events, collectively suggesting a recurring presence. The phraseology zooms in and back out again, in the virtual "lens" our our cognitively figuring the discourse presented to us — all in just three or four words. Even if "student after student" is said just for rhetorical effect — which is contextually possible — *how* it stages this effect still introduces a subjective coloring to the report.

Another factor in (138) and (139) is the various possible meanings of "come out against". This could be read as merely expressing a negative opinion, or as a more public and visible posturing. In fact, a similar dual meaning holds also for "protesting". Context, again, would dictate whether *protesting* means actual activism or merely voicing displeasure. Nonetheless, the choice of words can shade how we frame situations. To *come out against* connotes expressing disapproval in a public, performative forum, inviting the contrast of inside/outside (the famous example being *come out of the closet* to mean publicly identifying as LGBTQ). Students may not literally be standing outside with a microphone, but — even if the actual situation is just students complaining rather passively — using *come out against* paints the situation in an extra rhetorical hue. The students are expressing the *kind* of anger that can goad someone to make their sentiments known theatrically and confrontationally. Similarly, using "protest" in lieu of, say, "criticize" — whether or not students are actually marching on the quad — impugns to the students a level of anger commensurate with politicized confrontation.

All these sentences are of course *also* compatible with literal rioting in the streets; but for sake of argument let's imagine (138)-(141) spoken in contexts where the protesting is more like a few comments to a school newspaper and hallway small-talk. The speakers have still chosen to use words whose span of meanings includes the more theatrical readings: "come out against" and "protest" overlap with "complain about" or "oppose", but they imply greater agency, greater intensity. These lexical choices establish subtle conceptual variations; for instance, to *protest* connotes a greater shade of anger than to *oppose*.

Such conceptual shading is not itself unlogical; one can use more facilely propositional terms to evoke similar shading, like "very angry" or "extremely angry". However, consider *how* language like (138)-(141) conveys the relevant facts of the mater: there is an observational, in-the-midst-of-things staging at work in these latter sentences that I find missing in the earlier examples. "The majority of" sounds statistical, or clinical; it suggests journalistic reportage, the speaker making an atmospheric effort to sound like someone reporting facts as established knowledge rather than observing them close-at-hand. By contrast, I find (138)-(141) to be more "novelistic" than "journalistic". The speaker in these cases is reporting the facts by, in effect, *narrating* them. She is building linguistic constructions that describe propositional content through narrative structure — or, at least, cognitive structures that exemplify and come to the fore in narrative understanding. Saying "a substantial number of students", for example, rather than just (e.g.) *many* students, employs semantics redolant of "force-dynamics": the weight of student anger is described as if a "substance", something with the potency and efficacy of matter.

This theme is also explicit in "critical mass", and even *tipping point* has material connotations. We can imagine different versions of what lies one the other side of the tipping point — protests go from complaining to activism? The school forced to reverse course? Or, contrariwise, the school "cracking down" on the students (another partly imagistic, partly force-dynamic metaphor)? Whatever the case, language such as "critical mass" or "tipping point" is language that carries a structure of story-telling; it tries to tie facts together with a narrative coherence. The students' protests grew more and more strident until ... the protests turned aggressive; or the school dropped its plans; or they won public sympathy; or attracted media attention, etc. Whatever the situation's details, describing the facts in force-dynamic, storylike, spatialized language (e.g. "come *out* against") represents an implicit attempt to report observations or beliefs with the extra fabric and completeness of narrative. It ascribes causal order to how the situation changes (a critical mass of anger could *cause* the school to change its mind). It brings a photographic or cinematic immersion to accounts of events and descriptions: *student after student* and *come out* invite us to grasp the asserted facts by *imagining* situations.

The denouement of my argument is now that these narrative, cinematic, photographic structures of linguistic reportage — signaled by spatialized, storylike, force-

© 2021/05/24 Nathaniel Christen

dynamic turns of phrase — represent a fundamentally different way of signifying propositional content, even while they *do* (with sufficient contextual grounding) carry propositional content through the folds of the narrative. I don't dispute that hearers understand logical forms via (138)-(141) similar to those more "journaslistically" captured in (131)-(133). Nor do I deny that the richer rhetoric of (138)-(141) play a logical role, capturing granular shades of meaning. My point is rather that the logical picture painted by the latter sentences is drawn via (I'll say as a kind of suggestive analogy) *narrative structure*.

I argued earlier that elements of propositional structure — for example, the set/type selective operator efficacious in *students polled* — can have relatively clean morphosyntactic manifestation in structural elements in language, like the verb-to-adjective mapping on *polled* (here denoted, in English, by unusual word position rather than morphology, although the rules would be different in other languages). Given my subsequent analysis, however, I now want to claim that the map between propositional structure and linguistic structure is often much less direct. I'm not arguing that "narrative" constructions lack logical structure, or even that their rhetorical dimension lies outside of logic writ large: on the contrary, I believe that they use narrative effects to communicate granular details which have reasonable logical bases, like degrees of students' anger, or the causative interpretation implied in such phrases as *critical mass*. The rhetorical dimension does not prohibit a reading of (138)-(141) and (57)-(60) as expressing propositional content — and using rhetorical flourishes to do so.

I believe, however, that *how* they do so unzips any neat alignment between linguistic and propositional structure. Saying that students' protests "may have reached a critical mass" certainly expresses propositional content (e.g., that enough students may now be protesting to effectuate change), but it does so not by mechanically asserting its propositional idea; instead, via a kind of mental imagery which portrays its idea, in some imaginative sense, iconographically. "Critical mass" compels us to read its meaning imagistically; in the present context we are led to actually visualize students protesting *en masse*. Whatever the actual, empirical nature of their protestation, this language paints a picture that serves to the actual situation as an interpretive prototype. This is not only a conceptual image, but a visual one.

Figurative language — even if it is actually metaphorical, like "anger boiling over" — has similar effect. Alongside the analysis of metaphor as "concept blending", persuasively articulated by writers like Gilles Fauconnier and Per Aage Brandt, we should also recognize how metaphor (and other rhetorical effects) introduces into discourse language that invites visual imagery. Sometimes this works by evoking an ambient spatiality (like "come out against") and sometimes by figuring phenomena that fill or occupy space (like "students protesting" — one salience of this language is that we imagine protest as a demonstrative gesture expanding outward, as if space itself were a theater of conflict: protesters arrayed to form long lines, fists splayed upward or forward). There is a kind of visual patterning to these evocations, a kind of semiotic grammar: we can analyze which figurative senses work via connoting "ambient" space or via "filling" space, taking the terms I just used. But the details of such a semiotic are tangential to my point here, which is that the linguistic structures evoking these visual, imagistic, narrative frames are not simply reciprocating propositionl structure —- even if the narrative frames, via an "iconic" or prototype-like modeling of the actual situation, *are* effective vehicles for *communicating* propositional structure.

What breaks down here is not propositionality but *compositionality*: the idea that language signifies propositional content *but also* does so compositionally, where we can break down larger-scale linguistic elements to smaller parts *and* see logical structures mirrored in the parts' combinatory maxims. In the later examples, I have argued that the language signifies propositional content by creating narrative mock-ups. The point of these imagistic frames is not to recapitulate logical structure, but to have a kind of theatrical coherence — to evoke visual and narrative order, an evolving storyline — from which we then understand propositional claims by interpreting the imagined scene. Any propositional signifying in these kinds of cases works through an intermediate stage of narrative visualization, whose structure is holistic more than logically compositional. It relies on our faculties for imaginative reconstruction, which are hereby drafted into our language-processing franchise.

This kind of language, in short, leverages its ability to trigger narrative/visual framing as a cognitive exercise, intermediary to the eventual extraction of propositional content. As such it depends on a cognitive layer of

© 2021/05/24 Nathaniel Christen

narrative/visual understanding — which, I claim, belongs to a different cognitive register than building logical models of propositional content directly.

3.5 *Procedures and Integration*

So far I have criticized paradigms which try to account for linguistic meaning via concordance between linguistic and propositional structure. The logical dimension to language is real, but it is not neatly ordered in the syntax/semantics interface. My slogan is that *syntactic composition does not recapitulate logical form.*

I believe that a lot of linguistics and philosophies of language obscure this point by virtue of commitments to two reasonable theses: *one*, that logical form is an important (and often the central) dimension of meaning, and *two* that language is compositional. But while language expresses propositional content — i.e., in many cases language "means" the way that propositions "mean" — language is not *compositional*, in general, the way that propositions are compositional. So compositionality and logicality have to be disentangled, and I contend linguistics and philosophies of language have failed to do so. Moreover, as I hinted above, I think this can be explained by the historical roots of philosophy of language — and by extension linguistics — in Analytic Philosophy, which in turn has a specific history and collaboration with formal logic.

This critique has two dimensions: first, although a predicate structure, a predicative specificity, does indeed permeate states of affairs insofar as we engage them rationally, such logical order is not modeled by language itself so much as by cognitive pictures we develop via interpretive processes — processes *triggered* by language details but, I believe, to some not insubstantial degree pre- or extra-linguistic. Moreover, second, insofar as we *can* develop formal models of language, these are not going to be models of predicate structure in any conventional sense.

Cognitive-interpretive processes may have formal structure — structure which may even show a lot of overlap with propositional forms — but these are not *linguistic* structures. Insofar as language triggers but does not constitute interpretive "scripts", the scripts themselves (i.e., conceptual prototypes and perceptual schemas we keep intellectually at hand, to interpret and act constructively in all situations, linguistic and otherwise) are not linguistic as such — and neither is any propositional order they may simulate. Language *does*, however, structure the *integration* of *multiple* interpretive scripts, so the structure of this integration *is* linguistic structure *per se* — and formally modeling such integration can be an interesting tactic for formally modeling linguistic phenomena. However, we should not assume that such a formal model will resemble or be reducible to formal logic in any useful way — formalization does not automatically entail some kind of *de facto* isomorphism to a system of logic (if not first-order then second-order, modal, etc.).

Instead, I want to focus in on branches of computer science and mathematics (such as process algebra, which I have already referenced) as part of our scientific background insofar as the *structural integration* of diverse "processes" (computational processes in a formal sense, but perhaps analogously cognitive processes in a linguistic sense) can be technically represented.

In "truth-theoretic" semantics, artifacts of language are intuitively pictured in terms of propositional structure. The guiding intuition compares word meanings to logical predicates; e.g. *John is married* is typically said when the speaker believes that, in fact, the predicate *married* applies to the speaker *John*. Switching to a more "procedural" perspective involves intuiting word-meaning more in terms of interpretive procedures than logical predicates. The change in perspective may not yield substantially different analyses in all cases, but I believe it does affect many cases.

Even a simple predicate like "married" reveals a spectrum of not-entirely-logical cases in ordinary language:

- (142) John is married to a woman from his home country, but he had to get the marriage legally recognized here.
- (143) John married his boyfriend in Canada, but they live in a state that does not recognize same-sex marriages.
- (144) John has been married for five years, but in many ways he's still a bachelor.
- (145) Every married man needs a bachelor pad somewhere, and wherever yours is, you need a mini-fridge.

We can make sense of these sentences because we do not conceptually define lexemes like *married* or *bachelor* via exhaustive logical rules, like *a bachelor is an unmarried*

© 2021/05/24 Nathaniel Christen

man. Instead we have some combination of prototypical images — familiar characteristics of bachelors, say, which even married men can instantiate — and a conceptual framework recognizing (and if needed distinguishing) the various legal, societal, and dispositional aspects of being married.

Intuitionwise, then, we should look beyond potential logical glosses on meanings — even when these are often accurate — and theorize semantics as the mapping of lexical (as well as morphosyntactic phenomena) to interpretive scripts and conceptual or perceptual/enactive schema — which we can collectively designate, for sake of discussion, as "cognitive procedures".

In the case of mass versus count, for example, the choice of which plural form to use — *some beers* or *some beer* — triggers an interpretation shaping how the plurality is pictured and conceived (which is itself triggered by the use of a plural to begin with). So "multiplicity" is figured in language according to two different overall frameworks or systems. Restricting attention just, say, to count plurals, there are any number of techniques for expressing sets/collections, and these are often context-sensitive. Recall the apparent logical anomaly among sentences like:

- (146) New Yorkers live in one of five boroughs.
- (147) New Yorkers reliably vote for Democratic presidential candidates.
- (148) New Yorkers constantly complain about how long it takes to commute to New York City.

where the first refers to those who live *in* the city of New York (as does presumably the second, if referring to electoral districts), while the third primarily to those who do *not* live in the city (but the designation "New Yorker" works in both cases). Moreover, (146) means literally "all" within its intended referent group, whereas the implication in (147) and (148) is that "all" is being used actually to mean "many" or "a majority of". These variations can be expressed in logical terms, and could be explicit in more "journalistic" language, but in less formal discourse may well be unstated. We can imagine sentences such as I above that hew more crisply to a logical articulation:

- (149) All New York City residents live in one of five boroughs.
- (150) The majority of New York voters support Democratic presidential candidates.
- (151) Many New York metro area residents complain about how long it takes to commute to New York City.

According to truth-theoretic semantics, sentences compel addressees to believe (or at least consider) logically structured propositions *by virtue of* linguistic shape replicating the architecture of the intended propositional complexes, as these would be represented in first-order logic. This view on linguistic meaning is consistent with the last three sentences, which are designed to map readily to logical notations (signaled by quasimathematical phrases like *the majority of*). But most sentences do not betray their logical form so readily: these latter sentences actually sound less fluent, more artificial, than their prior equivalents.

3.6 *Triggers vs. Compositionality*

In the absence of a compelling analysis of *compositionality* in the structural correspondence between narrative-framed language and logically-ordered propositional content, I consequently think we need a new theory of how the former signifies the latter. My own intuition is that language works by triggering *several different* cognitive subsystems. Some of these hew closely to predicate logic; some are more holistic and narrative/visual. Cognitive processes in the second sense may be informed and refined by language, but they have an extralinguistic and prelinguistic core: we can exercise faculties of narrative imagination without explicit use of language (however much language orders our imaginations by entrenching some concepts more than others, via lexical reinforcement).

I'm not just talking here about "imagination" in the sense of fairy tales: we use imaginative cognition to make sense of any scenario described to us from afar. When presented with linguistic reports of not-directly-observable situations, we need to build cognitive frames modeling the context as it is discussed. In the terms I suggested earlier, we build a "doxa inventory" tracking beliefs and assertions. Sometimes this means internalizing relatively transparent logical forms. But sometimes it means building a narrative/visual account, playing an imaginary version of the situation in our minds. Language could not signify in its depth and nuance without

© 2021/05/24 Nathaniel Christen

triggering this *interpretive-imaginative* faculty. Cognitively, then, language is an *intermediary* to this cognitive system. Using terms from Olin Vakarelov's "Interface Theory of Meaning" [73], [74] we might say that language is an *interface to* interpretive-imaginative cognitive capabilities.

My argument is then that extra-linguistic cognition — e.g., narrative construals, or situational understanding — can supply crucial steps in the emergence of sentences' intended propositional content; so relying on logical assessment of propositional content and correlated linguistic patterns is, at best, incomplete. I have presented numerous sentences which, in my opinion, evince patterns whose rational contributions lie outside linguistic cognition proper, but which become attached as interpretations or refinements of given linguistic content; for example the semantic model of *critical mass* or *tipping point* encapsulating a complex situational presenting.

The extra-linguistic source of many signifying contributors can be observed directly in reference to more formalized (say, Dependency-Grammatic) analyses (or indeed what they leave out). In the (56) case, the modifier *actually* implies that the asserted fact is somehow surprising and counterintuitive; it's easy enough to conclude that the speaker senses a certain lexical frisson, particularly in how *camping* is described as *lodging*. So the work of (56) is not just to point out that people camp on the beach, but also that this arrangement is the most popular "lodging" even though in its usual sense *lodging* applies more to hotels and inns. This latter pattern circuits especially through the words *actually*, *lodging*, and *camping*, none of which are word-relations identified by the Universal Dependency parse (refer back to Figure 1) but which can be seen as a second-order network superimposed on the sentence's syntactic core.

For other examples, in (55) the *colonial ambience* and *tropical climate* disjunction partly shifts the sense of *the city* from its status as something constructed and historical (the *architecture* is colonial) to a geographic location. We of course read *the climate* as *its* (the city's) climate, which establish a link between the two *the*s. This is almost anaphoric — if the phrase were *its climate* we would resolve *it* backward to *the city* — but the pattern is varied, by replacing *its climate* with *the climate*, which conceptually foregrounds the aspect of climate as an ambient phenomenon (amenable to a definite article) rather than just a possession (viz., a property of geographic places). So there are several intersecting patterns — both referential commonality and interpretive changes (see also the "Liverpool" examples for city qua architecture and qua locale) — between the two sentence-halves, which are not described in the parse-graph. And in, say,

▾ (152) We understand that the pipe fitters are also planning to picket the Lake Worth, Florida project as well. `(en_gum-ud-train)`

there is a two-layer scope implication: given *also* we read the hearer as knowing that they are picketing some *other* project. But we also hear *the pipe fitters* as referring to some specific group of pipefitters, with which the speaker has some contractual relation; we do not interpret this as a comment about pipe-fitters in general.

The details of (152), in short, bound the scope of *the pipe fitters* from both above and below. The importance of *also* for this interpretive understanding is elided from the corpora parse-graph, which just sees *also* as a common auxiliary. In general, the extra "layers" of meaning I have identified for (56), (55), and (152) require some interpretive (and seemingly extralinguistic) reasoning, so they are not explicitly traced in purely formulaic Intermediate Representations involving core linguistic aspects, such as parse-graphs vis-à-vis syntax.

The extralinguistic dimensions of language are not, of course, completely apart from language: extralinguistic reasoning is appropriate to fill in the gaps left by straightforward syntactic or logico-semantic intellection alone, but those gaps exist *because* language leads (via intra-linguistic structures) to signifying complexes. The logic and conceptual detail achieved through intra-linguistic processes set forth the parameters on extra-linguistic cognition that supplements them; so analytic coverage of the intra-linguistic aspect is still requisite for good analysis of the extra-linguistic. This formal necessity, however, should not be mistaken for a belief that logically-inflected models of syntax and semantics are complete or self-sufficient. Full implementation of a rigorous, logical syntactic-semantic system still does not get us to *language*.

This then raises the question of what kinds of formal systems *are* appropriate models for language. In a

© 2021/05/24 Nathaniel Christen

truth-theoretic semantics, the founding analogy is that meanings are propositional, so that language builds up aggregate content much as logical expressions build up complex propositions. This analogy allows the resources of mathematical logic to transfer, with some degree of rigor, over to linguistic representation — not only predicate logic but also modal logic, typed lambda calculus, and type theory in general. On the other hand, if we accept the picture that linguistic structures are often triggers to extralinguistic cognitive activity — that this is their contribution to language comprehension — then we may need to draw inspiration from a different suite of formal theories.

Probably influenced by early Analytic Philosophy, linguistics' formalizing paradigms have gravitated to mathematical logic, reflecting the influence of early 20th-century symbolic logic on a broad swath of Anglo-American science and intellectual life. Mathematics — and by extension formal logic as a theory of mathematical foundations — came to be seen as a template for rigorous analysis to which philosophy should aspire; on that hypothesis the speculative styles of phenomenology and Continental Philosophy were widely seen as methodologically flawed, compared to the systematized argumentation presented by the likes of Gottlob Frege, Bertrand Russel, Rudolf Carnap, or Willard Van Orman Quine. It is not a stretch to see the contemporary split between Cognitive and Computational linguistics as retracing, reiterating, or in a sense derived from last century's divorce between Analytic and Continental Philosophy.

On the other hand, in the digital age, a host of new formalizing projects offer an alternative to this mathematics-centric ideology. These probably include AI research to some degree, but also many other technical disciplines associated with the inner workings of computers and software — compilers and compiler design; stack machines and other low-level computing frameworks; applied type theory; database implementation; process algebras and other theories of communicating, interdependent, or concurrent processes. This technical ecosystem is arguably a better mine for linguistic inspiration precisely because all of these fields grapple with the problem of how to wrangle rationalistic behavior out of stubbornly unintelligent machines. When trying to coax a compiler to identify the correct functions for function-calls notated at each source-code location, for example, it does not directly matter whether the function-resolution is theoretically transparent — i.e., that a correct and unique resolution can be demonstrated by formal type analysis. A compiler (or other software) ecosystem does not internally manifest type theory; at best *humans* engineer software to abide by regulations which guarantee program behaviors through type-theoretic (and similar) analysis. But engineering computational systems in conformance with mathematical, or otherwise formalizable, abstractions is itself a significant technical project. The mere presence of abstractly provable structural properties is of only indirect relevance to computer implementations.

Likewise, it seems, the mere presence of system invariants which can be formalized from mental activity and given their own eidetic analysis is only indirectly relevant to cognitive science *per se*. Presumably brains have evolved under evolutionary pressures to undergird reasonable behavior; and perhaps structuralistic laws alongside physical ones are inevitable for any intelligent systems manifest in physical substance. So it is likely that logical and other formal theories have some explanatory value in identifying regulative parameters influencing the gestation and physical realization of intelligence. But unlike mathematics, where the mere demonstration of the necessary truth of a theorem guarantees its existence as mathematical fact, in the realms of both cognition and computation abstract formal truths are not manifested merely by being true. Organized systems must provide a dynamic space whose states are constrained by the relevant formal properties, in order for these formal certainties to be empirically influential.

Ultimately, computers have a limited collection of kernel functions, and software stacks are architectures that strategically call these functions in turn — through several layers of translation and interpretation — to achieve holistic effects (like displaying a textual document on a computer screen, or creating a three-dimensional animation from a set of geometric instructions). While some logical considerations structure the process by which high-level computer code is translated to low-level system functionality, the manifest ground of this system is not any logical structure but the low-level functions themselves, along with their physical side-effects (like rendering colors on-screen). This is a useful analogy for natural language: the material unfolding of language as an empirical phenomenon is not a direct concretization of any logical structure, but rather the inventory

© 2021/05/24 Nathaniel Christen

of cognitive processes which language may trigger. Extralinguistic cognition is then analogous, in the human mind, to a computer's "system kernel". While "mind as computer" analogies are usually reductionistic — software has no human subtlety or (inter)subjectivity — this architectural metaphor perhaps has some merit. At least it can point to topics in computer science that could be placed alongside predicate logic as formalizing intuitions for the study of natural language.

4 Cognitive and Computational Processes

Any attempt to bridge Computational Linguistics and Cognitive Grammar or Phenomenology must solicit one or several "founding analogies", linking phenomena on the formal/computational side with those on the cognitive/computational side. Here, I will start from the analogy of *cognitive* and *computational process*, or generically "process" (of either variety). Processes, per se, I will leave undefined, although a "computational" process can be considered roughly analogous to a single procedure implemented in a computer programming language. The story I want to tell goes something like this: understanding language involves many cognitive processes, many of which are subtly determined by each exact language artifact and the context where it is created. Properly understanding a piece of language depends on correctly weaving together the various processes involved in understanding its component parts, and the structure of the multi-process integration is suggested by the grammar of the artifact. Grammar, in a nutshell, uses relationships between words to evoke relationships between cognitive processes.

My formal elaboration of this model will be inspired at an elementary level by process *algebra* in the computational setting, as well as applied *type theory*. Inter-process relations are the core topic of Process Algebra, including sequentiality (one process followed by another) and concurrency (one process executing alongside another). In practice, detailed research around Process Algebra seems to focus especially on concurrency, perhaps because this is the more complex area of application (designing computer systems which can run multiple threads in parallel). It is likewise tempting to imagine that cognitive-linguistic processes exhibit some degree of parallelism, so that the various pieces of understanding "fall into place" together as we grasp the meaning of a sentence (henceforth using *sentence* as a representative example of a mid-size language artifact in general). Nevertheless, I will focus more on *sequential* relations between processes, suggesting a language model (even if rather idealized) where cognitive processes unfold in a temporal order.

On both the cognitive and computational side, temporality is relative rather than quantitative: the significant detail is not "before" and "after" in the sense of measuring time but rather how one process logically precedes another in effects and prerequisites. No theoretical importance is attached to *how long* it takes before processes finish, or how much time elapses between antecedent and subsequent processes (in contrast to subjects like optimization theory, where such details are often significant). We can set aside notions of a temporal continuum where subsequent processes occupy disjoint, extended time-regions; instead, one process follows another if anything affected by the first process reflects this effect at the onset of the second process. Time, in this sense, only exists as manifest in the variations of any state relevant to processes — in the computational context, in the overall state of the computer (and potentially other computers on a network) where a computation is carried out. Two times are different only insofar as the overall state at one time differs from the state at the second time. Time is *discrete* because the relevant states are discrete, and because beneath a certain scale of time delta there is no possibility of state change.

Analogously, in language, I suggest that we set aside notions of an unfolding process reflecting the temporality of expression. Of course, the fact that parts of a sentence are heard first biases understanding somewhat; and speakers often exploit temporality for rhetorical effect, elonging the pronunciation of words for emphasis, or pausing before words to signal an especially calculated word choice, for example. These data are not irrelevant, but, for core semantic and syntactic analysis, I will nonetheless treat a sentence as an integrated temporal unit, with no value attributed to temporal ordering amongst words except insofar as temporal order establishes word order and word order has grammatical significance in the relevant natural language/dialect.

While antecedent/subsequent inter-process relations

© 2021/05/24 Nathaniel Christen

© 2021/05/24 Nathaniel Christen

are among those formally recognized in Process Algebra, this specific genre of relation is implicit to other models important to computer science, such as Type Theory and Lambda Calculus. If $\mathfrak{t}$ is a type, then any computational process which produces a value of type $\mathfrak{t}$ has a corresponding ("functional") type (for sake of discussion, assume a "value" is anything that can be encoded in a finite sequence of numbers and that "types" are classifications for values that introduce distinctions between functions — e.g., the function to add two integers is different than the function to add two decimals; more rigorous definitions of primordial notions like "type" and "value" are possible but not needed for this paper). Similarly a process which takes as *input* a value of $\mathfrak{t}$ is its own type. If two processes have these two types respectively — one outputs $\mathfrak{t}$ and the other inputs $\mathfrak{t}$ — then the two can be put in sequence, where the output from the antecedent becomes the input to the subsequent. In this manner inter-process sequential relations become subsumed into "type systems" and can be studied using type-theoretic machinery rather than Process Algebras or Process Calculi as such.

The above comments apply to type theory applied in constructed environments like computer programming languages; but as I have discussed here there exists a robust type-theoretic tradition in (Natural Language) semantics, which is disjoint from but not entirely irrelevant to the type systems of formal and programming languages. Semantic types are recognized at several different levels of classification — I proposed the terms *macrotype*, *mesotype*, and *microtype*. Some of the most interesting type-theoretic effects involve medium-grained semantic criteria that are more general than lexical entries but more specific than Parts of Speech; this is the level where linguists have seemed to find the most fertile applications of sophisticated aspects to formal type theory, like dependent types and type-coercion. This perhaps reflects the situation of "mesotypes" as intermediary between and therefore interconnecting types in a mostly grammatic (Part of Speech) and a mostly semantic (lexical) sense; and therefore a topic or tool for both syntax and semantics.

Here my argumentation is informed particularly by the merged notion of *typed* processes. If we say that something (sticking with the mesotype level) has the *type* of a physical-body noun — that "Physical Body" is a type in the overall semantics of language — then I propose a corresponding type of cognitive (perceptual and conceptual) processes characteristic of perceiving and reasoning about physical things. A particular designatum — a bag of rice, say — is subsumed under the semantic type insofar as our perceptual encounters with that thing — and/or our conceptual exercises pertaining to its properties and proclivities (like being difficult to carry) — are roughly prototyped by a certain generic kind of cognitive process. This assumes that there is a similitude among processes of perceiving and thinking about physical bodies (at least the mid-sized, quotidian physical things that tend to enter nonspecialist language) sufficient to subsume them under a common prototype, which I then argue forms the cognitive substratum for the semantic type "Physical Object". Moreover, I contend a similar cognitive substratum can be found for other mid-scale semantic types that underlie analyses of semantic acceptability and metaphoricality, like "Living Thing", "Sentient Living Thing" (*flowers want sunshine* is metaphorical because it ascribes propositional attitudes to something whose type does not literally support them), and "Social Institutions" (*The newspaper you're reading fired its editor* exhibits a "type coercion" where *newspaper* is read first as an object and then as a company). One feature of semantic types is the lexical superposition of different types to produce what (in a slightly different context) Fauconnier calls a "blend": in *Liverpool, which is near the ocean, built new docks* (see 64), the city is treated as both a geographic region and a body politic.

Moreover "weighs", too, as a verb, can be given a typed-process interpretation. In its least metaphoric sense, "to weigh" connotes a practical action of measuring some object's weight by using something like a scale; as *cognitive* process the verb embodies an ability to plan, reflect upon, or contemplate this practice. So an "idea weighing 5 pounds" is anomalous because it is hard to play out in our minds a form of this practical act where the thing being weighed is mental. However, there are plenty of more figurative uses related to "weighing ideas", "heavy ideas", and so forth, so we are able to isolate the dimension of "judging" and "measuring" which is explicit in literal "weighing", and abstracting from the physical details use "weigh" to mean "measure" or "assess" in general. The phrase *weigh an idea* therefore connotes its own cognitive process — imagining someone thinking about the idea in an evaluative way — but

this figurative "script" is closed off by "5 pounds" which forces us to conceive the weighing literally with a scale, not figuratively as a kind of mental assessment. Once again, the type anomaly can be seen as a failure to map the linguistic senses evident in a sentence to an internally consistent set of cognitive procedures for dilating the semantic content seeded within each word.

Notice that I am treating cognitive processes, in themselves, as semantic more than grammatical phenomena. Literally, weighing something is a multi-step act (lifting it on the sale, reading the measurement), and even in our mental replay of hypothetical weighing-acts it seems impossible not to imagine distinct phases (just as it is impossible not to picture left and right sides of an imaginary cow). However, I assume that the cognitive script is figured by the lexeme "weighs" as a connotative unit: whatever internal structure our mental script of "weighing something" has, this structure is not a *linguistic* structure that must be encoded grammatically. Similarly, the concept *buttered toast* suggests a confluence of perceptual, physical-operational, and conceptual aspects — we are inclined to regard toast as *buttered* if it looks a certain way and also if we have seen someone apply butter to it (or have done so ourselves) and also if we are in a context where we expect to find toast that may be buttered (we are not disposed to call a piece of bread in a grocery store "buttered toast" even if it has that appearance). So the adjective *buttered* introduces multiple cross-modal parameters in addition to the underlying concept *toast*; but I feel that the lexeme aggregates these parameters into a single *linguistic* unit. In Langacker's terms, the various elements of "buttered" do not suggest *constructive effort*, as if deliberate *linguistic* processing were needed to unpack the linguistic entity to its constituent parts. Instead, "buttered" functions *semantically* as a unit (and likewise syntactically as the unit entering relations with other words — e.g. buttered-toast is an adjective/noun pair, not the noun *butter* at the root of the adjective) — even if its cognitive process is multi-faceted. Indeed, this is precisely the signifying economy of language: it captures complex cognitive procedures by iconic, repeatable lexical units.

On that theory, tieing specific word-senses to stereotyped cognitive processes is a matter of semantics, not grammar per se. Grammar, I contend, comes into play when multiple processes need to be integrated. The concept "buttered toast", for example, seems to start from a more generic concept (toast) and then add extra detail (the buttering, with all that implies conceptually, pragmatically, and sensorially). This is suggested by the substitutability of just *toast* for *buttered toast*:

- (153) I snacked on toast and coffee.
- (154) I snacked on buttered toast and iced coffee.

Because the first sentence is perfectly clear, it seems that the ideas expressed (at least in this context) by *toast* and *coffee* are reasonably complete in themselves, so the adjectives have the effect of starting with a complete idea and adding on extra detail. Procedurally, then, it seems like we have some process which takes us to "toast" and "coffee" and then, subsequent to that (logically if not temporally) we add the wrinkle of re-conceiving the toast as buttered and the coffee as iced. In short, the adjective-noun pairing is compelling us to run a pair of cognitive processes in sequence, one establishing the noun-concept as a baseline and one adding descriptive detail by an "adjectival", a specificational process.

Counter to that analysis, someone might judge that phrases like "buttered toast" and "iced coffee" are conventional enough that we don't interpret them through two meaningfully disjoint processes. This is entirely possible, given how erstwhile aggregate expressions become established units. Different phrases exhibit different levels of entrenchment:

- (155) I snacked on toast and instant coffee.
- (156) I snacked on toast and Eritrean coffee.

Arguably "instant coffee" is a de facto lexical unit, partly because reading it in terms of constituent parts is rather nonsensical (there's no non-oblique way to understand "coffee" being qualified as "instant"). Surely, however, *Eritrean coffee* is heard as a compound phrase (at least in 2021 — it is unlikely, but not impossible, that future Eritrean coffee growers will be so successful that we hear the phrase as a brand name or culinary term of art, like "Hershey's kisses" or "French toast"). The status of *iced coffee* is probably somewhere between these two. But to the degree that a language element (whether word or phrase) is entrenched and generally processed linguistically as a unit, I maintain, it tends to be governed by an integrally complete cognitive process — not necessarily one without inner structure, but where the elements of this structure piece together per-

© 2021/05/24 Nathaniel Christen

ceptually and situationally, rather than seeming to be *linguistically* disjoint conceptualizations that are brought together by the shape of phrases. Conversely, where a cognitive process has this integral character, discursive pressures nudge the language toward entrenching some descriptive phrase as a quasi-lexeme; what starts being heard as a compound designation evolves to the point where language users don't attend to constituent parts.

Obviously, this theory presupposes that there is an available distinction to be drawn between a "procedural" synthesis of disparate cognitive processes and a perceptual and/or conceptual synthesis constitutive of individual cognitive episodes. Phenomonology seems to back this up — there are some conceptual compounds that come across as more episodically fused than others. Buttered toast may evoke a temporally not-quite-instantaneous conceptualization — at the core of the concept is a practical activity that takes a few seconds to complete — but we also can imagine the buttering-act apprehended in one sole episode. On the other hand, "Eritrean coffee" ties together concepts of much more scattered provenance; the perceptual unity of *coffee* (in the sense of a specific liquid in a specific container) along with the geopolitical "background knowledge" implicit in the adjective *Eritrean*. As a cognitive synthesis *Eritrean coffee* is conceptual rather than perceptual. Provisionally we can treat this in the context of *buttered toast* being a partially-entrenched phraseology while *Eritrean coffee* is undeniably a phrasal compound, something whose constructive form must be parsed linguistically rather than figuratively.

This analysis, though, needs many caveats. After all, many bonafide *phrases* (not "quasi-lexemes") nevertheless exhibit significant phenomenological unity — i.e., they evoke integral perceptual complexes: *big dog*; *hot coffee*; *speeding car*; *red foliage*. Linguistically these seem like an underlying concept acquiring perceptual specificity via adjectival modification: "hot" was how the coffee came to my experience because I experienced it as hot (it wasn't like I experienced the coffee and then had to contemplate whether it is hot or cold). Coffee can't *not* be experienced as hot, cold, or lukewarm; it cannot be experienced without temperature (assuming I am coming into contact with it and not just looking at it). Similarly a car must be seen as at rest, moving slowly, or speeding along; foliage must be seen as having some color(s).

I have argued, however, that unless entrenched as idiomatic phrases adjective-noun pairs like *hot coffee* or *buttered toast* should be read as grammatical complexes and accordingly (in my theory) as junctures between distinct cognitive processes. On the other hand, I argued that "instant coffee" was effectively entrenched *because* there is no simplistic conceptual unity between *instant* and *coffee*, which makes it harder to hear the phrase as descriptive. Instead, the semantics of that particular adjective-noun connection are circuitous and a little hyperbolic: "instant" coffee is coffee as a substance (not a drink, in that state) from which coffee the drink can be quickly (but not instantaneously) prepared. There is a lot going on the seemingly simple "instant coffee": the shift from coffee-as-substance to coffee-as-drink; the "instant" exaggeration. In this case, the adjectival modification has *so many* moving parts that, I'm inclined to argue, it is hard to cover the whole scenario with a descriptive phrase; which in turn creates selective pressures for some pseudo-lexical unit to emerge (which turned out to be *instant coffee*) as a mnemonic for the whole conceptual multiplex. Indeed, conceptually intricate wholes tend to quickly acquire pithy conventional nominalizations simply for rhetorical convenience ("Brexit"; "Quantum Gravity"; "International Transfer Window"; "#metoo").

Notwithstanding these variations, I still find a certain logic in the relation between phenomenological unity and semantic entrenchment. Perceptually integrated wholes may correlate with linguistically aggregate constructions insofar as there is a transparent logical destructuring in the perceptual unity: in the case of substance-attribute pairs (like *hot coffee*) — deferring in the phenomenological context to Husserl's account of "dependent moments" — there is a basically unsubtle distinction between an underlying concept (like coffee) and the qualities which are its mode of appearance as well as conceptual predicates (like hot, cold, black, or light, describing sensory properties innate to the experience of a coffee-token as well as state-reports that can be propositionally attributed to it). Although the minimal sensate intention of the coffee and the predicative disposition toward ascriptions like *black* and *hot* are consciously intertwined, surely I am aware of a logicality in experience that gives the sensate and predicative dimensions different epistemic status. I don't think of my experience of the coffee's being hot as just a hot sensation qua medium of my sensorily apprehending

© 2021/05/24 Nathaniel Christen

the coffee, but rather as the sensate mechanism by which I observe the apparent fact that the coffee is hot, as a state of affairs and not just as subjective impression. We are constantly extrapolating our perceptual encounters to propositional content along these lines.

As such, I contend that such an (in some sense) innate perception-to-predication instinct grounds the procedural slicing of linguistic processes: *hot coffee* retraces in a linguistic construction the logical order of a coffee perception which on one level is a raw perceptual encounter but is simultaneously a predicative attribution. "Hot coffee" denotes a substance that can be experienced in the mode of a base concept (coffee) which is given predicative qualification (the coffee *is* hot). The fact that there may be no noticed temporal gap in *experience* between the sensate perception and the epistemic posture does not preclude a certain logical antecedent-subsequent ordering: the concept *coffee* is the predicative base for my propositional attitude that what I am dealing with here is hot coffee, not hot-sensations-disclosing-coffee or coffee-I-experience-as-hot (but who knows, maybe I'm hallucinating) or any other artificial skeptifying of my actual experience, which is of raw perception pregnant with propositional content.

So I wish to justify claims that (non-entrenched) phrase complexes like "hot coffee" are unions of disjoint cognitive processes by noting that while such phrases evoke a certain perceptual unity, they evoke a *kind* of unity which we habitually regard *conceptually* as divided into sensate givenness founding epistemic attitudes. Cognitive processes are not exclusively perceptual; they are some mixture of perceptual and conceptual (and enactive/kinaesthetic/operational). A perceptual unity can cover two conceptual aspects, like a sheet covering two mattresses. So the perceptual unity of hot coffee can become the conceptual two-step of coffee as substance and hot as attribute; committing this unity to cognition as an overarching lifelong faculty involves registering a thought-process of coffee as a substance which can, in acts of logical predication, be believed to be hot or cold, black or light, etc. The apprehension of the substance is a different cognitive process than the predication of the attribute, in terms of how these mental acts fit within our epistemic models, even if these two processes are experientially fused. Typically we see the coffee before we touch or taste it, so already the coffee has a logical status apart from the heat we predicate in it.

Likewise, even if the black color is inextricable from our perceiving (apart from odd situations where we drink the coffee without looking at it), we know the color will change if we add milk (even if just in principle because, preferring black coffee, I don't actually do so); so we know the coffee has a logical substrate apart from its color too. All of this ideation is latent in the coffee-perceptions notwithstanding whatever perceptual unities we experience that cloak logical forms such as substance/attribute under the inexorable togetherness of disclosure (the phenomenological impossibility of spatial expanse without color, say). In short, disjoint cognitive processes can be required to reconstruct a perceptually unified situation or episode, insofar as we are not just living through the episode but prototyping, logically reconstructing, signifying it — the perceptual unity in the moment does not propagate to procedural atomicity in absorbing the episode into rational exercises.

Experience, then, presents *both* perceptual unities and cognitive-propositional multiplicity; language can inherit both holism on the perceptual side and compositionality on the rational side, even in a single enactive/perceptual episode. Depending on how we via language want to figure and express experience, we can bring either unity or compositionality to the fore. Our linguistic choices will evoke perceptual unity if they select entrenched word-senses or quasi-lexical forms; they will evoke compositionality if they gravitate toward compound phrases and complex, relatively rare lexicalizations and modes of expression. To the degree that we are interested in a cognitive-phenomenological *semantics*, we can attend to the first part of this equation, to how the understood atomicity of a word sense or a conventionalized phrase often suggests an object or phenomenon consciously apprehended as an integral whole; we can trace phenomenologically the apperceptive unity that seems to drive the linguistic community's accepting lexical atoms in this sense. Conversely, to the degree that we are interested in a cognitive-phenomenological *grammar*, we can attend to how logically composite predication emerges even within perceptual unity, because our encounter with phenomena is not (save for exotic artistic or meditative pursuits) the "*dasein*" of irreflective sensory beings immersed in a world of pure experience but the deliberate action of epistemic beings carrying (modifiable but not random) propositional attitudes to perceptual encounters.

Modeling grammar as a coordination between cogni-

© 2021/05/24 Nathaniel Christen

tive processes may be an idealization, precisely because the compositive and integrative faces of consciousness are two sides of the same coin: it's not as if we work through a thought of *coffee* or *toast*, abstract and without sensory specificity, noticeably prelude to conceived/perceived attributes like *hot*, *cold*, or *buttered*. But we can still ascribe to linguistic-understanding processes an idealized, "as if" temporality, treating the elucidating of a sentence as a sequence of procedures leading from bare concepts to well-rendered logical tableau, suffused with some level of descriptive and situational particularity. So we go from *coffee* to *iced coffee* to *buttered toast and iced coffee* to *snacking on buttered toast and iced coffee*; each link in the chain stepping up toward propositional totality.

My point is not that the logical form of the sentence is composed from logically primitive and abstract parts, which is fairly trite; my point is that such logical composition is only apparent after a pattern of cognitive integration that is more subtle and exceptional. Extramentally, buttered toast is just toast with butter on it, a fairly simplistic logical conjunction. Read as a baton passed between two acts of mind, however — conceiving toast and then conceiving it buttered — the conjunction is more elaborate; the cognitive resources of *buttered* are not just "something with butter on it" but the implication of a sensory summation (the flavor, color, scent) and operational narrative (we have seen or performed the deliberate act of applying the butter). Similarly a person dressed up is not just someone whose torso is encircled by articles of clothing; a barking dog is not just an animal making random noises; a stray cat is different from a lost cat. In their interpenetration, cognitive processes develop (in the photographer's sense) narrative and causative threads that are latent in worldly situations but reduced out of logical glosses; that is why it seems incomplete, lacking nuance, or beside the point to explicate semantic meanings in logical terms, like "bachelor" as "unmarried man" (we can certainly imagine a sentence like *My best friend has been married for years but he's still a bachelor*, to imply he still has the habits and attitudes of his single days).[19]

[19] As Bradley Rives points out, connections such as *bachelor* implying *unmarried* may be better understood as logical entailments authorized (and conventionally recognized) as implied by (at least non-metaphorical) uses of the word — analogous to *kill* implying someone has *died* — rather than trying to fix *meanings* by defining "bachelor", say, as (just) the predicate of being male and unmarried, or "kill" as *cause to die* (see [63, p. 24]). I will return to this issue below in the context of "conceptual roles".

A theory of sentences building from conceptual underspecification to logical concreteness does not preclude there being different scales of specificity. *I snacked on toast and coffee* is just as acceptable as *I snacked on buttered toast and iced coffee*. The communication conveys as much situational detail as warranted in the conversational, pragmatic context. Language always has the *capability* to push further and further into specificity; how exhaustively the language user avails of this capability is a matter of choice. As theorists of language we must then analyze how language possesses the *latent* capacity to draw ever finer pictures; the adjectival *buttered* toast and *iced* coffee takes the granularity of signifying at one level (the level of the *I snacked on toast and coffee* sentence) and layers on (or really layers *within*) a yet more specific level. The architecture of how this happens is well addressed by type-theoretic methods (both coarse and mid-grained).

The remainder of this section will try to expand on this type-theoretic intuition. My central thesis is that language understanding involves integrating diverse "cognitive procedures", each associated with specific words, word morphologies (plural forms, verb tense, etc) and sometimes phrases. The form of type theory appropriate in this context is therefore one closely associated with procedural typing and resolution: that is, assigning types to procedures, and differentiating procedures based on the types of their "arguments" or "parameters" (input and output data).

4.1 Interpretive Processes and Triggers

The type-theory/procedural perspective I will mostly work from here contrasts with and adds nuance to a more "logical" or "truth-theoretic" paradigm which tends to interpret semantic phenomena via formal logic — for example, singular/plural in Natural Language as a basically straightforward translation of the individual/set distinction in formal logic. Such formal intuitions are limited in the sense that (to continue this example) the conceptual mapping from single to plural can reflect a wide range of residual details beyond just quantity and multitudes. Case in point are plurals: for each plural usage we have a conceptual transformation of an underlying singular to a collective, but how that collective is pictured varies in context. One dimension of this

© 2021/05/24 Nathaniel Christen

variation lies with mass/count: the mass-plural *coffee* (as in "some coffee") figures the plurality of coffee (as liquid, or maybe coffee grounds/beans) in spatial and/or physical/dynamic terms. So we have:

- (157) There's some coffee on your shirt.
- (158) There's coffee all over the table.
- (159) She poured coffee from an ornate beaker.
- (160) There's too much coffee in the grinder.
- (161) There's a lot of coffee left in the pot.
- (162) There's a lot of coffee left in the pot — should I pour it out?

These sentences use phrases associated with plurality (*all over*, *a lot*, *too much*) but with referents that on perceptual and operational grounds can be treated as singular — as in the appropriate pairing of *a lot* and *it* in the last sentence. With count-plurals the collective is figured more as an aggregate of discrete individuals:

- (163) There are coffees all over the far wall at the espresso bar.
- (164) She poured coffees from an ornate beaker.
- (165) There are a lot of coffees left on the table — shall I pour them out?

So, mass versus count — the choice of which plural form to use — triggers an interpretation shaping how the plurality is pictured and conceived (which is itself triggered by the use of a plural to begin with). Compare *I sampled some chocolates* (where the count-plural suggests *pieces* of chocolate) and *I sampled some coffees* (where the count-plural implies distinguishing coffees by virtue of grind, roast, and other differences in preparation) (note that both are contrasted to mass-plural forms like *I sampled some coffee* where plural agreement points toward material continuity; there is no discrete unit of coffee qua liquid). Or compare *People love rescued dogs* with *People fed the rescued dogs* — the second, but not the first, points toward an interpretation that certain *specific* people fed the dogs (and they did so *before* the dogs were rescued).

But if we restrict attention to just, say, count-plurals, there are still different schemata for intending collections. This is one dimension I left unexamined in the earlier "New York" examples (see 13 and 148):

- (166) New Yorkers live in one of five boroughs.
- (167) New Yorkers reliably vote for Democratic presidential candidates.
- (168) New Yorkers constantly complain about how long it takes to commute to New York City.

The first sentence is consistent with a reading applied to *all* New Yorkers — the five boroughs encompass the whole extent of New York City. The second sentence is only reasonable when applied exclusively to the city's registered voters — not all residents — and moreover there is no implication that the claim applies to all such voters, only a proportion north of one-half. And the final sentence, while perfectly reasonable, uses "New Yorkers" to name a population completely distinct from the first sentence — only residents from the metro area, but not the city itself, commute *to* the city.

The assumption that logical modeling can capture all the pertinent facets of Natural-Language meaning can lead us to miss the amount of situational reasoning requisite for commonplace understanding. In *People fed the rescued dogs* there is an exception to the usual pattern of how tense and adjectival modification interact: we read "people fed" in *People fed the rescued dogs* as occurring *before* the rescue; because we assume that *after* being rescued the dogs would be fed by veterinarians and other professionals (who would probably not be designated with the generic "people"), and also we assume the feeding helped the dogs survive. We also hear the verb as describing a recurring event; compare with *I fed the dog a cheeseburger*.

To be sure, there are patterns and templates governing scope/quantity/tense interactions that help us build logical models of situations described in language. Thus *I fed the dogs a cheeseburger* can be read such that there are multiple cheeseburgers — each dog gets one — notwithstanding the singular form on *a cheeseburger*: the plural *dogs* creates a scope that can elevate the singular *cheeseburger* to an implied plural; the discourse creates multiple reference frames each with one cheeseburger. Likewise the morphosyntax is quite correct in: *All the rescued dogs are taken to an experienced vet; in fact, they all came from the same veterinary college* — the singular on *vet* is properly aligned with the plural *they* because of the scope-binding (from a syntactic perspective) and space-building (from a semantic perspective) effects of the "dogs" plural. Or, in the case of *I fed the*

© 2021/05/24 Nathaniel Christen

dog a cheeseburger every day there is an implicit plural because "every day" builds multiple spaces: we can refer via the spaces collectively using a plural (*I fed the dog a cheeseburger every day — I made them at home with vegan cheese*) or refer within one space more narrowly, switching to the singular (*Except Tuesday, when it was a turkey burger*).

Layers of scope, tense, and adjectives interact in complex ways that leave room for common ambiguities: *All the rescued dogs are [/were] taken to an experienced [/specialist] vet* is consistent with a reading wherein there is exactly one vet, and she has or had treated every dog. It is *also* consistent with a reading where there are multiple vets and each dog is or was treated by one or another. Resolving such ambiguities tends to call for situational reasoning and a "feel" for situations, rather than brute-force logic. If a large dog shelter describes their operational procedures over many years, we might assume there are multiple vets they work or worked with. If instead the conversation centers on one specific rescue we would be inclined to imagine just one veterinarian. Lexical and tense variation also guides these impressions: the past-tense form (*...the rescued dogs were taken...*) nudges us toward assuming the discourse references one rescue (though it could also be a past-tense retrospective of general operations). Qualifying the vet as *specialist* rather than the vaguer *experienced* also nudges us toward a singular interpretation.

What I am calling a "nudge", however, is based on situational models and arguably flows from a conceptual stratum outside of both semantics and grammar proper; maybe it is even prelinguistic. Consider:

- (169) People fed the rescued dogs.
- (170) Vets examined the rescued dogs.

There appears to be no explicit principle either in the semantics of the lexeme *to feed*, or in the relevant tense agreements, stipulating that the feeding in (169) was prior to the rescue — or conversely that (170) describes events *after* the rescue. Instead, we interpret the discourse through a narrative framework that fills in details not provided by the language artifacts explicitly (that abandoned dogs are likely to be hungry; that veterinarians treat dogs in clinics, which dogs have to be physically brought to). For a similar case-study, consider the sentences:

- (171) Every singer performed two songs.
- (172) Everyone performed two songs.
- (173) Everyone sang along to two songs.
- (174) Everyone in the audience sang along to two songs.

The last of these examples strongly suggests that of potentially many songs in a concert, exactly two of them were popular and singable for the audience. The first sentence, contrariwise, fairly strongly implies that there were multiple pairs of songs, each pair performed by a different singer. The middle two sentences imply either the first or last reading, respectively (depending on how we interpret "everyone"). Technically, the first two sentences imply a multi-space reading and the latter two a single-space reading. But the driving force behind these implications are the pragmatics of *perform* versus *sing along*: the latter verb is bound more tightly to its subject, so we hear it less likely that *many* singers are performing *one* song pair, or conversely that every audience member *sings along* to one song pair, but each chooses a *different* song pair.

The competing interpretations for *perform* compared to *sing along*, and *feed* compared to *treat* or *examine*, are grounded in lexical differences between the verbs. I contend, though, that the contrasts are not laid out in lexical specifications for any of the words, at least so that the implied readings follow just mechanically, or on logical considerations alone. After all, in more exotic but not implausible scenarios the readings would be reversed:

- (175) The rescued dogs had been treated by vets in the past (but were subsequently abandoned by their owners).
- (176) Every singer performed (the last) two songs (for the grand finale).
- (177) Everyone in the audience sang along to two songs (they were randomly handed lyrics to different songs when they came in, and we asked them to join in when the song being performed onstage matched the lyrics they had in hand).

In short, it's not as if dictionary entries would specify that *to feed* applies to rescued dogs before they are rescued, and *to treat* applies after they are rescued. Or that *sing along* nudges scope interpretation in one direction and *perform* nudges in a different direction. These interpretations are driven by narrative construals narrowly specific to given expressions. The appraisals would be very different for other uses of the verbs in (lexically)

© 2021/05/24 Nathaniel Christen

similar (but situationally different) cases: to "treat" a wound or a sickness, to "perform" a gesture or a play. We construct an interpretive scaffolding for resolving issues like scope-binding and space-building based on fine-tuned narrative construals that can vary alot even across small word-sense variance: As we follow along with these sentences, we have to build a narrative and situational picture which matches the speaker's intent, sufficiently well.

And that requires prelinguistic background knowledge which is leveraged and activated (but not mechanically or logically constructed) by lexical, semantic, or grammatical rules and forms: *rescued dogs* all alone constructs a fairly comprehensive mental picture where we can fill in many details by default, unless something in the discourse suggests otherwise (we can assume such dogs are in need of food, medical care, shelter, etc., or they would not be described as "rescued"). Likewise *sing along* carries a rich mental picture of a performer and an audience and how the two interact, one which we understand based on having attended concerts rather than by any rule governing *along* as a modifier to "sing" — compare the effects of *along* in *walk along*, *ride along*, *play along*, *go along*. Merely by understanding how *along* modifies *walk*, say (which is basically straightforward; to "walk along" is essentially to "walk alongside") we would not automatically generalize to more idiomatic and metaphorical uses like "sing along" or "play along" (as in *I was skeptical but I played along (so as not to start an argument)*).

We have access to a robust collection of "mental scripts" which represent hypothetical scenarios and social milieus where language plays out. Language can activate various such "scripts" (and semantic as well as grammatical formations try to ensure that the "right" scripts are selected). Nonetheless, we can argue that the conceptual and cognitive substance of the scripts comes not from language per se but from our overall social and cultural lives. We are disposed to make linguistic inferences — like the timeframes implied by *fed the rescued dogs* or the scopes implied by *sang along to two songs* — because of our enculturated familiarity with such events as dog rescues (and dog rescue organizations) and concerts (plus places like concert halls). These concepts are not produced by the English language, or even by any dialect thereof (a fluent English speaker from a different cultural background would not necessarily make the same inferences — and even if we restrict attention to, say, American speakers, the commonality of disposition reflects a commonality of the relevant cultural anchors — like dog rescues, and concerts — rather than any homogenizing effects of an "American" dialect). For these reasons, I believe that trying to account for situational particulars via formal language models alone is a dead end. This does not mean that formal language models are unimportant, only that we need to picture them resting on a fairly detailed prelinguistic world-disclosure.

There are interesting parallels in this thesis to the role of phenomenological analysis, and the direct thematization of issues such as attention and intentionality: analyses which are truly "to the things themselves" should take for granted the extensive subconscious reasoning that undergirds what we consciously thematize and would be aware of, in terms of what we deliberately focus on and are conscious of believing (or not knowing), for a first-personal exposé. Phenomenological analysis should not consider itself as thematizing every small quale, every little patch of color or haptic/kinasthetic sensation which by some subconscious process feeds into the logical picture of our surroundings that props up our conscious perception. Analogously, linguistic analysis should not thematize every conceptual and inferential judgment that guides us when forming the mental, situational pictures we then consult to set the groundwork for linguistic understanding proper.

These comments apply to both conceptual "background knowledge" and to situational particulars of which we are cognizant in reference to our immediate surroundings and actions. This is the perceptual and operational surrounding that gets linguistically embodied in deictic reference and other contextual "groundings". Our situational awareness therefore has both a conceptual aspect — while attending a concert, or dining at a restaurant, say, we exercise cultural background knowledge to interpret and participate in social events — and also our phenomenological construal of our locales, our immediate spatial and physical surroundings. Phenomenological philosophers have explored in detail how these two facets of situationality interconnect (David Woodruff Smith and Ronald McIntyre in *Husserl and Intentionality: A Study of Mind, Meaning, and Language*, for instance). Cognitive Linguistics covers similar territory; the "cognitive" in Cognitive Semantics and Cognitive Grammar generally tends to thematize the

© 2021/05/24 Nathaniel Christen

conception/perception interface and how both aspects are merged in situational understanding and situationally grounded linguistic activity (certainly more than anything involving Artificial Intelligence or Computational Models of Mind as are connoted by terms like "Cognitive Computing"). Phenomenological and Cognitive Linguistic analyses of situationality and perceptual/conceptual cognition (cognition as the mental synthesis of perception and conceptualization) can certainly enhance and reinforce each other.

But in addition, both point to a cognitive and situational substratum underpinning both first-person awareness and linguistic formalization proper — in other words, they point to the thematic limits of phenomenology and Cognitive Grammar and the analytic boundary where they give way to an overarching Cognitive Science. In the case of phenomenology, there are cognitive structures that suffuse consciousness without being directly objects of attention or intention(ality), just as sensate hyletic experience is part of consciousness but not, as explicit content, something we in the general case are conscious *of*. Analogously, conceptual and situational models permeate our interpretations of linguistic forms, but are not presented explicitly *through* these forms: instead, they are solicited obliquely and particularly.

What phenomenology *should* explicate is not background situational cognition but how attention, sensate awareness, and intentionality structure our orientation *vis-à-vis* this background: how variations in focus and affective intensity play strategic roles in our engaged interactions with the world around us. Awareness is a scale, and the more conscious we are of a sense-quality, an attentional focus, or an epistemic attitude, reflects our estimation of the importance of that explicit content compared to a muted experiential background. Hence when we describe consciousness as a stream of *intentional* relations we mean not that the intended noemata (whether perceived objects or abstract thoughts) are sole objects of consciousness (even in the moment) but are that within conscious totality which we are most aware of, and our choice to direct attention here and there reflects our intelligent, proactive interacting with the life-world. Situational cognition forms the background, and phenomenology addresses the structure of intentional and attentional modulations constituting the conscious foreground.

Analogously, the proper role for linguistic analysis is to represent how multiple layers or strands of prelinguistic understanding, or "scripts", or "mental spaces", are woven together by the compositional structures of language. For instance, *The rescued dogs were treated by an experienced vet* integrates two significantly different narrative frames (and space-constructions, and so forth): the frame implied by "rescued dogs" is distinct from that implied by "treated by a veterinarian". Note that both spaces are available for follow-up conversation:

- (178) The rescued dogs were treated by an experienced vet. One needed surgery and one got a blood transfusion. We went there yesterday and both looked much better.
- (179) The rescued dogs were treated by an experienced vet. One had been struck by a car and needed surgery on his leg. We went there yesterday and saw debris from another car crash — it's a dangerous stretch of highway.

In the first sentence *there* designates the veterinary clinic, while in the second it designates the rescue site. Both of these locales are involved in the original sentence (as locations and also "spaces" with their own environments and configurations: consider these final three examples).

- (180) The rescued dogs were treated by an experienced vet. We saw a lot of other dogs getting medical attention.
- (181) The rescued dogs were treated by an experienced vet. It looked very modern, like a human hospital.
- (182) The rescued dogs were treated by an experienced vet. We looked around and realized how dangerous that road is — for humans as well as dogs.

What these double space-constructions reveal is that accurate language understanding does not only require the proper activated "scripts" accompanying words and phrases, like "rescued dogs" and "treated by a vet". It also requires the correct integration of each script, or each mental space, tying them together in accord with speaker intent. So in the current example we should read that the dogs *could* be taken to the vet *because* they were rescued, and *needed* to be taken to the vet *because* their needing rescue was associated with being injured or in poor health. Language structures guide us toward how we should tie the mental spaces, and the language segments where they are constructed, together: the phrase "*rescued* dogs" becomes the subject of the passive-voice *were treated by a vet*, causing the two narrative strands of the sentence

© 2021/05/24 Nathaniel Christen

to encounter one another, creating a hybrid space (or perhaps more accurately a patterning between two spaces with a particular temporal and causal sequencing; a hybrid narration bridging the spaces). It is of course this hybrid space, this narrative recount, which the speaker intends via the sentence. This idea is what the sentence is crafted to convey — not just that the dogs were rescued, or that they were taken to a vet, but that a causal and narrative thread links the two events.

I maintain, therefore, that the analyses which are proper to linguistics — highlighting what linguistic reasoning contributes above and beyond background knowledge and situational cognition — should focus on the *integration* of multiple mental "scripts", each triggered by different parts and properties of the linguistic artifact. The *triggers* themselves can be individual words, but also morphological details (like plurals or tense marking) and morphological agreement. On this theory, analysis has two distinct areas of concerns: identification of grammatical, lexical, and morphosyntactic features which trigger (assumedly prelinguistic) interpretive scripts; and reconstructing how these scripts interoperate (and how language structure determines such integration).

In the case of isolating triggers, a wide range of linguistic features can trigger interpretive reasoning — including base lexical choice; word-senses carry prototypical narrative and situational templates that guide interpretation of how the word is used in any given context. *Rescued*, for example, brings on board a network of likely externalities: that there are rescuers, typically understood to be benevolent and intending to protect the rescuees from harm; that the rescuees are in danger prior to the rescue but safe afterward; that they need the rescuers and could not have reached safety themselves. Anyone using the word "rescue" anticipates that their addressees will reason through some such interpretive frame, so the speaker's role is to fill in the details descriptively or deictically: who are the rescuees and why they are in danger; who are the rescuers and why they are benevolent and able to protect the rescuees. The claim that the word *rescue*, by virtue of its lexical properties, triggers an interpretive "script", is a proposal to the effect that when trying to faithfully reconstruct speaker intentions we will try to match the interpretive frame to the current situation.

The "script" triggered by word-choice is not just an interpretive frame in the abstract, but the interpretive *process* that matches the frame to the situation. This process can be exploited for metaphorical and figurative effect, broadening the semantic scope of the underlying lexeme. In the case of "rescue" we have less literal and more humorous or idiomatic examples like:

- (183) The trade rescued a star athlete from a losing team.
- (184) New mathematical models rescued her original research from obscurity.
- (185) Discovery of nearby earth-like planets rescued that star from its reputation as ordinary and boring and revealed that its solar system may actually be extraordinary.

Each of these cases subverts the conventional "rescue" script by varying some of the prototypical frame details: maybe the "danger" faced by the rescuee is actually trivial (as in the first three), or the rescuee is not a living thing whose state we'd normally qualify in terms of "danger" or "safety", or by overturning the benevolence we typically attribute to rescue events. But in these uses subverting the familiar script does not weaken the lexical merit of the word choice; instead, the interpretive act of matching the conventional "rescue" script to the matter at hand reveals details and opinions that the speaker wishes to convey. The first sentence, for instance, uses "rescue" to connote that being stuck on a losing team is an unpleasant (even if not life-threatening) circumstance. So one part of the frame (that the rescuee needs outside intervention) holds while the other (that the rescue is in danger) comes across as excessive but (by this very hyperbole) communicating speaker sentiment. By both invoking the "rescue" script and exploiting mismatches between its template case and the current context, the speaker conveys both situational facts and personal opinions quite economically. Similarly, *rescue a paper from obscurity* is an economical way of saying that research work has been rediscovered in light of new science; and "rescued from a reputation" is a clever way of describing, with rhetorical force, how opinions have changed about someone or something.

All of these interpretive effects — both conventional and unconventional usages — stem from the interpretive scripts bound to words (and triggered by word-choice) at the underlying lexical level — we can assess these by reference to lexical details alone, setting aside syntactic and morphological qualities. When morphosyntactic details *are* considered — e.g. plurals, as in (157)-(165) —

© 2021/05/24 Nathaniel Christen

we then have a spectrum of other linguistic "triggers", involving perceptual and enactive figurations (e.g. how plurality/multiplicities are conceived), alongside interpretive "scripts". My essential point however is that language needs to trigger a collection of interpretive, perceptual, and enactive/operational processes, for complete understanding; and moreover that that language structures need to signal the "routing" of information between such processes, establishing patterns of priority/sequencing and information-supplementation among them. These patterns, I believe, can be modeled to some approximation via computational type and process theories.

I contend, moreover, that these cognitive processes are not *themselves* linguistic: while they may overlap with some language-relevant concerns (like conceptualization, and doxic specificity) they are not woven from the cloth of syntactic, semantic, or pragmatic elements internal to language. It is not within the purview of linguistics then to analyze interpretive scripts (except as a subsidiary case-study), or perceptual understanding, or situationally-mediated action. What *can* be left for linguistics proper is identifying the *triggers* to these cognitive realities — insofar as content or formations in language, within our goal-directed attempt to understand others' linguistic performances, compels us toward these extra-linguistic registers.

Linguistics on this perspective is necessarily and properly incomplete: we should not look to linguistic analysis to materially or structurally explain the cognitive processes triggered by language. But we *can* analyze the triggers themselves — potentially via formal and even computational methods. So the formal models I reviewed earlier — such as the combination of Dependency graphs with typed S-Expressions encoded via double-indices — can be adopted in a basically non-formalized, cognitive-linguistic paradigm, insofar as we ascribe to the global picture of language as an "interface" or "trigger" to extralinguistic cognition.

Suppose, for instance, that morphosyntax confirms the presence of a pluralizing operation, like *New Yorkers*. Against this minimal linguistic "information", we need to invoke narrative frames, interpretive scripts, and prelinguistic background knowledge to understand what *sort* of plurality the speaker intends. The phrase *New Yorkers* can refer to those who live in New York City (166, 13); in the New York area but *outside* the city (168, 168); or some combination. It can also mean *all* (166, 13) or just *some* New Yorkers (168, 14, 167). Usually these alternatives are not explicitly spelled out, and have to be inferred from context and via background knowledge about which readings are plausible. As I argued from multiple angles earlier, in the typical case — i.e., stylistically neutral, day-to-day language — syntactic composition does not neatly recapitulate logical form.

My prior analysis demonstrated warrants for this idea by highlighting narrative and imagistic aspects of language used to convey ideas, like "come out against" providing the verb-phrases in reports of people criticizing something. The "New York" examples like (166)-(168) point to a similar conclusion, but from a more lexico-semantic orientation: words like *borough* and *commute* carry a space of logical details that tend to force logical interpretations one way or another (e.g., the detail that the territory of a city is fully partitioned by boroughs, *so*, it is *all* citizens who live in a borough). This part of the logic is however not reflected in sentence-structure; it is, rather, latent in lexical norms and assumed part of understanding relevant sentences only because linguistic competence is understood to include familiarity with the logical implications of the lexicalized concepts: e.g. that the quantification in "New Yorkers live in one of five boroughs" is *all*, but the quantification in "New Yorkers vote democratic" is *most*.

Here I'll also add the following: the current examples show how if addressees *have* the requisite background knowledge, linguistic structure does not need to replicate logical structure very closely to be understood. The content which addressees grasp may have a logical form, and language evokes this form — guides addressees toward considering specific propositional content — but this does not happen because linguistic structure in any precise way mimics, replicates, or reconstructs it or is otherwise organized propositionally. Instead, the relation of language to predicate structures is evidently oblique and indirect: language triggers interpretive processes which guide us toward propositional content, but the structure of language is shaped around fine-tuning the activation of this background cognitive dynamics more than around any need to model predicate organization architecturally. In the case of plurals, the appearance of plural forms like *New Yorkers* or *coffees* compels us to

© 2021/05/24 Nathaniel Christen

find a reasonable cognitive model for the signified multitude, and this model will have a logical form — but the linguistic structures themselves do not in general model this form for us, except to the limited degree needed to activate prelinguistic interpretive thought-processes.

I make this point in terms of plural *forms*, and earlier made similar claims in terms of lexical details (cf. *(live in a) borough* vs. *vote Democratic*). A third group of triggers I outlined involved morphosyntactic *agreement*, which establishes inter-word connections that themselves trigger interpretive processing. Continuing the topic of plurals, how words agree with other words in singular or plural forms evokes schema which guide situational interpretations. So for instance:

- (186) My favorite band gave a free concert last night. They played some new songs.
- (187) There was some pizza earlier, but it's all gone.
- (188) There were some slices of pizza earlier, but it's all gone.
- (189) There were some slices of toast earlier, but there's none left.
- (190) There was some toast earlier, but they're all gone.
- (191) That franchise had a core of talented young players, but it got eroded by trades and free agency.
- (192) That franchise had a cohort of talented players, but they drifted away due to trades and free agency.
- (193) Many star players were drafted by that franchise, but it has not won a title in decades.
- (194) Many star players were drafted by that franchise, but they failed to surround them with enough depth.
- (195) Many star players were drafted by that franchise, but they were not surrounded with enough depth.
- (196) Many star players were drafted by that franchise, but they did not have enough depth (around them).

Plurality here is introduced not only by isolated morphology (like *slices*, *players*, *songs*), but via agreements marked by word-forms in syntactically significant pairings: was/were, it/they, there is/there are. Framing all of these cases is how we can usually schematize collections both plurally and singularly: the same set can be cognized as a collection of discrete individuals one moment and as an integral whole the next. This allows language some flexibility when designating plurals (cf. *Three times, students asked an interesting question*). A sentence discussing *slices of pizza* can schematically shift to treating the pizza as a mass aggregate in *it's all gone*. Here the antecedent of *it* is *slices* (of pizza). In the opposite direction, the mass-plural "toast" can be refigured as a set of individual pieces in *they're all gone*. The single *band* becomes the group of musicians in the band. In short, how agreements are executed invites the addressee to reconstruct the speaker's conceptualization of different referents discussed by a sentence, at different parts of the sentence: linking *it* to *slices* or *cohort*, or *they* to *the band* or *the toast*, evokes a conceptual interpretation shaped in part by how morphosyntactic agreement overlaps with "semantic" agreement. Matching *they* to *the band* presents agreement in terms of how we conceive the aggregate (as a collection of musicians); using "it" would also present an agreement, but one schematizing other aspects of the band concept.

In the last five above cases, *it* similarly binds (being singular) to *the franchise* seen as a single unit — here basic grammar and conceptual schema coincide — but *it* also binds to the *core of young players*. The players on a team can be figured as a unit or a multiple. The franchise itself can be treated as a multiple (the various team executives and decision-makers), as in *they failed to surround the stars with enough depth*. The last sentence is ambiguous between both readings: "they" could designate either the players or the franchise. Which reading we hear alters the sense of "have": asserting that the star *players* lack enough depth implies that they cannot execute plays during the game as effectively as with better supporting players; asserting that the *franchise* lacks depth makes the subtly different point that there is not enough talent over all. The variant which would include "around them" nudges toward the second reading; but it is still permissible to, according to speaker intent, parse the last *they* as designating the *franchise* and the last *them* as the *players* — i.e., that final *they/them* pair having different antecedents.

The unifying theme across these cases is that when forming sentences we often have a choice of how we figure plurality, and moreover these choices can be expressed not only in individual word-forms but in patterns of agreement. Choosing to pronominalize *slices of pizza* or *cohort of players* as *it*, or alternatively *they*, draws attention to either the more singular or more multitudinal aspects of the aggregate in question. But this effect is not localized to the individual *it/they* choice; it depends on tracing the pronoun to its antecedent and construing how the antecedent referent has both individuating and

© 2021/05/24 Nathaniel Christen

multiplicity-like aspects. Thus both individuation and plurality are latent in phrases like *slices of* or *cohort of*, and this singular/plural co-conceiving is antecedently figured by how subsequent morphosyntax agrees with the singular or, alternatively, the plural.

Moreover, these patterns of agreement invoke new layers of interpretation to identify the proper conceptual scope of plurals. In (197) we hear the scope of "they" as narrower than its antecedent "the band", because only the band's *musicians* (not stage crew, managers, etc.) typically actually perform:

- (197) The band planned a tour, where they debuted new songs.
- (198) The team flew to New York and they played the Yankees.
- (199) The city's largest theater company will perform "The Flies".

Likewise in (198), only the athletes are referenced via *they played*; but presumably many other people (trainers, coaches, staff) are encompassed by "the team flew". And in (199) we do not imply that the Board of Directors will actually take the stage (the President as Zeus, say).

Even in the course of one sentence, plurals are reinterpreted and redirected:

- (200) The city's largest theater company performed "The Flies" in French, but everyone's accent sounded Quebecois.
- (201) The city's largest theater company performed "The Flies"; then they invited a professor to discuss Sartre's philosophy when the play was over.

In the first sentence, the "space" built by the sentence is wider initially but narrows to encompass only the actual actors on stage. In the second, the "space" narrows in a different direction, since we hear a programming decision like pairing a performance with a lecture as made by a theater's administrators rather than its actors. I discussed similar modulation in conceptual schemas related to plurality and pluralization earlier; what is distinct in these last examples is how the interpretive processes for cognizing plurality are shaped by agreement-patterns (like *it* or *they* to a composite antecedent) as much as by lexical choice, or morphology, in isolation.

I have accordingly outlined a theory where lexical, morphological, and morphosyntactic layers all introduce "triggers" for cognitive processes, and it is these processes which (via substantially prelinguistic perception and conceptualization) ultimately deliver linguistic meaning. What is *linguistic* about these phenomena is how specifically linguistic formations — word choice, word forms, inter-word agreements in form — trigger these (in no small measure pre- or extra-linguistic) interpretations. But as I suggested this account is only preliminary to analysis of how multiple interpretive processes are *integrated*. Linguistic *structure* contributes the arrangements through which the crossing and intersecting between interpretive "scripts" are orchestrated. Hence at the higher linguistic scales and levels of complexity, the substance of linguistic research, on this view, should gravitate toward structural integration of interpretive processes, even more than individual interpretive triggers themselves.

I have thus far emphasized types' structural role, on the premise that type theory offers technical models that can be applied to the study of "cognitive process integration". Requirements on the transfer of information between processes — how processes complement one another and take others' details as givens — are suggestive analogies to latent interactions between mental "scripts" we are poised to apply to situations; and these interactions are the essential cognitive targets which linguistic structures have to trigger by semantic/lexical, and syntactic/morphosyntactic means. The point of analyzing process preconditions type-theoretically is to leverage types as constructional media for structural theories of inter-process integration.

This approach is largely syntactic: while many lexical and semantic details may be analytically relevant, the style of analysis I just proposed is ultimately a form of grammar, identifying how organizational patterns among language-elements trigger (and are thereby explained by the intention of) corresponding patterns among cognitive processes. For sake of completeness, though, I will also (at the end of this paper) offer a few thoughts on the more "semantic" manifestation of types (and type theory) as premises of organizing concepts.

4.2 *Procedural Networks and Link Grammar*

One of my goals in this section is to incorporate *Link Grammar* into a phenomenological and Cognitive Gram-

© 2021/05/24 Nathaniel Christen

mar perspective. I will not attempt to offer a neutral exposition of Link Grammar theory, so I will accept some terminology and exposition not necessarily originating from the initial Link Grammar projects (though influenced by subsequent research, e.g. computational models developed by Matt Selway and Kenneth Holmqvist [67]; [30], [29]). I also want to wed Link Grammar to my own semantic intuitions, set forth earlier, that word-meanings and morphosyntactic interpretations should be grounded on pre- or para-linguistic cognitive "scripts" that are activated (but not structurally replicated, the way that according to truth-thoretic semantics linguistic form evokes-by-simulating propositional form) by linguistic phenomena.

Link Grammar is, depending on one's perspective, either related to or a variant of Dependency Grammar (DG), [20], [19], [21], [52], [31], [3], which in turn is contrasted with "phrase structure" grammars, such as Head-Driven Phrase Structure Grammar (HPSG) [37], [36], [66], [82] (for example). Link and Dependency Grammars define syntactic structures in terms of word-pairs; phrase structure may be implicit to inter-word relations but is not explicitly modeled by DG formalisms — there is typically no representation of "noun phrass" or "verb phrases", for example. Phrase structure is instead connoted via how relations fit together — in *rescued dogs were fed*, for instance, the adjectival *rescued-dogs* relation interacts with the *dogs-fed* (or *dogs-were* plus *were-fed*) predication, an interaction that in a phrase-structure paradigm is analyzed as the noun-phrase *rescued dogs* subsuming the noun *dogs*. Dependency analyses often seem more faithful to real-world semantics because, in practice, phrases do not *entirely* subsume their constituent parts. Linguistic structure is actually multi-layered, where semantic and morphosyntactic connections resonate between units within phrases separate and apart from how linguistic structure is organized into phrasal units themselves.

Except for phrases that coalesce into pseudo-lexemes or proper names (like "United Nations" or "Member pf Parliament"), or indeed become shortened to single words (like "waterfall"), we perceive phrases both as signifying units and as aggregate structures whose detailed combinative rationale needs contectualization and interpretation. In short, phrases are not "canned" semantic units but instead are context-sensitive performances that require interpretive understanding. This interpretive dimension is arguably better conveyed by DG-style models whose constituent units are word-relations, as opposed to phrase-structure grammars which (even if only by notational practice) give the impression that phrases conform to predetermined, conventionalized gestalts.

While Link and Dependency Grammars are both contrasted with phrase-structure grammars, Link Grammar is also distinguished from mainstream DG in terms of how inter-word relations are conceived. Standard DG recognizes an asymmetry between the elements in word-relations — one element (typically but not exclusively a word) is treated as "dependent" on another. The most common case is where one word carries greater information than a second, which in turn adds nuance or detail — say, in *rescued dogs* the second word is more essential to the sentence's meaning. This potentially raises questions about how we can actually quantify the centrality of one word or another — in many cases, for instance, the conceptual significance of an adjective is just as trenchant as the noun which it modifies. In practice, however, the salient aspect of "head" vs "dependent" asymmetry is that any inter-word pair is "directed", and one part of the relation defined as dependent on another, however this dependency is understood in a given case.

By contrast, Link Grammar does not identify a head-dependent asymmetry within inter-word relations. Instead, words (along with other lexically significant units, like certain morphemes, or punctuation/prosodic units) are seen as forming pairs based on a kind of mutual incompleteness — each word supplying some structural or signifying aspect that the other lacks. Words, then, carry with them different varieties of "incompleteness" which primes them to link up with other words. Semantic and grammatical models then revolve around tracing the *gaps* in information content, or syntactic acceptability, which "pull" words into relation with other words. This approach also integrates semantic and syntactic details — unlike frameworks such as Combinatory Categorical Grammar, which also treats certain words as "incomplete" but identifies word connections only on surface-level grammatical terms — Link Grammar invites us to explore how semantic and syntactic "completion" intersects and overlaps.

Words can be incomplete for different reasons and in different ways. Both verbs and adjectives generally need to pair up with nouns to form a complete idea.

© 2021/05/24 Nathaniel Christen

On the other hand, nouns may be incomplete as lexical spotlights on the extra-linguistic situation: the important point is not that people feed dogs in general, but that *the rescued* dogs were fed prior to their rescue. So *dogs* "needs" *rescue* for conceptual specificity as much as *rescue* needs *dogs* for anchoring — while also *dogs* needs *the* for *cognitive* specificity, because the discourse is about some particular dogs (presumed known to the addressee), signaled by the definitive article. In other cases, incompleteness is measured in terms of syntactic propriety, as in:

- (202) We learned that people fed the rescued dogs.
- (203) No-one seriously entertained the belief that he would govern in a bipartisan manner.

In both cases the word *that* is needed because a verb, insofar as it forms a complete predicate with the proper subject and objects, cannot always be inserted into an enclosing sentence. *People fed the rescued dogs* is complete as a sentence unto itself, but it is not complete as a grammatical unit when the speaker wishes to reference the signified predicate as an epistemic object, something believed, told, disputed, etc. A connector like *that* transforms the word or words it is paired with syntactically, converting across part-of-speech boundaries — e.g. converting a proposition to a noun — so that the associated words can be included into a larger aggregate.

The interesting resonance between Link Grammar and Cognitive Grammar is that this perspective allows us to analyze how syntactic incompleteness mirrors semantic incompleteness, and vice-versa. "Incompleteness" can also often be characterized as *expectation* (which was how I framed this line of theory earlier): an adjective "expects" a noun to produce a more tailored and situationally refined noun (or nominal "idea"); a verb expects a noun, to form a proposition. Analogously,when we have in discourse an adjective or a verb we expect a corresponding noun — so via syntactic norms language creates certain expectations in us and then communicates larger ideas by how these expectations are met. Is a noun-expectation fulfilled by a single noun or a complex phrase? The notion of semantic and syntactic expectations also coordinates nicely with type-theoretic semantics; for example, the verb "believe" pairs with a semantic unit that can be interpreted in epistemic terms — not any noun but a noun of a kind that can be the subject of propositional attitudes (beliefs, opinions, assertions, arguments, etc.).

Kenneth Holmqvist, whom I discussed earlier as integrating Conceptual Space Theory with Cognitive Grammar, made a study of "expectation" in a similar sense a central feature of his doctoral dissertation [30]; I'll point out that there is an implicit resonance between *expectations* and link grammar; so Holmqvist's research actually potentially triangulates between Conceptual Space Theory, Cognitive Grammar, and Link Grammar.

Continuing analysis of *that* qua subordinator: the syntactic incompleteness of propositional phrases modified by *that* can therefore be traced to the semantic expectations raised by *believe*, and analogous verbs (opine, argue, claim, testify). The object of *testify*, say, is a statement of potential fact which we know not to take as necessarily true or honestly made (part of the nature of testimony is that it may be deliberately or accidentally fallacious). But to properly pair with *testify*, then, phrases must be semantically reinterpreted as nominalizations of propositions, rather than as mere linguistic expression of propositional content via complete sentences. The "epistemic" context transforms sentential content into nominal content available for further refinement:

- (204) The Trump campaign colluded with Russia.
- (205) Several witnesses testified that the Trump campaign colluded with Russia.
- (206) Reputable newspapers have reported that the Trump campaign colluded with Russia.
- (207) Most Democrats believe that the Trump campaign colluded with Russia.

The grammatical stipulation that a modifier like *that* is often necessary in such formulations correlates with the semantic detail that the "claimed", "testified", or "believed" content is not being directly asserted by the speaker as if in an unadorned declarative expression; this points to the contrast between (204) and the subsequent three sentences.

Morphosyntactic transformation similarly models the correlation between semantic and syntactic expectation — as can be demonstrated by a variant of the "believe" forms, via the phrase "believe in":

- (208) I believed in Father Christmas.

© 2021/05/24 Nathaniel Christen

- (209) I believed in Peace on Earth.
- (210) I believed in Obama.
- (211) I believed in lies.

Whereas *that* (after *believe*) "nominalizes" propositions, *in* reconceives (type-theoretically we would say "coerces") ordinary nouns into epistemic nouns (compatible with propositional attitudes). Obama is not an *idea*, but the connector *in* triggers an interpretation where we have to read "Obama" as something believed — effectively a type-theoretic tension resolved by understanding *Obama* in this context to designate either his platform or his ability to implement it. Interpretive *tension* is a natural corollary to a mismatch in expectations: *believe* expects something epistemic, but the discourse gives us a proper name. Analogously *budge* expects a brute physical entity in its simplest meaning, but in:

- (212) Obama wouldn't budge on reproductive rights.

we get a "sentient" noun, and have to read *budge* metaphorically. In short, *expectation*, *interpretive tension*, and *incompleteness* are interlocking facets of semiotic primitives that gestate into discursive maneuvers via which ideas are communicated economically and context-sensitively.

Link Grammar, proper, represents only the most immediate (mostly grammatical) facet of word links. For sake of discussion, I will discuss links in general as markers of *mutual dependency* between words, so a "link grammar" is essentially a "bi-directional" Dependency Grammar. Mutual dependencies manifest syntactic norms and contextual details that make individual words inadequate signifying vehicles for a particular communicating content. This overall principle becomes concrete in one form via grammatical relations, which is the layer modeled by Link Grammar proper. I have mentioned several ready examples — the syntactic dependency of verbs and adjectives on nouns for them to enter correctly into discourse (correlate with a semantic dependency in the other direction, to shape noun-ideas to the proper context and signifying intent); also part-of-speech or "subtyping" dependencies reflecting mandates that (in my examples) propositional phrases are coerced to nouns or nouns coerced to "epistemic" nouns. Technical Link Grammars recognize a broad spectrum of "link relationship" — between 50 and 100 for different languages. Parsing a sentence is then a matter of identifying all of the mutual dependencies — at least those evident on a syntactic level — that appear as inter-word links in the sentence. Phrase structure may be implicit in some links in combination — for example verb-subject plus verb-object links generate propositional phrases — but the technical parse is a "graph" of inter-word links rather than a "tree" of phrases ordered hierarchically. The parse-graph itself is only a provisional reading of the sentence, and linguistic understanding exists only insofar as its skeletal outline is filled out with semantic and situational details. But the graph layer articulated by a Link Grammar still provides a useful intermediate representation, showing mutual dependencies in their syntactic manifestations that then point toward their deeper semantic and situational origins.

For each *syntactic* bi-dependency, on this theory, there is a concordant semantic and signifying bi-dependency, partly conventionalized as a feature of the language and partly hewn to the current discourse context. To leverage Link Grammar in an overall Cognitive Linguistic environment, then, we need to examine the semantic relations that drive syntactic bi-dependency: how the grammatical structure of one word completing another is a codification of *semantic* mutual dependency. The *semantic* bi-dependencies operate on both abstract and concrete levels. Abstractly, it is obvious that an adjective or a verb depends on a linked noun to complete an idea. This abstract prototype of bi-dependency then takes concrete forms in each specific discourse, acquiring detail and specificity from situational contexts.

The crucial dimension in this theory is neither abstract nor concrete bi-dependency but the intersection of the two. The conventionalized lexical, syntactic, and morphosyntactic norms of a language present abstract prototypes of mutual word dependencies. The concrete instantiation of these forms — via word-pairs whose surface presentation indicates the presence of specific link relations (often with the aid of morphology, agreement, spoken inflection, and other morphosyntactic cues) — invites us to consider how an abstract bi-dependency becomes situationally concretized in the present, momentary context. In essence, abstract mutual dependencies are the raw materials from which situational appraisals are created. A pairing like "rescued dogs" uses a certain abstract-bidependent prototype — here the double-refinement of a noun and adjective grounding

© 2021/05/24 Nathaniel Christen

each other — to trigger the listener's awareness that the speaker's discourse is centered on one specific aspect of the dogs (that they were at some point rescued) with its conceptual corollaries and unstated assumptions (that, being in need of rescue, they were abandoned, in danger, and so on). Similarly the further link to the definite article — " *the* dogs" — evokes the prototype of a definite article grounding a noun, which in turn communicates the speaker's beliefs about the current state of the discourse.

This last "bi-dependency" deserves further comments, because nouns more often than not reveal some dependency on an article: *some dog(s)*, *the dog(s)*, *a dog*, *many dogs*. These pairings paint a picture both through the choice of article and the presence or absence of a plural. This picture is partly situational — obviously whether the speaker is talking about one or multiple dogs is situationally important — but it is also meta-discursive: selecting the definite article indicates the speaker's belief that the listeners know which dogs are on topic. The lexeme "the" thereby signifies a meta-discursive stance as well as a cognitive framing — both that the collection has enough coherence to function as a conceptual unit in context, as *the* dogs (and not something less specific like "*some* dogs") *and also* that the speaker and listeners share compatible cognitive pictures as a result of the prior course of the discourse. This also introduces several avenues for future discursive evolution — the listeners can respond to the speaker on both cognitive and meta-discursive levels. A direct question like *which dogs?* signifies that the first speaker's meta-discursive presuppositions were flawed — the referent of "the dogs" has not been properly settled by the discourse to that point. Or questions for clarification — like *how many dogs are there?* — indicate the listeners' sense that all parties' respective construal of the situation needs to be more neatly aligned for the discourse to proceed optimally.

The point I particularly want to emphasize here, though, is that these discursive/cognitive effects inhere not only in the word "the" but in its pairing with other words, like *the dogs*. We tend to see the lexical substratum of a language as the ground level of its signifying potential, but we should perhaps recognize bidependent prototypes as equally originating. The communicative content of *the dogs* is ultimately traced not only to the lexical potentialities of *the* and *dogs* as word-senses, but to the abstract prototype of the definite-article bidependence, which becomes concretized in the *the dogs* pairing at the same time as the individual words do.

In order to bring this account full-circle, I would then add that lexical units mutually completed by an instantiated bidependence can also be seen as a tie-in between interpretive procedures. Lexical interpretive scripts — the cognitive processing solicited by *the* and *dogs* in isolation — are themselves open-ended and ungrounded or incompletely grounded.

We can speak metaphorically of "words" being incomplete, or carrying expectations, but it is really the cognitive scripts associated with words that are lacking in detail. The resolution of a merely schematic cognitive outline to a reasonably complete situational construal can be likened to a rough sketch filled out in color — but we have to imagine that a sketch can be completed by pairing it with a second sketch, and the content in each one, crossing over, allows a completed picture to coalesce.

So in the current running example, *the* in itself evokes an interpretive process that on its own logic cannot be completed, and likewise *dogs*; but each script supplies the content missing from the other. In this sense the bidependent form concretized by the pair is actually evoking an interpretive phenomenon of mutual completion — the language structure here is guiding us toward an interpretive interpenetration where the two scripts tie each other's open ends. Whereas lexical items can be associated with single "scripts", prototypes of mutual dependency model patterns in how script-pairs can become mutually complete.

But unlike lexemes, which are notated directly by language, the instantiation of bidependent script-pairs occurs indirectly. Some paired-up words are of course adjacent, but adjacency between words does not have the same binding determinacy as sequencing among morphemes in *one* word. Instead, word adjacency is only one signal among many others suggesting that some prototypical inter-word relation applies between two words (which might be widely separated in a sentence). Morphology and syntax also point towards the pairings operative in a sentence — they are to bidependency prototypes what word-choice is to the lexicon.

© 2021/05/24 Nathaniel Christen

4.3 Bi-Dependency and Conceptual Roles

Word-pair relations can be considered from the point of view of syntax (insofar as larger-scale syntactic aggregates can be modeled as complexes whose building-blocks are individual word pairs) and also semantics (for example, by identifying plausible type systems for language and applying these to word-pairs such that a lexeme with "function-like" type transforms a suitably typed "target" word, yielding a suitably-typed phrase or "compound idea"). The overall hypothesis I have tried to iconify with the label "Cognitive Transform Grammar" is that each word-pair sets into place some further element of grounding and specificity, so that the cumulative whole of a sentence contains (prototypically) a relatively grounded and complete propositional content (John *has arrived*, Everybody *had already left*, Some invitees *are not coming*, and so forth) whereas individual words are (prototypically) generic and underdetermined. All language is the transformation of underspecified lexemes to relatively complete expressions of belief (or related propositional attitudes) via syntactically-constrained semantic combinations. This, I believe, is the essential dynamic or "momentum" of language, and it is necessary to situate different branches of linguistic analysis (grammar, semantics, pragmatics, prosody, morphology) under the aegis of this flow toward specificity — each subfield concerns how a different dimension of linguistic structure/convention contributes particular cognitive and processing norms/faculties to the overall dynamic of linguistic specification.

Many semantic effects in the "emergence of specificity" are tied closely to lexical senses, and reflect partial entrenchment which supports neither simplistic compositionality nor unstructured atomism. That is, *red card*, *pink ribbon*, *defense lawyer*, or *defensive tackle* (for example) are compounds whose meaning is more specific/restricted than the simple concatenation of concepts of *red* and *card* (and so on), but surely *red* and *card* retain some autonomy in the phrase — we can ask *was that card red?* or *are they getting another card?* showing a dimension of the word-pair as a two-word phrase and not a de facto compound lexeme (but compare how it would be awkward to ask of a member of parliament, say, *what is she a member of?*). In short, analysis of the semantic effects of word-pairing has to take on different commitments in different contexts. Sometimes an adjective/noun pair, for instance, is interpreted fairly closely as the former being a property-ascribing modifier of the latter (*black dog*, say) while other times the pairing of two words yields a "blend" whose contours are shaped by relatively narrow word-senses (as in the *red* of *red card*).

To the degree that they are complementary sources of meaning, compositionality (the fact of *black dog* being basically a blend of the two words, say) compared to entrenchment (*red card* as a phrase with a narrow constrained/conventionalized meaning, say) both tend toward extreme cases where semantics *per se* seems secondary. Either we have combinations which appear to operate on a conceptual/perceptual level (there isn't a lot of *linguistic* reasoning involved in construing a dog as black) or phrases reified to quasi-lexical units (without knowing the rules of football there would be no way to infer the meaning of *red card*, though it is still conceptually relevant that red cards are, in fact, cards, and red). The more interesting semantic cases lie between the two extremes, where we have mixtures of convention and novelty such that addressees interpret pairings/constructions (such as adjective-noun) partly by considering how the concepts exposed by each word can be combined, and partly also by consulting idioms or popular word-uses which signal that the relevant linguistic community tends to blend the constitutive concepts *in certain ways*.

In Jens Allwood [4], a form of Conceptual-Role Semantics (but one which is, I would argue, distinct from an "inferential" role semantics that is often conflated with conceptual roles) points to another strategy for providing semantic detail to the analysis of word-pairs. Allwood points out that different entities play distinct roles in the processes, states, or relations which are profiled by verbs (and thus form the central substratum of sentence-level meaning). If, as I suggest, the essential dynamic of language is the accumulation of detail such that sentences acquire determinate meaning, then what is thereby determinate is some form of relation, state, process, or event, i.e., the sort of concept tokened by verbs. The "things" profiled by nouns can play various *roles* in relation to verbs, i.e. in relation to the events/processes/states profiled by verbs and the situations/contexts where these obtain and occur. Nouns can profile *agents*, *patients*, *locations*, *purposes*, *instruments*, or *benefactors* of an event,

© 2021/05/24 Nathaniel Christen

process, or state (among other roles). These nicely map to case systems in some languages, of course, which suggests that there is something cognitively fundamental in conceptual roles (specifically the roles of concepts within *situations*, which in general are multi-part relations whose relata have distinct roles in the situation overall) analogous to how verb-tenses point to a cognitive essentiality of different modes of temporality relative to a linguistic act (time spans as open-ended or completed and as past, present, or future).

Allwood's discussion is also interesting because it points to how claims of Conceptual Role Semantics can be related to Cognitive Grammar: effects that we may be inclined (under the influence of Langacker) to figure as effects of construal and profiling may be seen to have an additional dimension involving situational roles relative to a central verb-profiled event, process, or state. Allwood focuses for examples on notions of "war" and "peace". From a cognitive perspective, note that *war*, say, can be figured in multiple ways — e.g., as a period of time, or as a place (where the war is being fought), or as a state, or a cause or result. The following examples are mostly derived from (but modified versions of) Allwood's cases:

- (213) War could break out any time.
- (214) The war has lasted many years.
- (215) A tenuous peace has held for over a year now.
- (216) The cruel war is raging, Johnny has to fight.
- (217) War is a cause of misery.
- (218) The use of war to score political points is immoral.
- (219) Most of the refugees are fleeing civil war.
- (220) Before the war, things were different.
- (221) Historians will study this war for a long time.
- (222) Diplomats met from many countries to discuss the civil war.
- (223) The international community must end this war.
- (224) Diplomats met to discuss how to maintain the peace.
- (225) Don't underestimate our resolve — this is war!
- (226) Secretly the governments met to discuss peace.
- (227) We must first and foremost work for peace.
- (228) Environmental degradation was a cause of the war.
- (229) All this destruction was a result of the war.
- (230) There is more war to come.
- (231) A little peace would make a big difference.
- (232) There is war between moderates and radicals.
- (233) The moderates made peace with the radicals.
- (234) Some countries pulled out of the war.
- (235) The countries are now officially at war.

In each of these examples "war" (or peace) as an entity/phenomenon is construed in different ways, as far as how we mentally orient towards it via cognitive patterns of space, time, institutionality, state (war as a condition), and so forth, albeit with varying levels of metaphoric tinge (e.g. "pull out of" figures war as a kind of constrictive substance).

Of course, speakers can *construe* war (and peace) in different ways depending on what specific aspect they are thinking about; to *flee* a war, say, is to leave a location where a war is happening. This does not mean that "war" simply figures as a denotation of place, however; we usually understand *flee a war* or *escape a war* to imply that the war itself hinders normal movement of people from one place to another. Obviously the nature of war adds conceptual shading to whatever concept is designated by a particular use of the "war" lexeme, whether it is figured as a place, time, state, tool (e.g. for political gain), subject (for historians), etc. Nevertheless *war* as such can present itself in these different cognitive guises, and one necessity for a speaker is to elucidate in what sense she construes *war* for that specific occasion.

What emerges from Allwood's readings however is the further point that this construal also often implicates how war is figured in a conceptual role vis-à-vis the surrounding sentence; so *strive for peace* or *movements toward peace* figure "peace" as a destination (something playing a "locative" conceptual role) whereas *prepare for war* figures "war" as a condition which persists over time. Similarly *destruction caused by war* figures war as an agent (something acting/causative) whereas *treaty to end the war* figures war as a patient (something acted upon). Allwood implies at least thematically (albeit without explicitly appealing to Cognitive Grammar) that both cognitive construals (which we can see as speaker-relative) and conceptual roles (which are more lexeme-focused) contribute to words' meanings in their given linguistic context. Such effects can be observed in sentences such as (see [4, page 9]):

- (236) The peace was finally achieved.
- (237) ?The war was finally achieved.

© 2021/05/24 Nathaniel Christen

- (238) The peace lasted for five years.
- (239) The war went on for five years.
- (240) ?The peace went on for five years.

I think we hear (237) and (240) as dubious partly because of connotations, since war is rarely discussed as something desirable (implied by (237)) and peace rarely undesirable (connoted by (240) insofar as "went on" probably has, when not neutral, negative polarity in English; it suggests what is "going on" constitutes a disruption to an ordinary state, whereas peace, we hope, *is* the normal state). Allwood suggests, however, that these are also differences in the construal situations where the terms are more likely to be used; " 'peace' seems [more] to be focusing on the end state resulting from processes leading to peace" (page 9), i.e. peace is either an indefinitely continuing state or a momentary event (such as a cease-fire; the temporal end of war) whereas war is usually figured as an extended but bounded stretch of time (though there is also a rarer sense where "war" might be used to name the moment when a conflict *begins*). These are certainly differences in cognitive construal, but they also reflect patterns in how the concepts *war* and *peace* are likely to be conceptualized as inhabitants of various roles (agent, patient, location, instrument, etc.) in relevant situations.

In short, when attending to how profilings of events, states, and/or processes become more detailed as complete linguistic meanings emerge, we often should pay attention to how some details come into place by connecting something to a verb-profile *according to the role it plays*. Minimally a verb usually has a subject and object, but the event/process/state designated by each verb can be given further description as well (where/when it occurs, why or for whom, how or with what, and so forth) such that various sentence-clauses find their place in an overall sentence by providing one of these auxiliary details. Conceptual Role Semantics, then, can add a level of semantic specificity to the vague notion that components of a sentence contribute to meaning overall by "adding details" — by isolating particular roles which concepts play in situational construals, we can model words and sentences as encoding descriptions of these roles through linguistic constructions, so that an overall sentence aggregates multiple details (including but not limited to subject/verb/object). A system of conceptual roles can provide a theoretical basis for identifying structures and organization in *how* situational details supplemental to a subject/verb/object core are accumulated.[20]

I have argued that the essential thrust of linguistic meaning is toward propositional specificity, a process of supplying whatever details a speaker deems necessary for conveying situations in ways that match her beliefs and observations (or via word and structural choices which are conducive to whatever goals she has for the given linguistic act). Meaning, in short, is an *accretion of detail*, but this accumulative process happens in a speaker-relative fashion, since the details included (and how they are described) reflect the speaker's unique epistemic and practical stance — the speaker wants to achieve something by speaking (/writing). As such, accretion of detail tends to overlap with *grounding*, insofar as filling in details of a speaker's construal of relevant things and situations tends to result in a clarification of how the relevant situational content is grounded relative to the speaker's perspective — the sense in which a specific event and/or process/state, noticed and in some sense intellectually and perceptually individuated for the speaker, stands in relation to the speaker as an integral whole that can be an object of thought and, via language, of communication. We can then proceed from a generic theory of "accretion of detail" at smaller scales (e.g. word-pairs) to a "conceptual role" based theory which has more semantic posits (in the sense of roles being conceptual and linguistic universals in contrast to word-sense peculiarities which call for lexical rather than semantic investigations) and then to notions of "grounding" and speaker-relative epistemics and motivations, which I will consider in the next section. (A more detailed discussion of conceptual roles and this general progression toward a "theory of grounding" is sketched out in the second and third parts of this essay-set.)

[20] This way of approaching Conceptual Role Semantics places somewhat different emphasis than most accounts self-identifying with that term (see e.g. [35], [65], or [58] as interesting treatments as well as overviews of that mainstream) but I think overlaps therewith. I primarily want to emphasize conceptual roles in describing — or observing, or being *disposed to* observe — situations, i.e. to site concept-semantics against the background of situational cognition. I would argue that logical inferences, or other kinds of mental activity against which accounts of conceptual roles are often defined, take place in the context of situations, so examining situational reasoning and its influence on conceptual roles is more fundamental — and also more general, in that cognizing situations involves not only logical judgments but also, for instance, inferring causal explanations, or selecting points of focus and attention, or planning one's own actions within a situation.

© 2021/05/24 Nathaniel Christen

Focusing on conceptual *roles* provides a rigorous approach to semantics which does not reduce to "truth-theoretic" semantics (in terms I used earlier) or other paradigms that (I have argued) overstate the "logic" of meaning. In the case of "bachelor", for instance (see (144) and (145) above) the definition of bachelor as "unmarried man" is one of the most ubiquitous talking points in linguistics and philosophy, mostly seen as self-evident, or so it seems. I find this worth some curiosity because at least in modern English, or perhaps just modern (North) American English, "bachelor" (as far as I'm aware) is in fact rarely used simply to denote a man who is unmarried (in conversation one is more likely to say something like *John is single* or *John? He's not married*). Instead, when *bachelor is* used it tends to be for emphasizing "bachelor-like" qualities. In short, the conceptual *role* for bachelor is to characterize someone as having certain attributes we think of as typical of unmarried men, or to dispose us to notice such attributes, not to serve as a pithy shortcut for "unmarried man". This is why describing a *married* man as a "bachelor" is (I would argue) relatively permissible, even if perhaps a bit of a rhetorical flourish.

As mentioned above (footnote 19) we can consider logical entailments, such as *bachelor* implying (when used non-metaphorically) *unmarried man*, as pragmatic maxims for how concepts should be used, which is not the same as regarding such entailments as *constitutive* or *definitional* either of concepts or of lexemes that embody them. Certainly it is reasonable to assume in most cases that someone referred to as a bachelor is in fact male and unmarried. I would argue that implying such logical conclusions is part of the *role* of words (like bachelor) which do in fact have some sort of affiliated logical substratum (another example would be the implication that *to kill* entails, non-metaphorically, to *cause to die*). When, for instance, someone is described as an "eligible bachelor" we would indeed (notwithstanding any subtleties I've mentioned about the word's usages) typically hear the speaker as intending to convey (amongst other details) that said person is single and male (this is arguably the most familiar literal usage of "bachelor" in modern English — but note that the context where the word's simplest inferential role, at least when used literally, is most often activated only within a semi-entrenched phrasal unit; which may arguably be an evidentiary example of claims to the effect that bare inferential roles are not that common as the *primary* conceptual role of word-senses in their freest and most flexible form, ready for combination into different constructions and contexts).

However, *how* inferential roles relate to a word's overall *linguistic* meaning can be complex and indirect; for reasons cited I don't believe that the lexeme *bachelor* primarily serves to condense the double-predicate "male and unmarried" into a single word. The fact that *bachelor* implies *male and unmarried* is certainly, I would argue/concede, a factor in how the word is used in different contexts, but can be a factor in different ways. The speaker may connote *male and unmarried* as a side-effect of characterizing an unmarried man, but the primary goal would more likely be characterizing some man's lifestyle or demeanor — as Bradley Rives points out (see the footnote), an unmarried man with a long-time (maybe cohabitating) romantic partner is less likely to be referred to in these terms — or may do so for ironic effect when the word is used *metaphorically* about a *husband*.

In short, conceptual roles in the lexical context seem to *play off of* logical entailments in various ways, incorporating the expectation or conventionality of such entailments as a source for semantic and/or rhetorical effects, without being definitionally reducible to the entailments involved — particularly in conversational, quotidian speech (meanings may be more formalized and rigid in special registers, such as legal or scientific language). Granted that *lexical* "roles" in a rhetorical/idiomatic sense are not necessarily *conceptual* roles; but I would argue that such usages are indeed, often, linguistic manifestations of a properly conceptual stratum — e.g. the role of *bachelor* in how we are disposed to synthesize impressions of an adult man: to cue us for certain attributes (as much as to logically assert singleness) can I think be deemed a *conceptual* (not just *rhetorical*) role.

When conceptual roles are bound to (especially everyday) linguistic constructions, they take on further determinateness related to dialogic contexts (as compared to private thoughts). I would argue that one reason *brother* or *father* are more common than *sibling* or *parent* is that a common usage for these words refers to someone's *own* brother/father, and *this is my brother/father* sounds more conversational than *this is my sibling/parent*. Presumably these are lexical terri-

© 2021/05/24 Nathaniel Christen

tories where a gender-neutral and gender-specific term overlap, and given significant overlap ("competition", in a sense) between word alternatives one tends to evolve into a rarer, more formal role and the other to take on a more everyday register (cf. *attorney* and *lawyer*). There may be no explanation for why the gender-specific alternatives evolved to be more quotidian for the nuclear family (if anything precedents in English trend in the opposite direction, with gender-neutral terms, e.g. for animals and professions, becoming more commonplace). But given this background, we should certainly take a speaker's sense of formality/informality and differences of register as factors in their understanding of words. In everyday English, for instance, we refer to our male siblings as *brothers* because this is conventional usage; it may be appropriate thereby to summarize that "brother" *means* "male sibling", but such a gloss tends to elide relevant usage details. Logically *brother* and *sister* may then be seen as subtypes of *sibling*, where the former two differ from the latter only in their degree of detail (they carry more information because they allude to a classification based on gender as well as kinship relation-to-speaker) but such a purely logical picture fails to represent the contrasts in tone and register between them.

These cases illustrate a general idea that words do not often serve a conceptual *role* as simple logical constructions ("brother" as "male sibling" probably exemplifies such logical transparency better than "bachelor" as "unmarried man", but even here the logical structure is a *gloss* on meaning rather than an exhaustive exposition of meaning). And if a lexeme's *role* is rarely to simply coalesce an otherwise-compound logical construction, then word *meanings* are rarely logically transparent (to that degree) either.

I would make similar points vis-à-vis word pairings such as *red card*, or (as a case in point) *red hair*. Interestingly, I think the shades of red acceptable for *red hair* are far from prototypical of the red spectrum — someone whose hair is *dyed* a bright or pure reddish shade (e.g. the shade of a football *red card*) probably would not be said unhesitantly to have *red hair*. In this particular use-case (or so are my intuitions) we describe hair in terms of *natural* color, which excludes a lot of color space (in contrast to phrases like *red-dyed hair* or the like). Once again, I think these usage-patterns make sense in the context of conceptual roles — the *role* of describing hair in terms of a color is to describe a natural quality of someone's appearance, akin to a biological feature (we are less likely to say *he is tall* of a short person on stilts, or *he is hirsute* of a bald person wearing a wig). Similar comments could apply to *red wine*, which is a popular case-study in discussions about Conceptual Space theory; the range of red-like hues which are *naturally occurring* colors of wine falls outside the prototypical *red* spectrum.[21] Insofar as *red hair* or *red wine* restricts the extension of "red" to a narrower color-space than how this word is usually employed, this apparently reflects situational context which is latent in conventional usage (but not explicitly marked anywhere with *red* or *hair*/*wine* or with the generic adjective-noun construction). The specific adjective-plus-noun form embodied by *red hair* is, in short, tied by convention to the situation of characterizing "natural" hair color, and excludes (for typical usage, that which is not marked by an aura of oddness or rhetorical extravagance) the situation of *dyed* or otherwise altered hair.

This specific *red hair* example illustrates how a certain form of conceptual role semantics may be seen as an alternative (though also in a sense complementary) to conceptual *space* theories originating with the work of Peter Gärdenfors. In conceptual *space* terms, the pairing of two words (or two concepts) can result in a "blend" which may be modeled in terms of a quasi-figurative "space" demarcated by individual concepts. Colors are an almost canonical example of this theory because (via the color-pyramid) colors really do form a geometric space, and individual color-concepts (like *red*) can be, without too much idealization, represented as ("convex") "regions" within that space.

There is in short a subspace of the color pyramid (or red/green/blue cube or one of several other mathematical

[21] The case of *white* wine brings up further difficulties since white wine is never a color close to actual white, whereas red wine is at least somewhat reddish. The only way to preserve some sort of "prototype-based" theory in this case is apparently to take *white* here as essentially a metaphor or idiom (or maybe a metonym premised on "white" somewhat evoking an absence of color, creating at least an indirect plausible association between "whiteness" and *transparency*) such that *white wine* is entrenched phrasally to the point of overriding the constituents, i.e. once familiar with this usage we expect no entailment of *white wine* actually being *white*. I am inclined to point out that this whole analysis somewhat begs the question of why *white wine* demands such a purely idiomatic reading whereas *red wine* has a different analysis; is it pure happenstance which phrases evolve in meaning according to one pattern or another? If so the pressures leading to how phrases get entrenched (and why not others, e.g. "yellow" wine) would seem to be at least as important and interesting a semantic question as accounting for our conceptualization of meanings such as *red wine* or *white wine* in their status as conventionalized.

© 2021/05/24 Nathaniel Christen

formulations of color-space) which is prototypically red, and a somewhat larger space of reddish hues that would in many circumstances be deemed "red". This provides a "geometric" picture of the concept, which becomes active in various combinations, such as adjective-noun contexts (*little red box*, say, implies a range of color values and also a range of spatial dimensions for the box *qua* rectangular object). In "little red box" the color-dimensions from *red* and spatial dimensions implied by *box* are simply added on to one another (independently), but in other contexts a combination may somehow restrict the full range of subspace otherwise associated with a concept (*red card* implies something with both a specific shade of red and a specific shape, less flexible than "card" in general). In the case of *red hair* there is a similar narrowing effect; but I would argue that this modulation of the underlying "conceptual space" geometry (to the extent that we accept such a quasi-mathematical intuition as appropriate for meaning in general) is not latent in the combination of *red* and *hair* as a simple phrase, but rather in the unique role which has evolved for color terms in the specific context of hair-descriptions.

To put it differently, I feel that "conceptual spaces" by themselves do not fully explain the semantic dynamics at play, even if they capture some of the parameters involved. The deeper theory, I would argue — which is more explanatorily generative, more keyed to the causative factors in usage-patterns — would come from analyses of conceptual *roles*. (See [7], [44], [43], or [72] for recent, mathematically rigorous enhancements to Gärdenfors's original theory which deviate from the original model in thought-provoking ways, notably in terms of how one should formalize "convexivity".)

Conceptual Space theory also overlaps with (or can serve as one way to formalize) "prototype" semantics, and we can consider how roles relate to the prototypical meanings of their corresponding word-senses — for instance, the *prototype* of "bachelor" is presumably an unmarried (adult) male who *also* bears certain "bachelor-like" traits and habits. Cases outside this prototype (including a *married* man coyly described as a bachelor) are still compatible with the word-sense to the degree that they are close enough to be credible (a "bachelor-like" married man may be *similar enough* to a prototypical bachelor to justify the designation notwithstanding its *logical* incompatibility). In other words, resemblance-to-prototype can arguably explain the propriety of uses which explicitly violate logical constraints (for whatever rhetorical or metaphorical effect). This seems correct as far as it goes, but an undifferentiated appeal to "(prototype) resemblance" doesn't capture the full semantic dynamics, I would argue, any more than a simple mathematization of the relevant conceptual space. In the current example, there are different ways that a non-prototypical bachelor (viz., a person described as such) can resemble a (prototypical) bachelor enough that the choice of words is acceptable (enough to have communicative force). One *way* this resemblance can work is by someone having bachelor *traits* but not fit the "technical" logic (notably by being, in fact, married), so that the usage is heard as (communicative but) metaphorical or somehow ironic. Other partial-resemblances would work differently, e.g. a *single* man in a long-term relationship is technically a bachelor but does not have the attributes normally connoted by the word, so usage in *that* sort of context would presumably be heard as accurate and non-metaphorical but somehow vacuous; we would want to ascertain the *point* of the speaker choosing *that* word in context (perhaps to highlight for whatever reason his *not* marrying said hypothetical long-term partner, for instance). Such contrasts between different "directions" of generalization from, or deviation-from-but-also-resemblance-to, a prototype core (perhaps formalized via a sort of Conceptual Space model) can best be analyzed, I would argue, by identifying how concepts can have different *roles* wherein each role tends to permit usages which tolerate exceptions to a prototype in different ways (and likewise rests on resemblance to a prototype in *other* ways).

With that said, I think the combination of conceptual *role* and conceptual *space* theories (both of which I will discuss in more detail in the later essays) offers intriguing possibilities for semi-formal semantics, and for whatever reason has not (so far as I am aware) been investigated in detail, even if there are occasional references to conceptual roles within the Conceptual Space literature, and vice-versa. I will not develop this thesis further here, except indirectly insofar as, reiterating comments from a few paragraphs ago, I think *roles* piece together rigorously with notions of *grounding* and speaker-relative epistemics, which are themes in the next section.

© 2021/05/24 Nathaniel Christen

5 Conceptual and Discursive Aspects of Grounding

To this point, my core argument in this paper is that sentences are constructed in stages leading up to a concrete proposition (or, more precisely, a speech-act which refers to but may do more than merely assert a proposition). This is conventional wisdom and so hardly a theoretical insight on its own, but I have attempted to fill in several nontrivial details: first, that we can model the stage-wise unfolding of a sentence toward propositional completion in terms of individual word-pair; second, that type-theoretical models can assert formal guidelines on how stacks of word-pairs (where either word in a pair is also part of a new pair in the progressive stages toward completion); and, third, that the structure of this accumulation does not necessarily reciprocate the structure of the propositional content which, at the level of the entire sentence, becomes enunciated. Instead, propositions are often invoked indirectly, and the architecture of intermediate sentence-stages serves to trigger perceptual or otherwise extralingustic cognitive processes, rather than to define logical form as such. In totality, though, sentence-understanding does involve a progression to a complete propositional content, with the caveat the this progression (for the addressee) may involve the integration of multiple cognitive systems, and cannot be analyzed in purely linguistic terms.

Fully working through this theory, then, implies the need for an account of "progression to complete propositional content" which can in fact cut across diverse cognitive registers. A full treatment of such issues is outside the scope of this paper, but I will note several important details, as I see them. In particular, the "accretion of detail", as sentence-parts build up to a complete idea, is always centered around the speaker's epistemic relation to the relevant propositional content. (Of course, this centering may be indirect — a speaker may certainly discuss someone *else's* beliefs or propositional attitudes — but such discourse carries with it an unstated familiarity with that third party's beliefs or thoughts.) To understand how sentences aggregate to propositions, we should therefore pay particular attention to the connection which indirectly ties each element of a sentence to the speaker — and in particular to the speaker's epistemic stance vis-à-vis each element. For nouns: what is the space where the object is located (perhaps an abstract space) and how does it relate to spaces where the speaker is or has been perceptually situated, or observing somehow. For verbs: how does the speaker observe or become aware of the state or change embedded in the verbal semantics? For adjectives: how does the speaker observe the quality which an adjective asserts as inhering in a noun? And so forth.

Langacker's Cognitive Grammar is particularly oriented to the effects of epistemic "grounding" along these lines on sentence structure. The point is not simply that every element of a sentence is aligned vis-à-vis the speaker in some manner, and addressees need to unpack these alignments to understand (and evaluate) what the speaker is saying. More substantially, epistemic grounding creates a pattern within the sentence itself, including designating content as more or less peripheral to each sentence's grounding effects. The core subject of a sentence is also, in many cases, the anchor for how the sentence is grounded in the speaker's epistemic attitudes — this is the object which is the speaker's focus of attention, and the space and time where the speaker observes/encounters the object is a baseline against which other space/time frames (e.g. in subordinate clauses) may be offset. One essential contrast is whether the noun/object at the core of a sentence is a specific entity or something merely virtual — *the winner of the last election* compared to *the winner of the next election*. Langacker often uses a notation where a closed circle stands for a concretized discourse "anchor", whereas an open circle marks that the corresponding anchor is in some sense virtual, hypothetical, insubstantiated, or in some sense "conjured" in the course of a conversation (triggered by constructions such as "she wants", "he intends to get ...", "anyone who ..." and so forth).

Alongside this examination of anchoring, we can also consider how the anchored "process" or "state" — the head verb attached to the central noun — functions on the planes of virtuality/actuality and epistemic speaker-relation. The essential thrust of most sentences derives from the process or state which a speaker observes/reports in an anchored noun; subsidiary parts of the sentence then fill in details refining this central process. The relation of secondary clauses or phrases to the central nexus is often marked, for example, by different verb cases — a locative indicates where a process/event occurred; a benefactive for what or whom the action was performed (in cases where the central verb profiles

© 2021/05/24 Nathaniel Christen

a deliberately intended act); an instrumentive provides details about *how* the central action is, was, or will be carried out, and so forth. In general, specifications about *how*, *where*, *why*, *for whom*, and so forth can be (logically speaking) stripped away from a central verb/noun anchor. There's nothing theoretically deep or sophisticated about posing sentences' center/periphery distinction in these terms, but notice that such a theory, albeit simple and commonsensical, leads to a structural articulation within sentences that overlays and combines with type-theoretic and transform-pair relations I have discussed in this paper. For example, case-marked clauses (e.g. locative or benefactive phrases) serve at the "macro-type" level as adverbs, type-theoretically mapping between verbs insofar as adding detail to a verb-concept transforms it into a new verb-concept which is more cognitively complete: from *she bought a new dress*, say, which is complete in itself as far as having requisite verb, subject, and object, to her buying the dress *at* a store, *for* a person (or perhaps an event, etc.), *with* a gift card, and so on. All of these extra details modify the "buy" concept. Of course, the also modify the overall predicate of her buying the dress, but they do so through the verb; *buy* is the "site" of the added detail being semantically operative (contrast with *buy an expensive dress*, or *an elderly customer bought an expensive dress*, where added detail is sited through the subject or object and not the verb per se).

Given these ideas, we begin to see how words or phrases have facets which derive from the these units' respective relations to the underlying anchoring framework of a sentence — the central verb and noun profiled at the core of the sentence — as well as the nature of the sentence's core anchoring itself. Collectively these facets construct an epistemic tableau which unfolds across the components of a sentence, classifying details into various levels of logical centrality. The fact that a dress is *expensive*, for instance, may well be less central than that it is *bought* (compare *stolen*, *made*, *saw*, etc.). In order to properly interpret a sentence, hearers have to discern central/peripheral distinctions as part of the sentence's logical structure. This process happens to some degree via lexical coding alone, but extra-semantic features such as intonation and gestures often play a role as well. In this sense analysis of intonation and speech patterns — no less than syntax, semantics, discourse, or morphology — can be an organic element of the cognitive-linguistic architecture. I will in this section briefly review how prosody has been treated by several authors in the pragmatics and/or cognitive-grammar traditions.

5.1 *Discourse, Intonation, and Grounding*

Speakers tend to enunciate more quickly, and in lower tones, in conjunction with peripheral components of a sentence compared to centralized ones; higher-pitched accents and emphasis tend to mark central items such as anchor nouns and verbs. Such observations give rise to a kind of cognitive-linguistic framework for studying prosody and speech, which attends both to tonal variation (as notated, represented, and analyzed by conventional prosodic analysis) and to the cognitive/epistemic articulation of sentences. As a specific example, consider sentences such as:

- (241) // **The book she read** // was very interesting. //
- (242) // I recommended **a book** // **that she read immediately**. //
- (243) // **A book** just appeared // **that she read immediately** //
- (244) I recommended // a book I read years ago // to her.
- (245) This time **Harry** didn't cause our defeat.

Here (245) is from the *Handbook* (page 527/550), and (241)-(243) are from Langacker [41, pages 102-103]. These examples illustrate how patterns of enunciation reinforce anchoring and focusing effects in discourse: in the "book" examples the stress on that word marks it as the central point of attention in the sentence; and in (245) there are several potential stress patterns which cue the listener to different semantic interpretations.

Within this last group of examples, (242) stands out insofar as the technical noun/verb profiling anchor of the sentence does not align with stress patterns: overall the subject is *I* and the central verb is *recommended*, but the thematic emphasis of the sentence is on *she* reading the *book*. Langacker intimates that we are inclined to place spoken emphasis on the second-half clause in (242) precisely because we need to override the normal tendency to read subordinate clauses as supplying secondary detail rather than primary topical foci. In (244), for instance, which is my dreamt-up contrast to (242), I believe the natural pattern of intonation would vocally suppress and

© 2021/05/24 Nathaniel Christen

deemphasize the somewhat-tangential "that I read years ago" phrase, with the speaker aware that her expressed thought will not reach any level of completeness until concluding with "to her", so the vocal modulations are suggesting to the listener certain contours in primary and secondary phrasing. The (242) case is different — at least potentially, depending on context — insofar as the main idea the speaker wishes to convey is that "she" read the book "immediately", with the specific fact that said book was initially recommended by the speaker serving mostly to frame and introduce the topic of her actions vis-à-vis the book. In Langacker's other two examples "book" is also emphasized vocally, we assume, but naturally so because it also coincides with the actual subject of the sentence. In (243) the stress organically aligns with the *subject* (the book) but less so the profiled nexus *verb* (that it just appeared), which — as Langacker notates via bold highlighting — may be tonally deemphasized, with stress instead placed on the clauses stating her reading it immediately. In other words, here too the intonation shifts emphasis away from the central verb/object pair in part, but only in part, insofar as the subject now keeps its emphasis but the profiled verb loses the emphasis, at least partially.

Meanwhile, in the context of (245), chapter author Julia Hirshberg observes that the sentence has multiple (mutually exclusive) readings, and that prosody tends to signal that one or another is intended. If *Harry* is stressed, for instance, the implication is that someone or something *else* caused the relevant team's defeat, so the sentence asserts (or is performed in a context where it is common knowledge that) that team did actually lose, without blaming Harry. Conversely, if *defeat* is stressed, we might infer that the team was not in fact defeated; because the sentence overall is introduced via a negation, some element present on the surface level in the sentence is intended to be negated, and vocal stress can point listeners to where this negation should be applied. If "defeat" is negated, this connotes that the team had in fact won. Conversely, if *cause* is stressed, one could read that the team still lost, only not due to Harry. Hirshberg does not specifically address a scenario where no word acquires noticeable stress, but my inclination on imaging (245) spoken with mostly level or falling tones would be to hear a stress on the last word — triggering the reading (or hearing) that the team had not lost — with the caveat though that such intonation patterns would have a more pronounced effect of implying a certain common discursive reference-frame for all conversants. One consequence of stressed items is to call attention to material which is discursively "new"; conversely, the *absence* of stress implies that all elements in the statement just made are already entrenched in the current discourse. In short, I would assume an even-toned enunciation of (245) to be most likely when the conversants were already discussing Harry, the team, and whatever game they had recently played.

Other examples in the same *Handbook* chapter also touch on the theme of stress patterns signaling discourse novelty, as in:

- (246) (example 20, page 526/549)It's odd that Clyde married Bertha.
- (247) (example 18, 526/549)Londoners most often go to Brighton.

In (246), we would probably stress *Clyde* if we were already talking about Bertha, but not previously about Clyde; whereas stressing *odd* would imply that the prior conversation was about both Bertha and Clyde already, with the speaker opining about the fact of their getting married (rather than about who it was that ended up marring Bertha, if the conversation to that point had been about her, or conversely who Clyde married, if the discourse had been about him). Likewise in (247) stressing *Londoners* suggests focusing on people from London as the largest group from among those who visit Brighton, whereas stressing Brighton suggests focusing on a place which disproportionately serves as a destination — on the first reading we focus on people from London as opposed to *from* elsewhere; on the second we focus on going to Brighton versus *to* elsewhere.

Langacker and Hirshberg also both discuss prosodic effects in phrasal grouping which can determine the literal meaning of sentences, not just their shades of interpretation. These effects take on several aspects, such as distinguishing when subordinate clauses are modifiers to entire other clauses or just parts therein (see (248, taken from page 527/550)) or distinguishing lists from compountd nouns (see (249, taken from page 524/547)):

- (248) Bill doesn't drink because he's unhappy.
- (249) Stir in rice wine and seasonings.
- (250) Alice bought a dog. A poodle.

© 2021/05/24 Nathaniel Christen

In (248) intonation has to clarify whether the speaker believes Bill *does not* drink, or *does* but not because of unhappiness — which turns on whether "because he's unhappy" modifies the entire *Bill doesn't drink* phrase, implying that the speaker accepts that phrase as a true proposition (needing to be explained), or just the verb *drink* (so that the second clause's negation applies only to the idea of Bill's drinking being caused by unhappiness, without denying the drinking itself). Likewise in (249) we need to know whether to group *wine* with *rice* to form a compound noun, or conversely whether to keep *rice*, *wine*, and *seasonings* as sibling elements in the same discursive register; and in (250) [41, page 119] we would tend to enunciate the second sentence quietly to signal that it serves as an emendment to the first (which serves indirectly to link the poodle mentioned in the second sentence with the dog of the first; without that structuring in tonality the hearer might be unclear of which poodle is being topicalized).

Langacker furthermore argues that variations in tone and accentuation depend on epistemic factors latent in overall sentences. He gives these example (page 115):

- (251) The **ROOF LEAKS** if a building hasn't been **MAINTAINED**.
- (252) The **TEAM** is **UNSUCCESSFUL** because the **COACH** is so **INCOMPETENT**.

Langacker traces patterns of emphasis to the "window" within a sentence where a listener is likely to focus on a phrasal scale, and also to the real-world context where the objects and situations relevant to that window are situated. He then argues:

> Barring special emphasis — where activation is greater than it would otherwise be — the most active element is the one profiled at the composite structure level by the expression appearing in the window; in *the cat on the roof*, the cat is made salient in this manner. Active to a lesser extent are elements profiled by component expressions at lower levels of organization, in this case the roof and the locative relationship. Active to a lesser degree are unprofiled elements required for the conception of the profiled entity, e.g. the building evoked by *roof*. In (7)(a) [which is (251) here], its weakly active status in the first clause (due to *roof*) is sufficient for *building* to be accentually reduced in the second ... Latent notions are not explicit in a window but are primed due to spreading activation, hence easily accessible. Since a team often has a coach, for instance, the latter notion is latent in the former, though not required for its conception. This latency is sufficient for *coach* to take the definite article in (7)(b) [(251) here], but insufficient for the noun to be unaccented. Latency is of course a matter of degree, so there is no sharp boundary between latent elements and those which are dormant in the sense that their current activation has no linguistic consequences. [41, page 115]

In effect, Langacker argues that when we hear stress patterns we tend to align the stressed elements with objects *in the context of* their grounding vis-à-vis a sentence's core situation: because a roof is (an essential) part of a building, a discussion centered on the building can latently pivot to considering the roof (and vice versa). If a *roof* is mentioned in discourse, we instinctively contextualize in terms of the building to which it belongs, and so this building itself acts as if it were already introduced into the "ledger" even if it is not explicitly mentioned; hence it acquires the intonational pattern of something already present in the conversation (without the stress that would often accompany newly introduced content). Langacker argues (backed with a small-caps notation to indicate "medium" prosodic emphasis) that (252) is different because the concept of *coach*, while closely associated with a team, is not latent to the same extent as the building beneath a roof (one can perceive and conceive of a team without a coach, but not a roof without a building). Such degrees of cognitive inter-association combine with "windowing" effects, in terms of how we hear multi-word units (influenced by intonation as well as phrase-structure), to determine where we tend to place spoken emphases.

Langacker justifies the idea of a "window" of attention, which encompasses both our disposition to hear words grouped into mid-sized phrases and our inclination to focus attention on particular aspects of environing situations, by invoking our natural tendency to break a sentence down prosodically into chunks, separate and apart from formal syntactic boundaries. This engenders examples like (from page 111):

- (253) // If you see a dress // and you really like it // you

© 2021/05/24 Nathaniel Christen

should buy it // while you can. //

- (254) // At least for Alice // that wonderful dress // is just too expensive. //

Langacker notates these sentences to indicate that we hear them in roughly three- or four-word groups, even though larger-scale phrasal readings are also possible.

Elsewhere Langacker analyzes a different intonation/anchoring pattern, specifically the genre of sentences where a stressed and prosodically isolated clause anchors a sentence thematically, though not as the technical nexus, as in the first two of these [42, page 250]:

- (255) **Last night**, she was painting Easter eggs.
- (256) **With a fine brush**, she was painting Easter eggs.
- (257) Open the carton L-H%. Now remove the monitor carefully.
- (258) The Johnsons are solid citizens. They *H pay their *H taxes H-L%. They *H attend *H PTA meetings L-H%.

Here I juxtapose Langacker's examples with two from Hirshberg's chapter (page 533/556) cited above. These are similar examples, I contend, sharing the feature of a preliminary clause or sentence grounding subsequent content, in a fashion reinforced prosodically. Langacker implies that the "space-building" clause[22] receives particular vocal stress, presumably because it is cognitively foregrounded for the speaker. Intonation emphasizing the preliminary clause suggests that this material is somehow changing the focus of the conversation; e.g., that in (255) the speakers were not already discussing anything occurring "last night"; or in (256) the speaker is particularly focused on what may seem to be a granular detail. In the second *Handbook* example, Hirshberg argues that a pattern of relatively sustained stress — leading through an introductory/framing sentence into subsequent content — connotes an "open-ended set", in her words, as if the speaker wishes to imply that the discourse only mentions *some* ways the Jones are reputable. In (257) the discourse traces a mental grounding act that coincides with a linear sequence of concrete actions which the addressee either literally carries out, or contemplates: the "open the carton" instruction both mentally frames — and "activates" thought of — what is *in* the box, and also the discourse is staging or anticipating a specific enactive sequence, which starts with the box being opened and follows to engaging (literally, not just mentally) with its contents. This particular cognitive framing is connoted prosodically by the intonation across the sentence boundary, if (as would be natural were the example to be spoken aloud) the speaker evinces a "*continuation rise* contour, which is represented by a low phrase accent and high boundary tone (L-H%), generally interpreted as conveying that there is 'more to come'" (Hirshberg, page 533).

The prosodic details I have reviewed in the last few paragraphs are conventionalized in language to some degree, but not codified with the precision of lexical or syntactic norms, for example. This suggests for me that we tend to form intonation patterns instinctively rather than by rote adherence to linguistic rules; by extension, the cognitive drivers of such patterns may well emerge in large part from prelinguistic or extralinguistic intellectual strata. This idea seems quite plausible in relatively simple and straightforward examples of prosody; for instance, it seems speculatively plausible that we tend to subconsciously place greater vocal stress on a word or pair of words (such as a nexus noun/verb) which is the central organizing motif or topical focus of the sentence we are saying. Such a tendency could well emerge not through a linguistic medium per se but rather because these are the entities which figure most prominently in our mind; the greater weight they carry in the layering of our cognitive attention becomes subconsciously translated to a habit of enunciating the corresponding words, when we come to them, more slowly and emphatically. Similarly, we may be disposed to vocalize tangential clauses quietly and with fewer stress tones, and relatively central phrases slower and louder, because our thought processes in speaking itself — together with the larger phenomenon of our cognizing the overall scenario in which speaking and conversing is one part — requires us to mentally organize different aspects of the situation, sorting them into central and peripheral (to varying degrees), so each clause emanates from a cognitive register which we have already tagged as more or less foregrounded in our situational reasoning. If this speculation is accurate, patterns of inclination are a window onto contextual cognition and conceptualization as much as onto language itself — or, more precisely, they are a window onto language *because* they reveal the

[22] Note this is not a term Langacker uses in this context, though elsewhere he approvingly references Gilles Fauconnier's analyses of space-building.

© 2021/05/24 Nathaniel Christen

cognitive operations from which linguistic meaning often originates.

This line of reasoning, which here I have sketched only briefly, points to a potential for centralizing prosodic features and analysis more firmly within cognitive linguistics. It calls in particular for hybrid theoretical models which are attuned both to speaking/intonation patterns and to effects of epistemic grounding, anchoring, center/peripheral clausal patterns, and the general organizing principles through which sentences build up propositional content through an accretion of detail. Part of this theoretical synthesis may be suggested through notational conventions — for instance, it is reasonable to suggest a parse-representation that highlights anchoring nouns and verbs apart from subsidiary ones, and points to relations in terms of how peripheral (or "semi"-peripheral) elements connect to the nexus. In (242), for example, the sentence's thematic focus stands apart from what is technically the root verb/noun, so that the example is marked by a kind of double focus, a competition between centers which are each posed by cognitively impactful desiderata (syntactic versus thematic/topical), such that *book* and *I* (and *recommend*/*read*) have a different anchoring/peripheralizing relation than would normally be manifest between principle and subordinate clauses. In the case of (251) there is another kind of extra cognitive detail suggested by how our prosodic instincts, as we speak such a sentence and in the course of so doing instinctively thread intonational contours, emerges from patterns of cognitive activation and partial activation: a sentence (or clause or "thought") which profiles a build has latent within it the possibility of (perhaps via a subordinate clause or lexical antecedent) profiling its roof, and vice-versa. Alongside such anchoring-related inter-word relations we can also mention the more straightforward effects of secondary clauses and phrases adding (semi-tangential) detail to a nexus verb, particularly through different case-forms as I mentioned above (locative, benefactive, and so forth). These anchoring-related details can all be potentially notated in a representation parallel to parse-graphs. If we then *also* consider prosodic patterns, which in turn are represented via different specialized languages or markup formats endemic to speech technology, the idea of a hybrid cognitive/prosodic theory can be evoked through a representational paradigm. This idea is explored a little more in the code supplementing this paper.

5.2 *Intonation Patterns in the Context of Dyanamic Patterns of Enunciation*

Over the next few pages I will sketch a brief argument for how prosodic features can be analyzed alongside syntax, semantics, and pragmatics — that is, how the "selective pressures" which guide choices of which words to use, or which morphological and syntactic patterns to apply, and which discursive norms to obey, also draw in questions of how we should prosodically emphasize (or de-emphasize) certain words. On this argument, there is no *a priori* reason for excluding intonation from other aspects of languages. There are prosodic conventions which we internalize as language users no less than entrenched word-uses or grammatic forms.

On the other hand, I would argue in general that conventions at any layer or scale of linguistic phenomena are applied only indirectly, and that formulating spoken (or by extension written) utterances is a dynamic process whereby conventional patterns are at best only guidelines. In the case of word-choice or syntactic form, we are compelled (I would argue) to clothe our expressions in their precise form because we generally use the unfolding nature of language as a tool to simulate the unfolding nature of our thoughts, so that we construct the various inflection points of our discourse — the nonobvious word choices, the speech pauses and other grouping effects, the phrase-boundaries, and so forth — to describe a *dynamic structure* within each sentence. We can then use the details of this dynamic structure as a kind of second-order reference conveying facts about what we are saying — *why* we believe something to be the case, or why we feel it is of interest to our listeners. Individual moments of our enunications should be analyzed against the background of this epistmic/dialogic continuum, through which speakers do not only assert judgments but recapitulate their rationale and (relative to current discourse/dialog/situations) importance. In *that* context, I think we can make a strong case for prosody being governed by the same selective/signifying pressures as lexical/semantic, syntactic/morphological, and discursive/dialogic norms, so that *in principle* intonation should always be considered as folded in to Cognitive Grammar.

© 2021/05/24 Nathaniel Christen

For sake of discussion, I motivate my comments via two examples similar to ones presented by Michaelis and Lambrech in [49, e.g. page 223]:

- (259) It's AMAZING, the things she says sometimes.
- (260) It's RIDICULOUS, what he's saying.

Presumably the speaker of (259) heard something "she" said and was moved to respond. This situation implies several factors, including (1) that there was something noteworthy about what she just said (which could be that it was unusually intelligent, if "she" is a young girl, or else somehow bizarre or unexpected, and so forth) and (2) that the current situation reminds the speaker of other situations which resemble it to some degree. That is, the comment-worthy intelligent/bizarre (or whatever) nature of something she has just said is only one in a series of comparably attention-getting comments of hers. The speaker implies that this is not the first time he or she has noticed this tendency; for whatever reason the speaker chooses this occasion to verbally call attention to it.

Upon hearing (259) we infer that the speaker has noticed such a pattern, probably also on occasions prior to the current context, but is now choosing to verbalize his or her thoughts. This implies among other things that the speaker feels these thoughts are at least minimally interesting to other conversants, that they are not too mundane to warrant being registered as topics of collective discussion. Things which are out of the ordinary do in general tend to engage collective interest more than things which are expected; describing "the things she says" as "amazing" suggests that they are unusual enough to be worth thinking about. The characterization "amazing" seems to comment both on the speaker's own impression but also, more meta-discursively, on the speaker's belief that others would agree that "things she says" are indeed interesting enough to justify someone calling attention to them.

Our interpretation of (259) moreover intuits that there are several stages in the speaker's formulation prior to choosing to say (259) in particular and believing that (259) would be at least marginally interesting to others. We gather that the speaker has heard something "she" recently said and that it particularly caught the speaker's attention; that the attention-grabbing force of the observed comment reminded the speaker of other occasions; and that the unusualness of the new comment along with the recurrence of that pattern seems newsworthy enough to the speaker to warrant registering a comment in the collective dialogue. Interestingly, (259) seems less likely to warrant enunciation *without* the recurrence of the pattern — if "she" had not on other occasions said something "amazing" — but the speaker does not begin the enunciation by making reference to those prior occurrences, and indeed they are connoted only obliquely. The commencement of (259) — "It's amazing" — certainly describes first and foremost a property of whatever she has *just* said. Other "amazing" things she has said are implied to exist by the use of the present tense ("things she says" rather than, say, "things she just said"). So the enunciation restages a thought-process where the speaker notes first the "amazingness" of what she *just* said and then connects it mentally with prior similar occasions. That whole mental process was triggered by that first sense of there being something "amazing", and the force of that impression for the speaker presumably warrants the guess that the situation is worth talking about. So that "local" amazingness seeds the speaker's formulation and also motivation to vocalize the resulting thoughts, and only subsequent to that initial amazingness-impression does the speaker seem to recall other comparable impressions from the past. Of course, the existence of those past occasions further warrants the comment-worthiness of the speaker's impressions. So we intuit that the speaker undergoes a thought process starting from being struck by an initial amazingness, mentally connecting it to a pattern evinced in prior occasions, judging the combined weight of the recurrence of the pattern and the degree of remarkableness of the current situation to warrant saying these thoughts aloud, and then *circling back* to the initial amazingness-impression to retrace these thoughts for the listener, acknowledging at some level that any speaker has a conversational responsibility to clarify why their contribution to the current dialog warrants collective attention.

We understand thought processes to be generally complete prior to expressing them verbally — at least in canonical cases, though certainly it is common for people to "search for the right words" mid-sentence — but we also perceive that speakers feel obliged to justify their time on the dialog floor, so part of conversing is connoting why we think what we say is important. This often

© 2021/05/24 Nathaniel Christen

means retracing for others the origin of our beliefs or impressions — I judge that x because I have observed y; or to be specific to the current example I am struck by what she just said (partly because I've been struck by things she says in the past). What we say tends to unfold in ways which, alongside agreeing to syntactic norms, also reciprocates the process of our formulating beliefs/impressions relevant to the utterance, presumably because retracing the origins of assertions or observations lends them greater persuasive force.

This is the cognitive/dialogic context where I believe we should analyze the unfolding of (259): the initial *amazing* is *stressed* prosodically because it has a specific role in the overall setup of the speaker's retracing and warranting their coming public with their thoughts. The "amazingness" triggered the thought-process to begin with, so it is plausible to begin the public registering of the thoughts there. At the same time, this formation deviates from English conventions, since we're hearing an adjective without knowing what noun it modifies; presumably the speaker realizes this and compensates by stressing the adjective, signalling that the noun itself will come afterward. Perhaps at some level we are inclined to intone stress on a case like *amazing* because we wish to clarify that this word is not the end of the sentence ("it's amazing" by itself could be registered as a complete thought, albeit context-bound and possibly heard as a kind of verbalization of private thoughts maybe intended to make the speaker's sentiment known but not to elicit further attention). The stress on *amazing*, on such a reading, is in part a way of the speaker "keeping the floor", which among other things deflects confusion vis-à-vis the adjective ground; overall then the performance to that point creates a pronounced expectation (*what* is amazing, in fact?), helping drive listeners' attention to the content which completes the expectation (*the things she says sometimes*). The fact that this construction — in the present tense and adumbrated by "sometimes", both with the effect of connoting the present situation as part of a recurring pattern — fulfills an expectation helps to register the pattern referenced via the expression-fulfilling phrase as a center of attention, the phrase as a whole being to some extent "framed" by the expectation it fulfills. Overall, then, the specifics of (259) as an expression — the intonation on *amazing*, the abnormal syntax, and so forth — maneuver to highlight a pattern extending back into the past while justifying an act of commenting on the pattern by connecting it to the present context.

In short, I believe that the specific details of how (259) is constructed and enunciated represents an unfolding process shaped by the speaker's private though-processes and also by his or her intuitive sense of conversational propriety. Exact word (and prosodic) choices are driven by both epistemic and dialogic pressures, with speakers feeling a general need to explain their beliefs/judgments and also to warrant why their thoughts might be of interest within a conversation. These pressures yield subtle inflection points where our enunciations evolve in specific ways, marked as specific lexical, syntactic, or prosodic choices. This enunciatory unfolding though specific choice-points applies even if we are not consciously selecting words (or intonations) while in fact speaking; whoever says (259) could very well have formed a deliberate intention to express (259) exactly as written before starting to speak. But we hear enunciations as retracing a formulative process whose patterns and branches convey selective meanings (significance attached to selecting a word, morphology, or intonation in lieu of other possibilities) even if the process is "rehearsed" in the speaker's mind ahead of time, so that however much there is a meaning in its unfolding-pattern we also hear the speaker's enunciation as a completed thought, which the speaker has judged to be complete prior to speaking.

The specific patterns analyzed by Michaelis and Lambrecht (from which I took (259) and (260) with modifications) — in particular the unexpected and prosodically emphasized initiation of a sentence with an adjective like *amazing*, and then supplying a ground for the adjective in present-tense form — implies that there is a time frame or patterning which extends beyond the current moment of enunciation. Michaelis and Lambrecht point out that plural (lack of) agreement along with tense help craft our sense of the designated pattern; we do not appear to hear any anomaly in the non-plural *it's amazing* compared to hypothetical *they're amazing*. Examining similar cases, Michaelis and Lambrecht suggest that the designation is not the specific phenomena governed by the adjective but the overall facts of that adjective applying; in their example the phrase "it's unbelievable the people who are verbally abusive to fat people" is glossed as "not [the] people in themselves" being "unbelievable", but "the cardinality of the set to which they belong"; the fact that there are *many* such people. This is why we

© 2021/05/24 Nathaniel Christen

hear *it's unbelievable* instead of *they're unbelievable* as actually better-formed, even though the two parts of the sentence seem to disgree on number.

However, this specific kind of pattern — which includes a vocal stress on an adjective as well as choosing number and tense that calls attention to a *fact* (of her often saying amazing things, of *many* people abusing someone overweight) more generally than a particular instance — only applies in certain circumstances, and one criterion is that the "fact" suggested can be enlarged outside the scope of the immediate context. Presumably she has *also* said amazing things in the past, and presumably the speaker moved to comment on "fat shaming" has observed a *pattern* of abusive behavior toward overweight people. Both of these connoted patterns are understood to extend outside the frame imposed by the specific dialog currently under way. The phrasing in (260) has a similar implication of enlarged scope, but the implication here is more that the connoted scope is more continuous with the current context. Presumably in (260) the speaker can mentally frame a specific segment of the person's "ridiculous" comments that motivated (260) being said — perhaps the speaker and other conversants are watching a politician on television. The pressure to motivate one's dialog turns implies that a frame can be constructed within the overall situation of that politician speaking which is specifically referenced by, and triggered, the speaker's actually saying (260). But (260) also implies that there is a second scope containing that primary frame, because the word-choice is *things he is saying* (plural and progressive-tense) and not *what he just said* (singular and past tense); only in contrast to (259) (partly due to the tense being progressive and not simple present) the implication is that the larger scope applies not to a patterns of prior occasions stretching back into the past, but rather to the current situation (e.g., the current interview or speech which the conversants are witnessing) but apprehended on a broader scale.

Although the situations may vary somewhat in structure, the specific discursive and prosodic combinations evident in (259) and (260) (and in Michaelis and Lambrecht's examples) are only apparently acceptable in specific *kinds* of situations. Michaelis and Lambrecht thereby use these cases as motivating examples for treating these specific scenarios as a kind of situational prototype which has a corresponding sentence-level constructional pattern, something roughly sketched as "it's x, the y...". They analyze the situational prototype (without using that term itself) and its corresponding encoding in language in terms of Construction Grammar, implying that a certain conventionalized construction has evolved to serve as a kind of holistic expression of the corresponding situational pattern. From the point of view of prosody and cognitive grammar, this analysis is interesting because it emphasizes how language has to convey cognitively unified patterns in situational understanding, not just discrete referring or describing acts; and also because they assume (and use these examples among other as motivation for the claim that) intonations can be features of entrenched constructions as much as semantics and syntax.

I would take issue with the Construction Grammar analysis as presented if we find that it implies a purely holistic mapping of the pattern to the situation — i.e., that the construction comes to a speaker's mind more or less integrally once presented with a situation roughly warranting the "it's x, the y..." kind of precis, with the same immediateness of a lexeme to a referent, say — so as to undermine how situation-to-construct associations should always be treated as unfolding and reciprocative of thought processes. In other words, I would find it less helpful to imply that construction like "it's x, the y..." are static recipes which we as speaker hold in our mind and trot out whenever a situation is sufficiently close to some prototype to match the construction as its preferred linguistic articulation (which is not to imply that Construction Grammar commits its advocates to such a static implication). I would prefer to say instead that what we *linguistically reconstruct* as a Construct is more of a recurring process whereby certain patterns of formulating some thought about a situation, then rehearsing one's vocal comments embodying those thoughts, and then actual speaking (in dialog contexts) tend to recur in similar ways, under the relevant epistemic and dialogic pressures. The repetition of these patterns, which we may observe in our speech-community, can presumably make certain constructions somewhat entrenched, so we are motivated in contexts like (259) to choose this specific manner of speaking in part because of (259)'s resemblance to a "it's x, the y..." pattern, which we assume will be familiar to our listeners and therefore prompt their interpretation of what we say. But even when matching the situational constraints on our utterance-formulation against a conventionalized construction, I

© 2021/05/24 Nathaniel Christen

would still argue that the construction serves as only a rough guide, and we still have to deem our utterance in the listener's mind as an unfolding process which should reciprocate in some manner our rationale for expressed beliefs/judgments and how they connect to the current conversation. Certain choices such as prosodically stressing *amazing* arise from our sense of *unfolding* of our utterance from the listener's point of view, as well as potentially from construction-patterns which may be present in the background. Still, with that caveat, these analysis can give a concrete example of how prosodic choices form the performative and communicative substrate of discursive unfolding — of the inflection points where our utterances present specific intonation, lexical, and/or syntactic choices that collective build up meanings and supply sentences with illocutionary force by demonstrating/warranting assertions through a reciprocation of their epistemic provenance.

According to Miriam Fried and Jan-Ola Östman,

> The primary motivation for Construction Grammar is the insight that the juxtaposition of two or more forms seldom results in a simple concatenation of the meanings those forms might have in isolation. Consequently, Construction Grammar sees linguistic units as particular associations between form and meaning that must be represented as such, rather than leaving such associations to the operation of a set of rules for how to combine individual forms. [25, page 3]

I would qualify this paradigm merely by suggesting that we should not discount the compositionality of language solely on the basis of linguistic form not seeming to reciprocate *logical* compositions, because — as I have stressed numerous times here — there are other ways for aggregate meaning to be built from component parts than via mimicking predicate composition.

My own theoretical intuitions are that the evident shape of linguistic structures is due to particular sequencing in the "accretion of detail" that unfolds in a certain evolutionary sense in the course of formulating a sentence, and that conventionalized patterns are rough guides to dynamic processes rather than static templates that are passively filled in. Such "unfolding" actually, I think, occurs at several levels, first in our coming to cognitive awareness of some propositional attitude, then in our rehearsing how to assert (and warrant) that illocutionary content (in part by restaging for others' benefit the logically prior phenomenon of our own internal coming to that awareness), and finally in actual speech (where we may be influenced by observing others' reactions). The fact that "unfolding" occurs on these three distinct scales and time-frames complicates the picture of sentence-construction as a dynamic process, but as a thumbnail we can say that the typical sentence is a *complete idea* — formulated in advanced — but it is formulated in such as manner as to reciprocate, describe, or otherwise signify an unfolding (largely epistemic) process; to communicate the equivalent of *this is why I believe* x, or to derive how I can present it with a rhetorically persuasive ordering. The unfolding of linguistic structure is "dynamic" not only (or not always) in the incomplete sense that a sentence may be changed mid-speech, but rather that even utterances fully planned in advance have a dynamic structure constructed with the acknowledgment that a complete idea is built up in temporal stages, a partiality and accretion-pattern which can be used for discursive effect. This is an indirect dynamism, since we have a chance to "edit" the dynamic unfolding so as to form a fully worked-through expression-intention, but the dynamic nature of semantics and grammar, however indirect and rehearsed, is (I claim) essential to observed syntactic and semantic structures. In light of arguments such as those from Construction Grammar, I would argue also that prosody can be bundled into these claims as well.

I claim moreover that the dynamic/unfolding aspect of linguistic structure (however much we need to qualify this dynamism, because it is also wrapped in a certain prestaging of enunciatory planning) is not adequately modeled by most linguistic formalizations (being better treated with only semi-formal methods such as Cognitive Grammar). I have implied earlier, however, that a more rigorously formal account of linguistic structure which *does* recognize language's dynamic contours can be developed through type theory, which in its syntactic form captures landmarks in unfolding structure (since partially-complete phrases can be given type signatures) and in its semantic form captures conceptual refinements and situational expectations which are part of language's unfolding patterns. I have only barely sketched a thorough treatment of type-theoretic semantics and syntax along these lines, but perhaps the material I have presented conveys my rationale for believing that type-

© 2021/05/24 Nathaniel Christen

theoretic methods have potential to unify the formal and semi-formal branches of linguistic analysis in productive ways.

To that end, in any case, I will conclude with a brief discussion of type theory which returns to some arguments I left incomplete earlier.

5.3 *Types, Sets, and Concepts*

So far I have considered types mostly in their syntactic role of marking the *order of accretion* of detail; linguistic form cannot just present predicate content as a set of disjoint elaborations on a central situation. To be sure, sometimes different details are mutually independent enough that there is no specific order germane to their articulation — *I bought wine for your party in Ontario* and *I bought wine in Ontario for your party* both work. Nevertheless, phrases often *are* semantically interdependent and, in general, syntactic norms need to be entrenched so that there is a pattern in how phrases are pieced together into larger wholes, since many applications of syntax rules *will* apply to interdependent parts. The "accretion of detail" at the heart of any sentences' invocation of propositional content has to unfold in an ordered fashion, where some phrasal or word-pair units logically, vocally, and cognitively follow/precede others. This forms the organizational premise behind part-of-speech rules, I would argue, and, by extension, behind the system of linguistic types at the coarsest ("macrotype") level. At more granular levels, however, types (or so I have argued) hue more closely to concepts and word-senses, their role being to provide semantic to lexical conventions and/or to negotiate mid-scale conceptual manipulations such as guiding how to conceive of Liverpool (geographically, politically, etc.) or how to employ metaphors that attribute erstwhile beliefs or desires to nonsentient things. In short, to understand how type-theoretic ideas can influence semantics as well as syntax, it is necessary to extend the "system" of types to more granular scales. My argumentation has been premised, however, on the idea that the partition of types into different scales or fineness is at least to some extent fuzzy, so that all these different layers of type-inspired analysis potentially interpenetrate. In this sense syntactic analysis of types in terms of "organizing accretion of detail" is interconnected with analysis in terms of ontological shifts in register, for example, to a degree which would be less theoretically foregrounded if instead of type-theoretic framing we were instead to conceptualize "macrotypes" as parts of speech, say, and "mesotypes" as criteria for word-sense disambiguation.

In formal/computational contexts, types can be defined as sets of both values and "expectations" [13] (meaning assumptions which may be made about all values covered by the type); alternatively, we can (perhaps better) consider types as *spaces* of values. Types' extensions have internal structure; there can be "null" or "invalid" values, default-constructed values, and so forth, which are "regions" of the conceptual space spanned or encompassing types. There is definitional interdependence between types and functions: a function is defined in terms of the types it accepts as parameters and returns — rather than its entire set of possible inputs and outputs, which can vary across computing environments.[23] These are some reasons why in theoretical Computer Science types are not "reduced" to underlying sets; instead, extensions are sometimes complex spaces that model states of, or internal organization of comparisons among, type instances.

Moreover, despite mathematical set theory's influence on formal logic (and philosophy, by extension), philosophers have long been troubled by the seeming arbitrariness and minimal structure of the essential notion of *sets* themselves, and have explored other constructs to play analogous intellectually foundational roles, such as Michael Jubien's Property Theory [34] or Jiri Benovsky "modal perdurants" [9, p. 8] (here citing two treatments I find particularly persuasive or thought-provoking). This is one *philosophical* interest in type theory as types have more formal structure and also, in a sense, more metaphysical coherence than "sets" (in the classical sense allowing set-composition from unrelated members). In a certain intellectual role, however, types are still somehow "set-like" in that we can consider the totality of values that may inhabit each type, and we need some theory or model of that totality — which in turn engenders

[23] Moreover, expectations in a particular case may be more precise than what is implied by the type itself — it is erroneous to assume that a proper type system will allow a correct "set of values" to be stipulated for each point in a computation (the kind of contract enforced via by documentation and unit testing). So state-space in a given context may include many "unreasonable" values, implying that within the overall space there is a "reasonable" subspace, except that this subspace may not be crisply defined.

© 2021/05/24 Nathaniel Christen

different versions of type theory.

An obvious paradigm is organizing type-extensions around prototype/borderline cases — there are instances which are clear examples of types and ones whose classification is dubious. I contend, however, that common resemblance is not always a good marker for types being well-conceived — many useful concepts are common precisely because they cover many cases, which makes defining "prototypes" or "common properties" misleading. Also, sometimes the clearest "representative" example of a type or concept is actually not a *typical* example: a sample letter or model home is actually not (in many cases) a real letter or home. So resemblance-to-prototype is at best one kind of "inner organization" of concepts' and types' spaces of extension.

Sets, concepts, and types represent three different primordial thought-vehicles for grounding notions of logic and meaning. To organize systems around *sets* is to forefront notions of inclusion, exclusion, extension, and intersection, which are also formally essential to mathematical logic and undergird the classical interdependence of sets, logic, and mathematics. To organize systems around *concepts* is to forefront practical engagement and how we mold conceptual profiles, as collections of ideas and pragmas, to empirical situations. To organize systems around *types* is to forefront "functions" or transformations which operate on typed values, the interrelationships between different types (like subtypes and inclusion — a type can itself encompass multiple values of other types), and the conceptual abstraction of types themselves from the actual sets of values they may exhibit in different environments. Sets and types are formal, abstract phenomena; whereas concepts are characterized by gradations of applicability, and play flexible roles in thought and language. The cognitive role of concepts can be discussed with some rigor, but there is a complex interplay of cognitive schema and practical engagements which would have to be meticulously sketched in many real-world scenarios, if our goal were to translate conceptual reasoning to formal structures on a case-by-case basis. We can, however, consider in general terms how type-theoretic semantics can capture conceptual structures as part of the overall transitioning of thoughts to language.

A concept does not merely package up a definition, like "restaurant" as "a place to order food"; instead concepts link up with other concepts as tools for describing and participating in situations. Concepts are associated with "scripts" of discourse and action, and find their range of application through a variegated pragmatic scope. We should be careful not to overlook these pragmatics, and assume that conceptual structures can be simplistically translated to formal models. Cognitive Linguistics critiques Set-Theoretic or Modal Logic reductionism (where a concept is just a set of instances, or an extension across different possible worlds) — George Lakoff and Mark Johnson, prominently, argue for concepts' organization around prototypes ([39, p. 18] [33, p. 171, or p. xi]) and embodied/enactive patterns of interaction ([39, p. 90], [33, p. 208]). Types, by contrast, at least in linguistic applications of type theory, are abstractions defined in large part by quasi-functional notions of phrase structure. Nevertheless, the *patterns* of how types may inter-relate (mass-noun or count-noun, sentient or non-sentient, and so forth) provide an infrastructure for conceptual understandings to be encoded in language — specifically, to be signaled by which typed articulations conversants choose to use. A concept like *restaurant* enters language with a collection of understood qualities (social phenomena, with some notion of spatial location and being a "place", etc.) that in turn can be marshaled by sets of allowed or disallowed phrasal combinations, whose parameters can be given type-like descriptions. Types, in this sense, are not direct expressions of concepts but vehicles for introducing concepts into language.

Concepts (and types also) are not cognitively the same as their extension — the concept *restaurant*, I believe, is distinct from concepts like *all restaurants* or *the set of all restaurants*. This is for several reasons. First, concepts can be pairwise different not only through their instances, but because they highlight different sets of attributes or indicators. The concepts "American President" and "Commander in Chief" refer to the same person, but the latter foregrounds a military role. Formal Concept Analysis considers *extensions* and "properties" — suggestive indicators that inhere in each instance — as jointly (and co-dependently) determinate: concepts are formally a synthesis of instance-sets and property-sets [83], [8], [80]. Second, in language, clear evidence for the contrast between *intension* and *extension* comes from phrase structure: certain constructions specifically refer to concept-extension, triggering a mental shift from thinking of the concept as a schema or prototype to

© 2021/05/24 Nathaniel Christen

thinking of its extension (maybe in some context). Compare:

- (261) Rhinos in that park are threatened by poachers.
- (262) Young rhinos are threatened by poachers.

Both sentences focus a conceptual lens in greater detail than *rhino* in general, but the second does so more intensionally, by adding an extra indicative criterion; while the former does so extensionally, using a phrase-structure designed to operate on and narrow our mental construal of "the set of all rhinos", in the sense of *existing* rhinos, their physical place and habitat, as opposed to the "abstract" (or "universal") type. So there is a familiar semantic pattern which mentally transitions from a lexical type to its extension and then extension-narrowing — an interpretation that, if accepted, clearly shows a different mental role for concepts of (other) concepts' *extension* than the concepts themselves.[24]

Concepts, in short, do not mentally signify sets, or extensions, or sets-of-shared-properties. Concepts, rather, are cognitive/dialogic tools. Each concept-choice, as presentation device, invites its own follow-up. *Restaurant* or *house* have meaning not via idealized mental pictures, or proto-schema, but via kinds of things we do (eat, live), of conversations we have, of qualities we deem relevant. Concepts do not have to paint a complete picture, because we use them as part of ongoing situations — in language, ongoing conversations. Narrow concepts — which may best exemplify "logical" models of concepts as resemblance-spaces or as rigid designators to natural kinds — have, in practice, fewer use-cases *because* there are fewer chances for elaboration. Very broad concepts, on the other hand, can have, in context, too *little* built-in *a priori* detail. (We say "restaurant" more often than *eatery*, and more often than *diner*, *steakhouse*, or *taqueria*). Concepts dynamically play against each other, making "spaces" where different niches of meaning, including levels of precision, converge as sites for one or another. Speakers need freedom to choose finer or coarser grain, so concepts are profligate, but the most oft-used trend toward middle ground, neither too narrow nor too broad. *Restaurant* or *house* are useful because they are noncommittal, inviting more detail. These dynamics govern the flow of inter-concept relations (disjointness, subtypes, partonymy, etc.).

Concepts are not rigid formulae (like instance-sets or even attributes fixing when they apply); they are mental gadgets to initiate and guide dialog. Importantly, this contradicts the idea that concepts are unified around instances' similarity (to each other or to some hypothetical prototype): concepts have avenues for contrasting different examples, invoking a "script" for further elaboration, or for building temporary filters. In, say,

- (263) Let's find a restaurant that's family-friendly.

allowing such one-off narrowing is a feature of the concept's flexibility.

In essence: no less important, than acknowledged similarities across all instances, are well-rehearsed ways vis-à-vis each concept to narrow scope by marshaling lines of *contrast*, of *dissimilarity*. A *house* is obviously different from a *skyscraper* or a *tent*, and better resembles other houses; but there are also more nontrivial *comparisons* between houses, than between a house and a skyscraper or a tent. Concepts are not only spaces of similarity, but of *meaningful kinds of differences*.

To this account of conceptual breadth we can add the conceptual matrix spanned by various (maybe overlapping) word-senses: to *fly*, for example, names not a single concept, but a family of concepts all related to airborn travel. Variations highlight different features: the path of flight (*fly to Korea*, *fly over the mountain*); the means (*fly Korean air*, *that model flew during World War II*); the cause (*sent flying (by an explosion)*, *the bird flew away (after a loud noise)*, *leaves flying in the wind*). Words allow different use-contexts to the degree that their various *senses* offer an inventory of aspects for highlighting by *morphosyntactic* convention. Someone who says *I hate to fly* is not heard to dislike hand-gliding or jumping off mountains.[25] Accordant variations of

[24] There is a type-theoretic correspondence between intension and extension — for a type t there is a corresponding "higher-order" type of *sets* whose members are t (related constructions are the type of *ordered sequences* of t; unordered collections of t allowing repetition; and stacks, queues, and deques — double-ended queues — as t-lists that can grow or shrink at their beginning and/or end). If we take this (higher-order) type gloss seriously, the extension of a concept is not its *meaning*, but a different, albeit interrelated concept. Extension is not definition. *Rhino* does not mean *all rhinos* (or *all possible rhinos*) — though arguably there are concepts *all rhinos* and *all restaurants* (etc.) along with the concepts *rhino* and *restaurant*.

[25] People, unlike birds, do not fly — so the verb, used intransitively (not flying *to* somewhere in particular or *in* something in particular), is understood to refer less to the physical motion and more to the socially sanctioned phenomenon of buying a seat on a scheduled flight on an

© 2021/05/24 Nathaniel Christen

cognitive construal (attending more to mode of action, or path, or motives, etc.), which are elsewhere signaled by grammatic choices, are spanned too by a conceptual space innate to a given word: senses are finer-grained meanings availing themselves to one construal or another.

So situational construals can be signaled by word- and/or syntactic form choice (locative, benefactive, direct and indirect object constructions, and so forth). Whereas conceptual organization often functions by establishing classifications, and/or invoking "scripts" of dialogic elaboration, cognitive structure tends to apply more to our attention focusing on particular objects, sets of objects, events, or aspects of events or situations. So the contrast between singular, mass-multiples, and count-multiples, among nouns, depends on cognitive construal of the behavior of the referent in question (if singular, its propensity to act or be conceived as an integral whole; if multiple, its disposition to either be divisible into discrete units, or not). Or, events can be construed in terms of their causes (their conditions at the outset), or their goals (their conditions at the conclusion), or their means (their conditions in the interim). Compare *attaching* something to a wall (means-focused) to *hanging* something on a wall (ends-focused); *baking* a cake (cause-focus: putting a cake in the oven with deliberate intent to cook it) to *burning* a cake (accidentally overcooking it).[26] These variations are not random assortments of polysemous words' senses: they are, instead, rather predictably distributed according to speakers' context-specific knowledge and motives.

I claim therefore that *concepts* enter language complexly, influenced by conceptual *spaces* and multi-dimensional semantic and syntactic selection-spaces. Concepts are not simplistically "encoded" by types, as if for each concept there is a linguistic or lexical type that just disquotationally references it — that the type "rhino" means the concept *rhino* ("type" in the sense that type-theoretic semantics would model lexical data according to type-theoretic rules, such as *rhino* as subtype of *animal* or *living thing*). Cognitive schemata, at least in the terms I just laid out, select particularly important gestalt principles (force dynamics, spatial frames, action-intention) and isolate these from a conceptual matrix. On this basis, we can argue that these schemata form a precondition for concept-to-type association; or, in the opposite logical direction, that language users' choices to employ particular type articulations follow forth from their prelinguistic cognizing of practical scenarios as this emerges out of collections of concepts used to form a basic understanding of and self-positioning within them.

In this sense I called types "vehicles" for concepts: not that types *denote* concepts but that they (metaphorically) "carry" concepts into language. "Carrying" is enabled by types' semi-formal rule-bound interactions with other types, which are positioned to capture concepts' variations and relations with other concepts. This is the relationship I would propose for integrating theories of concepts — as cognitive and lexicosemantic phenomena — with applications of type theory that, as I argued above, are fundamentally syntactic: types as structural invariants patterning the integration between processes operative in language understanding.

To express a noun in the benefactive case, for example, which can be seen as attributing to it a linguistic type consistent with being the target of a benefactive, is to capture the concept in a type-theoretic gloss. It tells us, I'm thinking about this thing in such a way that it *can* take a benefactive (the type formalism attempting to capture that "such a way"). A concept-to-type "map", as I just suggested, is mediated (in experience and practical reasoning) by cognitive organizations; when (social, embodied) enactions take linguistic form, these organizing principles can be encoded in how speakers apply morphosyntactic rules.

So the linguistic structures which I propose can be

airplane. The construction highlights the procedural and commercial dimension, not the physical mechanism and spatial path. But it does so *because* we know human flight is unnatural: we can poetically describe how the sky is filled with flying leaves or birds, but not "flying people", even if we are nearby an airport.

[26] We can express an intent to bake someone a cake, but not (well, maybe comedically) to *burn* someone a cake ("burn", at least in this context, implies something not intended); however, we *can* say "I burnt your cake", while it is a little jarring to say "I baked your cake" — the possessive implies that some specific cake is being talked about, and there is less apparent reason to focus on one particular stage of its preparation (the baking) once it is done. I *will* bake a cake, in the future, uses "bake" to mean also other steps in preparation (like "make"), while, in the present, "the cake *is* baking" emphasizes more its actual time in the oven. I *baked your cake* seems to focus (rather unexpectedly) on this specific stage even after it is completed, whereas *I baked you a cake*, which is worded as if the recipient did not know about the cake ahead of time, apparently uses "bake" in the broader sense of "made", not just "cooked in an oven". Words' senses mutate in relation to the kinds of situations where they are used — why else would *bake* mean "make"/"prepare" in the past or future tense but "cook"/"heat" in the present?

© 2021/05/24 Nathaniel Christen

formally modeled by a kind of type theory work communicatively as carriers and thereby signifiers of cognitive attitudes. The type is a vehicle for the concept because it takes part in constructions which express conceptual details — the details don't emerge merely by virtue of the type itself. I am not arguing for a neat concept-to-type correspondence; instead, a type system provides a "formal substrate" that models (with some abstraction and simplification) how properties of individual concepts translate (via cognitive-schematic intermediaries) to their manifestation in both semantics and syntax.

Continuing with declension as a case study, consider how an "ontology" of word senses can interrelate with the benefactive. A noun as a benefactive target most often is a person or some other sentient/animate being; an inanimate benefactive is most likely something artificial and constructed (cf., *I got the car new tires*). How readily hearers accept a sentence — and the path they take to construing its meaning so as to make it grammatically acceptable — involves interlocking morphological and type-related considerations; in the current example, the mixture of benefactive case and which noun "type" (assuming a basic division of nouns into e.g. animate/constructed/natural) forces a broader or narrower interpretation. A benefactive with an "artifact" noun, for example, almost forces the thing to be heard as somehow disrepaired:

- (264) I got glue for your daughter.
- (265) I got glue for your coffee mug.

We gather (in the second case) that the mug is broken — but this is never spelled out by any lexical choice; it is implied indirectly by using benefactive case. It is easy to design similar examples with other cases: a locative construction rarely targets "sentient" nouns, so in (reprising (3)-(5))

- (266) We're going to Grandma!
- (267) Let's go to him right now.
- (268) Let's go to the lawyers.
- (269) Let's go to the press.

we mentally substitute the person with the place where they live or work.

Morphosyntactic considerations are also at play: *to the lawyers* makes "go" sound more like "consult with", partly because of the definite article (*the* lawyers implies conversants have some prior involvement with specific lawyers or else are using the phrase metonymically, as in "go to court" or "to the courts", for legal institutions generally; either reading draws attention away from literal spatial implications of "go"). *Go to him* implies that "he" needs some kind of help, because if the speaker just meant going to wherever he's at, she probably would have said that instead.

Similarly, the locative in *to the press* forces the mind to reconfigure the landmark/trajector structure, where *going* is thought not as a literal spatial path and *press* not a literal destination — in other words, the phrase must be read as a metaphor. But the "metaphor" here is not "idiomatic" or removed from linguistic rules (based on mental resemblance, not language structure); here it clearly works off of formal language patterns: the landmark/trajector relation is read abstracted from literal spatial movement because the locative is applied to an expression (*the press*) which does not (simplistically) meet the expected interpretation as "designation of place". In short, there are two different levels of *granularity* where we can look for agreement requirements: a more fine-grained level where e.g. *locative* draws in a type-specification of a *place* or *location*; and a coarser level oriented toward Parts of Speech, and typologies of phrasal units. The former analysis addresses the level I have called "macrotypes", while the latter scale is at the "macrotype" level.

I envision the unfolding that I have just sketched out as something Phenomenological — it arises from a unified and subjective consciousness, one marked by embodied personal identity and social situation. If there are structural stases that can be found in this temporality of experience, these are not constitutive of conscious reality but a mesh of rationality that supports it, like the veins in a leaf. Stuctural configurations can be lifted from language insofar as it is a conscious, formally governed activity, and lifted from the ambient situations which lend language context and meaning intents. So any analytic emphasis on structural fixpoints threaded through the lived temporality of consciousness is an abstraction, but one that is deliberate and necessary if we want to make scientific or in any other manner disputatable claims about how language and congition works.

© 2021/05/24 Nathaniel Christen

Figure 3: Dependency-style graph with argument repetition

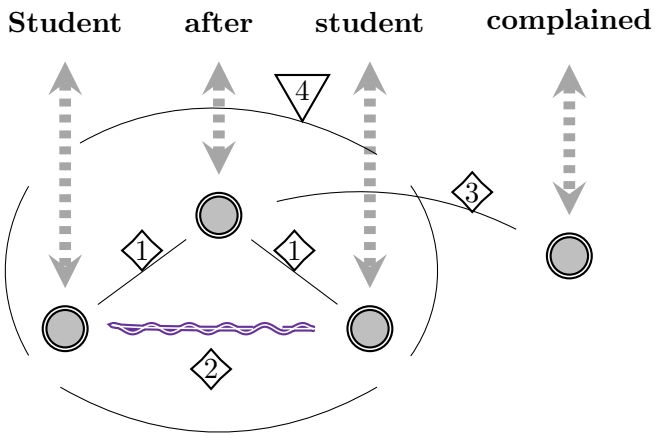

- ♦ 1 Head/dependent relation
- ♦ 2 Argument repetition (——`after`—— idiom)
- ♦ 3 Propositional completion ($V :: N \twoheadrightarrow \mathcal{P}rop$)
- ▼ 4 Phrase (modeled as applicative structure), typed as N^{+}

To return to the example of *Student after student*, I commented that designating one word to "represent" the phrase seemed arbitrary. If we consider functional-typing alone, *after* is the only non-noun, the natural conclusion is that "after" should be typed $N \twoheadrightarrow N \twoheadrightarrow N$ (which implies that "after" is analogous to the "functional" position, and in a lambda-calculus style reconstruction would be considered the "head" — Figure 3 is an example of how the sentence could be annotated, for sake of discussion). This particular idiom depends however on the two constituent nouns being the same word (a pattern I've also alluded to with idioms like *time after time*). Technically, this appears to be an example of *dependent types*: specifically, a type-theoretic model of *after* in this sense would seem to require that the $N \twoheadrightarrow N \twoheadrightarrow N$ signature be constrained so that the second noun matches the first — so the second N type is actually constrained to be a singleton type dependent on the first N's value. By way of illustration, Figure 3 shows a destructuring with implicit type annotations. Earlier I opined that the repetition in ——`after`—— is a rhetorical effect emphasizing that the recurrence of some phenomenon is tedious or noteworthy — in this one usage *after* becomes a pluralizing adjective that takes two arguments, but requires them to be the same. This sense of *after* can be given a function-like POS notation as, say, `after` :: $N^{\odot} \twoheadrightarrow N^{\odot} \twoheadrightarrow N^{+}$ (using $N^{\odot}$ and N^{+} to mean singular and count-plural nouns, respectively), but with a further stipulation that the second argument duplicates the first. Formally, then, the parameters for *after* are a dependent type pair ([10], [70]) satisfied by an identity comparison between the two nouns. This analysis captures a type-theoretic gloss on the structural contrast between *Student after student* and *Many students*, phrases which are similar but not identical in meaning (so whose differences need explaining).

I have offered a more cognitive account focusing on the implicit temporality of *Student after student*; this later type-oriented model is more formal, or at least leaves open the possibility that language is organically taking in structures engineered into artificial (e.g., computer programming) languages. It is certainly possible to witness formalizable structures in language patterns — Zhaohui Luo finds strong evidence for dependent types being a good model for semantic norms in [47], for example. Whether these kinds of formalisms have important causal influence on language acquiring its evident patterns, or are more like just convenient representational tools, is perhaps an open (and maybe case-by-case) question.

Consider alternatives for "many students". The phrase as written suggests a type signature (with "many" as the

© 2021/05/24 Nathaniel Christen

"function-like" or derivative type) $N^{\dagger} \dashrightarrow N^{\dagger}$, yielding a syntactic interpretation of the phrase; this interpretation also suggests a semantic progression, an accretion of intended detail. From *students* to *many students* is a conversion between two plural nouns (at the level of concepts and semantic roles); but it also implies relative size, so it implies some *other* plural, some still larger group of students from which "many" are selected. While rather abstract and formal, the $N^{\dagger} \dashrightarrow N^{\dagger}$ representation points toward a more cognitive grounding which considers this "function" as a form of thought-operation; a refinement of a situational model, descriptive resolution, and so forth. If we are prepared to accept a cognitive underpinning to semantic classification, we can make the intuition of Part of Speech signatures as "functions" more concrete: in response to what "many" (for example) is a function *of*, we can say a function of propositional attitude, cognitive schema, or attentional focus.

The schema which usefully captures the sense and picture of *students* is distinct (but arguably a variation on) that for *many students*, and there is a "mental operation" triggered by the *many students* construction which "maps" the first to the second. Similarly, *student after student* triggers a "scheme evolution" which involves a more explicit temporal unfolding (in contrast to how *many students* instead involves a more explicit quantitative *many/all* relation). What these examples show is that associating parts of speech with type signatures is not just a formal fiat, which "works" representationally but does not necessarily capture deeper patterns of meaning. Instead, I would argue, type signatures and their resonance into linkage acceptability structures (such as singular/plural and mass/count agreement) *point toward* the effects of cognitive schema on what we consider meaningful.

In *Student after student came out against the proposal*, to *come out*, for/against, lies in the semantic frame of attitude and expression (it requires a mental agent, for example), but, as per my prior analysis, its reception carries a trace of spatial form: to come out *to* a public place, to go on record with an opinion. Usually "come out [for/against]", in the context of a policy or idea, is similarly metaphorical. But the concrete spatial interpretation remains latent, as a kind of residue on even this abstract rendition, and the spatial undercurrent is poised to emerge as more literal, should the context warrant. However literally or metaphorically the "space" of the "coming out" is understood, however explicit or latent its cogitative figuration, is not something internal to the language; it is a potentiality which will present in different ways in different circumstances. This is not to say that it is something apart from linguistic meaning, but it shows how linguistic meaning lies neither in abstract structure alone, nor contextual pragmatics, but in their cross-reference.

References

1 Benjamin Adams and Martin Raubal, "A Metric Conceptual Space Algebra". https://pdfs.semanticscholar.org/521a/cbab9658df27acd9f40bba2b9445f75d681c.pdf

2 Benjamin Adams and Martin Raubal, "Conceptual Space Markup Language (CSML): Towards the Cognitive Semantic Web". http://idwebhost-202-147.ethz.ch/Publications/RefConferences/ICSC_2009_AdamsRaubal_Camera-FINAL.pdf

3 Alan Akbik and Jürgen Broß, "Wanderlust: Extracting Semantic Relations from Natural Language Text Using Dependency Grammar Patterns" http://ceur-ws.org/Vol-491/semse2009_7.pdf

4 Jens Allwood, "Semantics as Meaning Determination with Semantic-Epistemic Operations", in Jens Allwood and Peter Gärdenfors, eds., *Cognitive Semantics: Meaning and Cognition*, John Benjamis 199, pages 12-28. https://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.445.352&rep=rep1&type=pdf

5 Nicholas Asher and James Pustejovsky, "A Type Composition Logic for Generative Lexicon" https://www.cs.brandeis.edu/~jamesp/classes/cs216-2009/readings2009/TCLforGL.pdf

6 Chris Barker and Chung-Chieh Shan, "Types as Graphs: Continuations in Type Logical Grammar" https://www.nyu.edu/projects/barker/barker-shan-types-as-graphs.pdf

7 Lucas Bechberger and Kai-Uwe Kühnberger, "Measuring Relations between Concepts in Conceptual Spaces" https://arxiv.org/pdf/1707.02292.pdf

8 Radim Bělohlávek and Vladimı́r Sklenář, "Formal Concept Analysis Constrained by Attribute-Dependency Formulas" B. Ganter and R. Godin, eds.,: ICFCA 2005, LNCS 3403, pp. 176-191, Berlin, Springer-Verlag, 2005. http://belohlavek.inf.upol.cz/publications/BeSk_Fcacadf.pdf

9 Juri Benovsky, "Two Concepts of Possible Worlds { or Only One?". THEORIA, 2008, volume 74, pp. 318{330. http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.464.3904&rep=rep1&type=pdf

10 Jean-Philippe Bernardy, *et. al.*, "Parametricity and Dependent Types". http://www.staff.city.ac.uk/~ross/papers/pts.pdf

11 Ismaïl Biskri and Jean-Pierre Descles, "Applicative and Combinatory Categorial Grammar (from syntax to functional semantics)" https://www.researchgate.net/publication/232754402_Applicative_and_Combinatory_Categorial_Grammar_from_syntax_to_functional_semantics

12 Thomas Bittner, Barry Smith, and Maureen Donnelly, "The logic of systems of granular partitions." http://ontology.buffalo.edu/smith/articles/BittnerSmithDonnelly.pdf

© 2021/05/24 Nathaniel Christen

© 2021/05/24 Nathaniel Christen

13 Mathieu Bouchard, "A Type Theory for the Documentation of PureData" http://artengine.ca/~catalogue-pd/43-Bouchard.pdf

14 Line Brandt, "The Communicative Mind." Newcastle-upon-Tyne: Cambridge Scholars Publishing, 2013

15 Per Aage Brandt, "Spaces, Domains, and Meaning: Essays in Cognitive Semiotics". Newcastle-upon-Tyne: Cambridge Scholars Publishing, 2013 http://semiotics.au.dk/fileadmin/Semiotics/pdf/per-aage-brandt/Per_Aage_Brandt__Spaces__Domains__and_Meaning._Essays_in_Co.pdf

16 Stergios Chatzikyriakidis and Zhaohui Luo, "Individuation Criteria, Dot-types and Copredication: A View from Modern Type Theories." *Association for Computational Linguistics, Proceedings of the 14th Meeting on the Mathematics of Language (MoL 14)*, pp. 39{50, 2015. http://www.aclweb.org/anthology/W15-2304

17 Do Kook Choe and Eugene Charniak, "Parsing as Language Modeling". https://aclweb.org/anthology/D16-1257

18 Bob Coecke, *et. al.*, "Interacting Conceptual Spaces I: Grammatical Composition of Concepts". https://arxiv.org/pdf/1703.08314.pdf

19 Ralph Debusmann, "An Introduction to Dependency Grammar". Hauptseminar *Dependenzgrammatik*, SoSe 99. https://www.ps.uni-saarland.de/~rade/papers/dg.pdf

20 Ralph Debusmann, "Extensible Dependency Grammar: A Modular Grammar Formalism Based On Multigraph Description". Saarlandes University, dissertation 2006.

21 Ralph Debusmann, Denys Duchier, and Andreas Rossberg, "Modular Grammar Design with Typed Parametric Principles". James Rogers, ed., *Formal Grammar/Mathematics 2005*. CSLI Publications, 2009. http://web.stanford.edu/group/cslipublications/cslipublications/FG/2005/debusmann.pdf

22 Jean-Pierre Desclés, "Reasoning in Natural Language in Using Combinatory Logic and Topology: An Example with Aspect and Temporal Relations." https://www.aaai.org/ocs/index.php/FLAIRS/2010/paper/viewFile/1350/1736

23 Kit Fine, "Part-whole", in Barry Smith and David Woodruff Smith, *The Cambridge Companion to Husserl*, pp. 463-486. Cambridge University Press, New York, 1995.

24 Gilled Fauconnier, "Mental Spaces: Aspects of Meaning Construction in Natural Language". Cambridge, 1994

25 Mirjam Fried and Jan-Ola Östman, "Construction Grammar: A thumbnail sketch" https://www.researchgate.net/publication/280291354_Construction_Grammar_A_thumbnail_sketch

26 Peter Gärdenfors, "Does Semantics Need Reality?". Riegler A., Peschl M., von Stein A. (eds), *Understanding Representation in the Cognitive Sciences.*, pages 209-217. Springer, Boston, MA.

27 Peter Gärdenfors and Frank Zenker, "Theory Change as Dimensional Change: Conceptual Spaces Applied to the Dynamics of Empirical Theories". *Synthese 190(6)*, pp. 1039-1058, 2013. http://lup.lub.lu.se/record/1775234

28 Ben Goertzel, "Probabilistic Language Networks: Integrating Word Grammar and Link Grammar in the Framework of Probabilistic Logic". http://goertzel.org/ProwlGrammar.pdf

29 Kenneth Holmqvist, "Conceptual Engineering: Implementing cognitive semantics", in Jens Allwood and Peter Gärdenfors, eds., *Cognitive Semantics*, pp 153 - 171, Amsterdam, Philadelphia: John Benjamins, 1999.

30 Kenneth Holmqvist, "Implementing Cognitive Semantics: Image schemata, valence accommodation, and valence suggestion for AI and computational linguistics". PhD thesis, Dept. of Cognitive Science, Lund University, Lund, Sweden, 1993.

31 Juneki Hong and Jason Eisner, "Deriving Multi-Headed Planar Dependency Parses from Link Grammar Parses" https://www.cs.jhu.edu/~jason/papers/hong+eisner.tlt14.paper.pdf

32 Laurence R. Horn and Gregory Ward, "Handbook of Pragmatics" (Blackwell Handbooks in Linguistics). Wiley, 2004. https://www.felsemiotica.com/descargas/Horn-Laurence-R.-and-Ward-Gregory-Ed.-The-Handbook-of-Pragmatics.pdf

33 Mark Johnson, "The Body in the Mind: The Bodily Basis of Meaning, Imagination, and Reason". University of Chicago Press, 1990

34 Michael Jubien, "Ontology, Modality, and the Fallacy of Reference". Cambridge University Press, New York, 1993.

35 Carlos Montemayor, *et. al.*, "Implementation, Formalization and Representation: Challenges for the Integrated Information Theory". http://online.sfsu.edu/barros/publications/publications/files/MontemayorEtAl2019.pdf

36 Stephan Kepser and Uwe Mönnich, "(Un-)Decidability Results for Head-Driven Phrase Structure Grammar". Algebraic Methods in Language Processing AMiLP-3, Anton Nijholt and Giuseppe Scollo, eds., 2003. http://www.tcl-sfs.uni-tuebingen.de/kepser/papers/amilp3-hpsg.pdf

37 Lingpeng Kong, Alexander M. Rush, and Noah A. Smith, "Transforming Dependencies into Phrase Structures". http://www.aclweb.org/anthology/H01-1014

38 Yusuke Kubota and Robert Levine, "The syntax-semantics interface of 'respective' predication: A unified analysis in Hybrid Type-Logical Categorial Grammar" https://www.asc.ohio-state.edu/levine.1/publications/kl-resp.pdf

39 George Lakoff and Mark Johnson, "Philosophy in the Flesh: the Embodied Mind and its Challenge to Western Thought". New York: Basic Books, 1999.

40 Ronald Langacker, "Foundations of Cognitive Grammar, vol. 1". Stanford University Press, 1991

41 Ronald Langacker, "Interactive Cognition: Toward a Unified Account of Structure, Processing, and Discourse". *International Journal of Cognitive Linguistics*, Volume 3, Number 2, 2014. http://lchc.ucsd.edu/MCA/Mail/xmcamail.2014-08.dir/pdf8jEY9UCVVh.pdf

42 Ronald Langacker, "Investigations in Cognitive Grammar". De Gruyter, 2010. http://lchc.ucsd.edu/MCA/Mail/xmcamail.2014-08.dir/pdf8jEY9UCVVh.pdf

43 Martha Lewis and Jonathan Lawry, "Uncertainty Modelling for Vague Concepts: A prototype theory approach". https://www.sciencedirect.com/science/article/pii/S0004370209000903

44 Martha Lewis and Jonathan Lawry, "Hierarchical Conceptual Spaces for Concept Combination". https://www.sciencedirect.com/science/article/pii/S0004370216300492

45 Zhaohui Luo, "Type-Theoretical Semantics with Coercive Subtyping". https://www.cs.rhul.ac.uk/home/zhaohui/ESSLLI11notes.pdf

46 Zhaohui Luo, "Using Signatures in Type Theory to Represent Situations". https://www.researchgate.net/publication/268079019_Using_Signatures_in_Type_Theory_to_Represent_Situations

47 Zhaohui Luo and Sergei Soloviev, "Dependent Coercions". https://www.sciencedirect.com/science/article/pii/S1571066105803147

© 2021/05/24 Nathaniel Christen

48 Bruno Mery, Richard Moot, and Christian Retoré, "Plurals: Individuals and sets in a richly typed semantics". http://arxiv.org/pdf/1401.0660.pdf (archived) 3 Jan 2014

49 Laura A. Michaelis and Knud Lambrech, "Toward a Construction-Based Theory of Language Function: The Case of Nominal Extraposition" http://www.jstor.org/stable/416650

50 Erwan Moreau, "From link grammars to categorial grammars" https://hal.archives-ouvertes.fr/hal-00487053/document

51 Timothy Osborne and Daniel Maxwell, "A Historical Overview of the Status of Function Words in Dependency Grammar" https://www.aclweb.org/anthology/W15-2127

52 Joakim Nivre, "Dependency Grammar and Dependency Parsing." http://stp.lingfil.uu.se/~nivre/docs/05133.pdf

53 Esther Pascual, "Fictive Interaction: The conversation frame in thought, language, and discourse". Philadelphia, John Benjamins, 2014.

54 Efe Peker, "Following 9/11: George W. Bush's Discursive Re-Articulation of American Social Identity" Master's Dissertation, Linköping University, Sweden http://www.diva-portal.org/smash/get/diva2:21407/FULLTEXT01.pdf

55 Jean Petitot, "Syntax Topologique et Grammaire Cognitive". *Langages 25.103*, pp. 97-128, 1991.

56 Jean Petitot, "The morphodynamical turn of cognitive linguistics". https://journals.openedition.org/signata/549

57 Steven Pinker, "The Stuff of Thought: Language As a Window Into Human Nature". Penguin, 2007.

58 Joey Pollack, "Holism, Conceptual Role, and Conceptual Similarity". https://www.tandfonline.com/doi/abs/10.1080/09515089.2020.1729973?journalCode=cphp20

59 James Pustejovsky, "A Survey of Dot Objects" Unpublished manuscript, 2006.

60 Maxwell James Ramstead, "Transcendental Phenomenology and Naturalistic Epistemology." https://archipel.uqam.ca/8209/1/M14029.pdf

61 Siva Reddy, *et. al*, "Transforming Dependency Structures to Logical Forms for Semantic Parsing." https://aclweb.org/anthology/Q16-1010

62 Terry Regier, *The Human Semantic Potential: Spatial Language and Constrained Connectionism*. Cambridge, MA: MIT Press, 2014.

63 Bradley Rives, "The Empirical Case Against Analyticity: Two Options for Concept Pragmatists ". https://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.188.9556&rep=rep1&type=pdf

64 Aurélie Rossi, "Applicative and Combinatory Categorial Grammar: Analysis of the French Interrogative Sentences". *FLAIRS Conference, Association for the Advancement of Artificial Intelligence*, 2008. https://www.aaai.org/Papers/FLAIRS/2008/FLAIRS08-118.pdf

65 Dan Ryder, "Problems of representation II: Naturalizing content". In *The Routledge Companion to Philosophy of Psychology*, Routledge 2009, pages 251-279. https://www.semanticscholar.org/paper/Problems-of-representation-II%3A-naturalizing-content-Ryder/1073bf288423e3e9b1940628cb2deb7db8b8fae2

66 Gerold Schneider, "A Linguistic Comparison of Constituency, Dependency and Link Grammar". Zurich University, diploma, 2008. https://files.ifi.uzh.ch/cl/gschneid/papers/FINALSgeroldschneider-latl.pdf

67 Matt Selway, *et. al*, "Configuring Domain Knowledge for Natural Language Understanding" http://ceur-ws.org/Vol-1128/paper9.pdf

68 Daniel D. Sleator and Davy Tamperley, "Parsing English with a Link Grammar". https://www.link.cs.cmu.edu/link/ftp-site/link-grammar/LG-IWPT93.pdf

69 Kiyoshi Sudo, *et. al.*, "An Improved Extraction Pattern Representation Model for Automatic IE Pattern Acquisition" https://www.aclweb.org/anthology/P03-1029

70 Ribeka Tanaka, *et. al.*, "Factivity and Presupposition in Dependent Type Semantics". http://www.lirmm.fr/tytles/Articles/Tanaka.pdf

71 Miriam Taverniers, "Subjecthood and the notion of instantiation". https://semanticsarchive.net/Archive/DU0OGRk0/Taverniers-2005-Subjecthood-PREPRINT.pdf

72 Roman Urban and Megdalena Grzelinska, "A Potential Theory Approach to an Algorithm of Conceptual Space Partitioning". In *Cognitive Studies*, 2017, volime 1. https://ispan.waw.pl/journals/index.php/cs-ec/article/download/cs.1310/3065

73 Orlin Vakarelov, "Pre-cognitive Semantic Information". https://link.springer.com/article/10.1007/s12130-010-9109-5

74 Orlin Vakarelov, "The cognitive agent: Overcoming informational limits" https://philarchive.org/archive/VAKTCAv1

75 Jan van Eijck, "Dynamic Reasoning Without Variables" https://www.semanticscholar.org/paper/Dynamic-Reasoning-Without-EijckCWI/4fb3b23f81b8691fbdaf8d3eb24c086c5f826954

76 Rik van Noord, *et. al.*, "Exploring Neural Methods for Parsing Discourse Representation Structures" https://www.aclweb.org/anthology/Q18-1043

77 Anne Vilnat, "Dialogue et analyse de phrases" Habilitation Mémoire, Université Paris-Sud https://perso.limsi.fr/anne/HDR/MemoireHDR.pdf

78 Olav K. Wiegand, "A Formalism Supplementing Cognitive Semantics Based on a New Approach to Mereology". Ingvar Johansson, Bertin Klein and Thomas Roth-Berghofer, eds., *Contributions to the Third International Workshop on Philosophy and Informatics*, 2006.

79 ——— "On referring to Gestalts". Mirja Hartimo, ed., *Phenomenology and Mathematics*, pp. 183-211. Springer, 2010.

80 Rudolf Wille, "Conceptual Graphs and Formal Concept Analysis". *Conceptual Structures: Fulfilling Peirce's Dream Fifth International Conference on Conceptual Structures*, 1997.

81 David Woodruff Smith, "Mind World". Cambridge University Press, 2004.

82 Fei Xia and Martha Palmer, "Converting Dependency Structures to Phrase Structures". http://www.aclweb.org/anthology/H01-1014

83 Yiyu Yao, "A Comparative Study of Formal Concept Analysis and Rough Set Theory in Data Analysis". http://www2.cs.uregina.ca/~yyao/PAPERS/Rough_concept.pdf

84 Jordan Zlatev, "The Dependence of Language on Consciousness". http://www.mrtc.mdh.se/~gdc01/work/ARTICLES/2014/4-IACAP%202014/IACAP14-GDC/pdf/Language-Consciousness-Zlatev.pdf