Ruth Kempson,
King's College London
ruth.kempson@kcl.ac.uk
Eleni Gregoromichelaki
King's College London
eleni.gregor@kcl.ac.uk
Ronnie Cann
University of Edinburgh
r.cann@ed.ac.uk

Context and Compositionality: the Challenge of Conversational Dialogue

Abstract: In this paper we present an innovative linguistic framework, Dynamic Syntax, in which natural-language syntax is defined as procedures for context-dependent interpretation. The framework provides a formalism where both representations of content and context are defined dynamically and structurally, with time-linear monotonic growth across sequences of partial trees as the core structure-inducing notion. Application of this framework to the resolution of elliptical phenomena, both inter- and intra-sententially provides an integrated account of ellipsis construal. In addition, this intrinsically dynamic perspective extends naturally to the modelling of dialogue exchanges with free shifting of role between speaking and hearing (*split-utterances*). We shall argue in closing that the success in capturing dialogue patterns of ellipsis within an overall account of ellipsis leads to a novel composite concept of compositionality, and reconsideration of the relation between the language capacity and more general cognitive capacities.

Key words: dialogue, ellipsis, split utterances, compositionality, domain-specific

1. A syntax for dialogue modelling?

In this paper, we shall be promoting a concept of grammar as a set of mechanisms for proposition construction, and using it to model the interactions that take place in conversation. The very first questions to ask is: why such a grammar, and why conversational dialogue? The urgency of readdressing the question of what assumptions to make about

grammar is that, despite all the work done by semanticists and pragmatists in emphasising just how important context is to language understanding, we still have no clear understanding of what an appropriate notion of context consists in; and current sentence-based grammars fail altogether to address the challenge of explaining the systemic context-dependence of natural language (NL). Indeed, all conventional grammars side-step this challenge so that any unitary characterisation of context-dependent phenomena is in principle precluded.

The problem is this. We know now that it is diagnostic of an NL that all aspects of explaining meaning require context-dependent forms of explanation. Nevertheless, our concept of knowledge of language remains stubbornly sentence-based and static. In virtue of the established competence-performance distinction, nothing within the grammar itself makes any reference to concepts taken to pertain to performance; and all aspects of NL that relate to processing are thus excluded from all aspects of competence. Such relegation includes any reference to time-linearity or to the incrementality of processing: all phenomena that might indicate the need of such reference have to be captured in different (hierarchical) terms as expressible within the boundaries set by the sentence remit. One consequence of this is that any aspect of language which displays dependencies that can be established both within and across sentence-boundaries can only be captured as two separate phenomena: one internal to the sentence taken as "grammar-internal", and one that is across sentence boundaries seen as "grammar-external" and, therefore, a distinct, discourse-based phenomenon. This scenario arises both with anaphora and ellipsis, for these systematically display both structural, and so, by assumption, sentence-internal properties, and yet nevertheless discourse cross-sentential properties. But, as it turns out, this bifurcation is not a problem specific to anaphora and ellipsis and so a local problem needing only a structure-particular form of solution: the problem arises with all context-dependent phenomena like tense construal, domain restriction, etc. Furthermore, in addressing ellipsis, having separated grammar-internal ellipses from discourse ellipses, the argumentation is restricted to whether the generalisations expressible within the remit of grammar should be captured in semantic or syntactic terms: neither

syntactic nor semantic accounts make any reference to a general concept of context. Analyses either mention some devices for readjustments of indexings, with stipulations for distinct types of indexing apparently required, or of a range of operations of lambda abstraction and binding. Whatever parallelism there might be between anaphora and ellipsis in virtue of both being exemplars of the general phenomenon of context-dependence is simply not expressible. To make matters worse, neither syntactic nor semantic accounts are fully successful and so-called pragmatic ellipsis is invoked where these fail (Stainton 2006); but this additional phenomenon of pragmatic ellipsis is taken by all parties to this three-fold debate to be peripheral and not as undermining supposedly grammar-internal explications of ellipsis phenomena. The result is that there is a general consensus that ellipsis is "fractally heterogeneous" (Ginzburg and Cooper 2004; Merchant 2009); and the observation that ellipsis occurs when the intended construal can be recovered directly from context, far from being taken as the point of departure for articulating some requisite concept of context, is set aside as no more than a folk intuition, not expressible within our linguistic theories.

In this paper we address this impasse and take up the challenge of reflecting the folk intuition directly. We shall argue that, respecting this intuition, a unitary characterisation of ellipsis and indeed of context itself are both possible, despite the observable diversities. However, such an achievement will not take place without having to abandon a number of familiar assumptions. Syntax will be defined as constraints on the incremental process of building up propositions, with the concept of incrementality in real time that is diagnostic of language processing reflected directly in the grammar system. The concepts of competence and performance will then have to be reconsidered, and the boundary between them redrawn. The grammar system will not be encapsulated, as syntactic constraints will interact freely with general cognitive constraints to determine interpretation. The system will not be domain-specific either, since the terms in which the concept of structural growth will be expressed will be that of growth of concepts in a system of conceptual representations, in principle shared with other cognitive activities such as vision and hearing. The result will be dynamic, with

evolving concepts of structured content and context in which procedures for language processing take centre stage. Within this shifted perspective, a unified story of context-dependency will emerge, and an integrated account of ellipsis. Furthermore the dialogue ellipsis effects which are so problematic for conventional NL grammar assumptions will emerge as nothing more than a sub-part of the overall explanation.

2. Dialogue modelling: interactive structure-building

In free-running conversational data, utterances are often apparently collaboratively constructed so that the contributions made by individual contributors to the dialogue are highly elliptical, relying on the context in which the conversation takes place for their interpretation:

(1) Context: Friends of the Earth club meeting
A: So what is that? Is that er... booklet or something?
B: It's a book
C: Book
B: Just ... talking about al you know alternative
D: On erm... renewable yeah
B: energy really I think.....
A: Yeah [British National Corpus: D97]

It might be tempting to take this type of data as merely illustrating performance dysfluencies. However, this would not be a good move, as 20% of our conversations may be made up of such data (as a recent corpus account established, see Purver et al 2009); and it is these data which are the sole input for the language learning child. Such free extension of partial structures started by someone else may be as a way of helping some other party finish their utterance, guessing in so doing what they might have in mind; but this condition is by no means necessary, and such exchanges freely occur, whether or not the parties to the activity are actually agreeing with each other over either the result or the form of words:

(2) A: Well I do know last week thet=uh Al was certainly very …
B: pissed off [Lerner96]
(3) (A and B arguing:)
A: In fact what this shows is
B: that you are completely wrong.

The very first difficulty is to decide where utterance boundaries occur, given that contributors to a conversation can exchange roles well within what the grammar determines is a sentence boundary:

(4) Conversation from A and B, to C:
A: We're going to ...
B: Bristol, where Jo lives.
(5) A: I just returned
B: from ...
A: Finland.

The problem for linguistic explanation is that such sub-sentential switches (what we call *split utterances*) with speaker/hearer exchange of roles can take place at any point, and across all syntactic dependencies (Purver et al 2011):

(6) A: I'm afraid I burned the buns.
B: Did *you* burn
A: *myself*? No, fortunately not.
(7) A: *D'you* know whether *every waitress* handed in
B: *her* tax forms?
A: or even *any* payslips?
(8) Gardener: I shall need the mattock.
Home-owner: *The*...
Gardener: *mattock*. For breaking up clods of earth.
[British National Corpus]

(6) involves a split between a reflexive pronoun and its antecedent; (7) involves a split between a quantifying expression and some pronoun that it binds, and then across a disjunction and another shift of speakers

to a polarity item dependent on that initially presented quantifier and the affective questioning element indicated by the first part. (8) involves a split between determiner and noun. The result is that unless such data are included within the remit of our grammar formalisms, then no phenomenon at all will get complete coverage. And confirmation that such data should be addressed by linguists is that children interact in such activities from the earliest stages of language acquisition in quite as systematic a way as adults, as the nursery rhyme game vividly displays:

(9) Carer: Old McDonald had a farm... On that farm he had a
Child: cow.

Yet it is far from obvious how to address this phenomenon, given orthodox assumptions. The output of the grammar is standardly a set of structures inhabited by complete sentences, supposedly as input to some performance theory to be articulated. But none of these fragments will be included in the set of wellformed expressions. And the possibility that such data might be explicable within denotational semantics, by way of alternative, is a non-starter, as the denotational semantics favoured by formal semanticists is externalist, unrelated to all cognitive considerations. Furthermore, the concept of context that is needed is very much richer than any denotational concept of context: context in these situations involves incremental structural update, with structure derived from arbitrary sentence parts able to function as context for a subsequent elliptical fragment.

These data pose problems for specific proposals in influential accounts of ellipsis construal (e.g. Dalrymple et al 1991), over and above the issue of how to capture the intuitive notion of a shared utterance. The problem is that interruptions are possible at any point, and in some cases so early that no intended propositional content is as yet fixable:

(10) A: They X-rayed me, and took a urine sample, took a blood sample. Er, the doctor
B: Chorlton?
A: Chorlton, mhm, he examined me.....
[British National Corpus]

In (10), for example, an interpretation is assigned to the fragment Chorlton well before the proposition in which it appears has been constructed, so its understanding can hardly be said to be in virtue of recognising the speaker's intended propositional content. Indeed any account of the construal of the fragment in (10) as involving abstraction over some previous propositional content, as proposed by Dalrymple et al and similar accounts, would yield quite the wrong interpretation with the doctor interpreted as having done the x-rays and taken samples of blood and urine.

There is the further problem that the intentions of the parties to the dialogue may only emerge/develop during the exchange, and so cannot be intrinsic to all processes of communicative understanding:

(11) A: Oh. They don't mean us to be friends, you see. So if we want to be.
B: which we do
A: then we must keep it a secret. [natural data]

(12) (A mother, B son)
A: This afternoon first you'll do your homework, then wash the dishes and then
B: you'll give me £10?

Not only this. Utterances may be multi-functional, so that more than one speech act can be expressed in one and the same fragment:

(13) A: Are you left or
B: Right-handed

(14) M: It's generated with a handle and
J: Wound round?
M: Yes, wind them round and this should, should generate a charge [British National Corpus]

In short, the phenomenon of context-dependence as posed by the data of conversational dialogue is highly problematic for all orthodox assumptions of syntax, semantics and pragmatics. In this paper we turn to a novel framework, that of *Dynamic Syntax*, to explore the extent to

which its intrinsically dynamic perspective gets a better handle on these data, and in particular the *split-utterance* phenomenon.

3. Dynamic Syntax: addressing the context challenge

Dynamic Syntax (DS) is an avowedly representationalist model of interpretation of which the core notion is interpretation growth. It is a model of how interpretations, represented as binary tree-structures of predicate-argument form, are built up relative to context, and individual steps in this building process reflect the incrementality with which hearers (and speakers) progressively build up interpretations for strings using information from context as it becomes available. Central to articulating this concept of growth are the attendant concepts of underspecification and its update, with a range of types of underspecification: not merely of content, but also of structure, and of structural relations. Indeed, this process of building up structure is taken to be what constitutes the syntax of NL grammar. With the dynamics of structural growth built into the core grammar formalism, NL syntax by definition is a set of principles for articulating growth of such structures: syntactic mechanisms, being meta to the representations themselves, are procedures that define how parts of representations of content can be incrementally introduced and updated. Furthermore, all procedures for structural growth are defined relative to context; and context is just as structural and dynamic as the concept of content with which it is twinned. Context, thus, is a record not merely of the (partial) structures built up, with the typed formulae that decorate them, but also of the procedures – actions – used in constructing them (see Cann et al 2007; also Section 4 below). The bonus of such explicit adoption of representationalist assumptions and the shift into a perspective in which the grammar reflects key properties of the dynamics of how language processing takes place is, as we shall see, that we have a natural basis for articulating a novel grammar-pragmatics interface that is fully commensurate with an integrated account of ellipsis and context-dependency.

3.1. The tree logic and tree-growth processes

The general process of parsing is taken to involve building as output a tree whose nodes reflect the content of some uttered sentence – in the simple case of a sentence uttered in isolation, a complete propositional formula. The input to this task, in such a simple case, is a tree that does nothing more than state at the root node the goal of the interpretation process to be achieved, namely, to establish some propositional formula. For example, in the parse of the string John upset Mary, the output tree to the right of the → in (15) constitutes some final end result: it is a tree in which the propositional formula itself annotates the top-node, and its various subterms appear on the dominated nodes in that tree rather like a proof tree in which all the nodes are labelled with a formula and a type (see Section 4.6 below). The input to that process is an initial one node tree (as in the tree representation to the left of the → in (15)) which simply states the goal as the requirement to a formula of appropriate propositional type (shown as ?Ty(t), the ? indicating that this is a goal not yet achieved, the S as final argument indicating an event term, of which more later):

(15) John upset Mary.

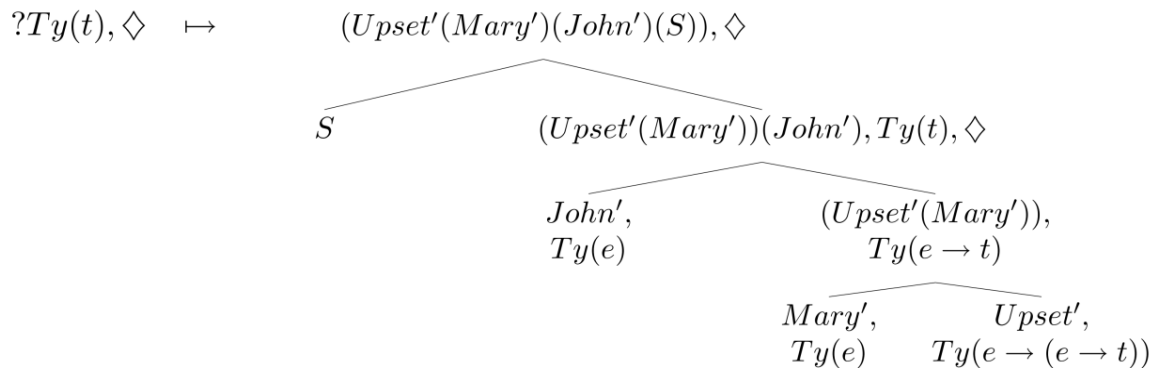

These DS trees are invariably binary, and, by convention, the argument always appears on the left branch, and the functor on the right branch (a pointer ◊ identifies the node under development). Each node in a complete tree is decorated not with words, but with terms of a logical language, these being subterms of the resulting propositional representation. The parsing task is to use both lexical input and

information from context to progressively enrich the input tree to yield such a complete output following general tree-growth actions.

In order to talk explicitly about how such structures grow, the trees need to be defined as formal objects; and DS adopts a (modal) logic of finite trees (LOFT: Blackburn and Meyer-Viol 1994).[1] The LOFT language makes available modalities not only for describing fixed tree relations, but also a basis for defining concepts of structural underspecification. Concepts of 'dominate' and 'be dominated by' are defined (using operators accompanied with the Kleene star), indicating some possible sequence of mother relations, or conversely a possible sequence of daughter relations; and these can be licensed even before there is some fixed number of such mother or daughter relations. For example, $\langle\uparrow^*\rangle Tn(a)$ is defined as a decoration on a node indicating that there must be at least one future development in which the node $Tn(a)$ bears a sequence of mother relations to the present node.[2] This structural underspecification is analogous to the more familiar underspecification displayed by anaphoric expressions, which are taken to project place-holding, meta-variable formula decorations, to be substituted by pragmatic substitution actions from context ensured by the second core concept of the framework: that of requirements for update. These are essential to get appropriate reflection of the time-linearity involved in building up trees in stages (partial trees). For every node, in every tree, all aspects of underspecification are twinned with a concept of requirement, ?X, for any annotation X on a node; and these are constraints on how the subsequent parsing steps must progress. Such requirements apply to all types of decoration, so that there may be type requirements, $?Ty(t)$, $?Ty(e)$, $?Ty(e \rightarrow t)$ etc; treenode requirements, $?\exists x(Tn(x))$, (associated with underspecified tree-relations ensuring that this undefined motherhood relation is determinatively fixed prior to completion of the emergent tree), and formula requirements $?\exists x Fo(x)$.

---

[1] There are two basic modalities, ways of describing node relations: $\langle\downarrow\rangle$ and $\langle\uparrow\rangle$. $\langle\downarrow\rangle\alpha$ holds at a node if α holds at its daughter, and the inverse, $\langle\uparrow\rangle\alpha$, holds at a node if α holds at its mother. There are also LINK ($\langle L\rangle$) relations between trees, with their inverse $\langle L^{-1}\rangle$.

[2] This is a standard tree-theoretic characterisation of dominate, used in LFG to express 'functional uncertainty' (see Dalrymple 2001 and references cited there).

This last form of requirement is associated with pronouns and other anaphoric expressions, which are lexically defined as projecting a place-holding metavariable with such a requirement ensuring that these are assigned a fixed term as value within the course of the derivation. Requirements thus drive the subsequent tree-construction process. Unless all requirements are eventually satisfied the parse will be unsuccessful.

Such structural underspecification and update can then be used to define core syntactic notions in a way that follows insights from parsing, and the time-linear dimension of processing in real time. In particular, the long-distance dependency effects which, since the late 1960's, have been taken by most to be diagnostic of a syntactic component independent of semantics are recast in terms of structural underspecification plus update. For example, when first processing the word *Mary* in (16) below, which is initially construed as providing a term whose role isn't yet identified, the parse is taken to involve the application of a computational action that introduces from the initial root node decorated with ?*Ty*(t), a relation to that top-node which is underspecified at this juncture, identifiable solely as dominated by the top-node, and requiring type *e*, i.e. with requirement ?*Ty*(e):

(16) Mary, John upset.

This enables the expression *Mary* to be taken to decorate this node: this is step (i) of (17).[3] Accompanying the under-specified tree relation is a requirement for a fixed treenode position: ?∃x.*Tn*(x). The update to this relatively weak tree relation becomes possible only after processing the subject plus verb sequence, which jointly yield the two-place predicate structure as in step (ii) of (17). The simultaneous provision of a formula decoration for this node and update of the unfixed node is provided in the unification step indicated there, an action which satisfies the update requirements of both nodes to be unified:

[3] *Tn*(0) decorating the top node of both partial trees is the tree-node identifier of the root node.

(17) Parsing: *Mary*, *John upset*:

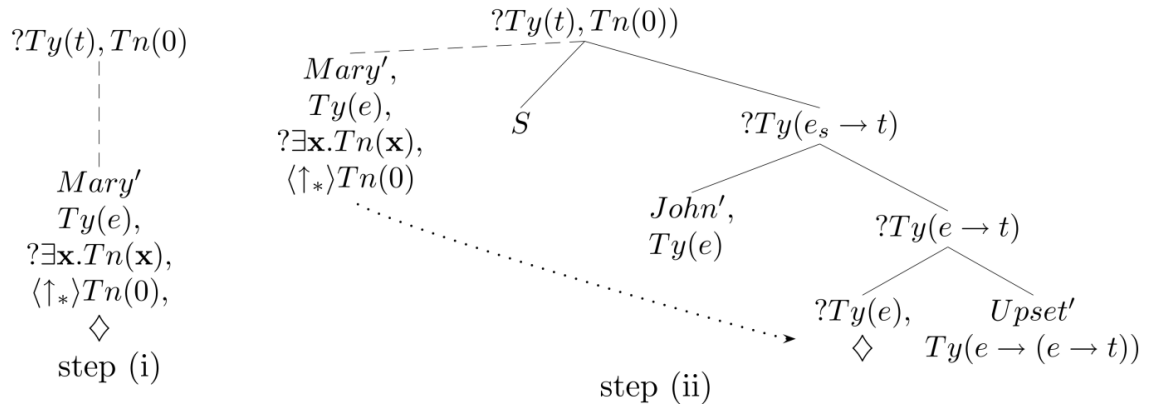

This process feeds into the ongoing development in which, once all terminal nodes are decorated, bottom-up application of labelled type deduction leads to the completed tree indicated in (15). Such an account of structural underspecification and update is not contentious as a parsing strategy: what is innovative is its application within the grammar-mechanism as the basic underpinning to syntactic generalisations.

This account might seem in principle skewed by focussing on parsing, but this is only superficial. Production also follows the very same processes, with but one further assumption – that at every step in production, there must be some richer tree, a so-called 'goal tree', which the tree under construction must subsume in the sense of being able to be developed into that goal tree by rules of the system. For the production of both (15) and (16), for example, each selected strategy for update has to be checked for subsumption with respect to the goal tree representing the content to be conveyed. These indeed share such a goal tree, illustrating how more than one sequence of strategies is licensed for any string-content pairing, both in parsing and production (to the advantage of real-time processing: Ferreira and Dell 2000). So parsers and producers alike use strategies for building up representations of content, either to establish interpretation for a sequence of words, or to find words which match the content to be conveyed.

To achieve the basis for characterising the full array of compound structures displayed in NL, DS defines in addition the license to build paired trees, so-called 'linked trees', linked together solely by the

sharing of terms, established, for example by encoded anaphoric devices such as relative-pronouns. Consider the structure derived by processing the string *John, who smokes, left* (omitting details of tense specification):

(18) Result of parsing *John*, who smokes, *left*:

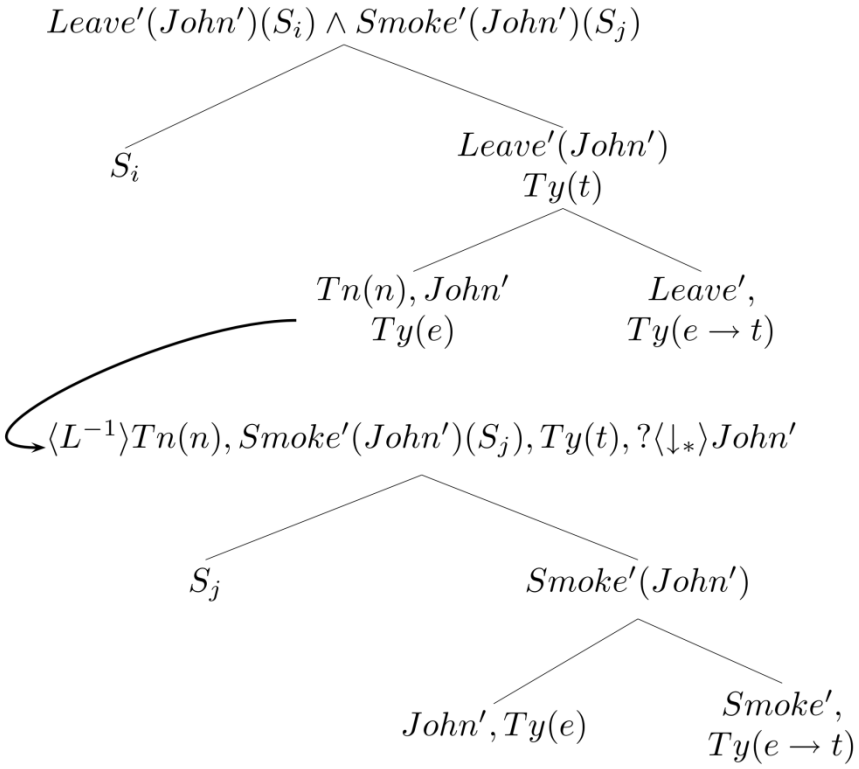

The arrow linking the two trees depicts the so-called LINK relation. The tree whose node is pointed by the arrow is the linked tree (read $<L^{-1}>$ as 'linked to'). Within any one such linked tree, the full range of computational, lexical and pragmatic actions remain available;[4] and with this flexibility to allow the incremental projection of arbitrarily rich compound structures, the result is a formal system combining lexical, structural and semantic specifications, all as constraints on the growth of trees. As argued in Kempson et al. (2001), Cann et al. (2005) and others, this leads to the comprehensive DS claim that the syntax of NLs does not involve a separate level of representation besides what is needed for

[4] The only pragmatic action formally defined in this framework is that of Substitution, presumed to apply in individual derivations to yield anaphora and ellipsis resolution.

semantics, not because there is no level of semantic representation, but because there is no independent level of syntactic representation. Despite the assumption that this progressive build up of a semantic representation is a basis for doing syntax, syntax in this model is not taken to include a level of representation where there is structure defined over a string of words. The trees in DS are not inhabited by words, but by the concepts words express in context, and there is no notion of linear ordering expressed on the tree. Furthermore, lexical specifications are defined in exactly the same terms of actions inducing tree growth, and these actions can take place only if the condition triggering these actions matches the decorations on the node at which the pointer has got to in the parse. So all structural restrictions are stated in terms of the interaction of constraints on tree growth.

A consequence of this methodology of incorporating the dynamics of incremental growth into the syntactic formalism itself is the way concepts of structural underspecification and subsequent update replace the need to postulate multiple levels of representation. The building of unfixed nodes and updating them replaces a multi-level account of syntax with progressive growth along a time line towards just one type of representation, hence a single representational level. The characterisation of lexical specifications in the same terms enables seamless integration of lexical and syntactic forms of generalisation, so that discrete vocabularies for lexical and syntactic generalisation are precluded. And constraints that, in other frameworks, are taken to be specific to NL syntax and not reducible to semantic generalisations are analysed as constraints on the same growth process. For example, the complex NP constraint, associated with a precluding of dependency of some expression outside a relative clause sequence with some site within that relative, is analysed in DS via the locality imposed by the licence to build linked-tree pairings. Any expression characterised as decorating an unfixed node, e.g. a relative pronoun,[5] has to be resolved within the tree which that unfixed node construction step initiates. Hence it cannot be

[5] Relative pronouns are lexically defined in English to induce a copy of the head at such an unfixed node (hence its position initiating the structure for relative clause construal).

resolved in some tree only linked to that tree, and the island constraint is captured, albeit in less familiar terms than is standard.

Such a system might appear to face the challenge of characterising quantification, often thought to constitute a second core case where the syntax of NLs is disjoint from what is required for the semantics of quantification. Notoriously, no NL overtly displays quantification following the pattern of predicate logic's propositional quantifying operators. But in this framework, this problem is addressed by grounding the account in the so-called *epsilon calculus* (Hilbert and Bernays 1939). This is a logic which provides the formal account of the so-called arbitrary names of natural deduction systems for predicate logic. The heart of such names is that their syntax is simple: they are a naming device like all other individual-denoting expressions of the logic.[6] It is the semantics for such names that is complex, for they are terms denoting witness sets for the entire proposition in which they occur; and this means that a rule of semantic evaluation is defined to determine, as output, that their internal structure reflects the environment in which they occur.[7] There is thus a concept of growth in this aspect too, in growth of the restrictor from what the incremental structural process provides (e.g. that projected by the nominal) and that of the predicate structure within which it is contained (see Kempson et al. 2001; Cann et al. 2005 for all details). Details aside, the bonus of this account in relation to multiplicity of levels of representation for NL grammar-writing is that another supposed dissymmetry between NL syntax and its

---

[6] As all students introduced to the natural deduction proof system for predicate logic swiftly find out, it is these arbitrary names which guarantee that the proof rules of predicate logic are notably easier to learn than the rules of semantic evaluation of predicate logic.

[7] The core of the calculus is its equivalence to predicate logic:

$$\frac{\exists x \varphi(x)}{\varphi(\varepsilon, x, \varphi(x))}$$

The epsilon term analogue of the existentially quantified formula by definition contains two occurrences of the predicate φ (the predicate which constitutes the open proposition that the existential quantifier binds), hence its reflection of its containing environment.

required semantics dissolves upon analysis, for the account presumes that these are terms of the same type *e* as all other argument expression.[8]

Overall then, the system involves but a single level of representation, the need for multiple levels replaced by the concept of growth of partial representations, these representations themselves being part of a denotationally interpretable system, the lambda calculus. In particular, the apparent multiplicity presented by all grammar formalisms which posit independent, statically defined, syntax and semantics is resolved through articulation of the dynamics of how just one type of representation is incrementally built up. An immediate consequence is that the system is un-encapsulated. In any one application to yield some derivation, application of general computational actions interact with lexical actions as driven by parsing the sequence of words. Crucially, actions of tree-building induced by grammatical and lexical processes may be interspersed with pragmatic actions of substitution and enrichment as driven by requirements associated with underspecified input. With wellformedness of a string of words defined as the availability of at least one possible sequence of actions from the initial goal to some completed propositional output with no requirements outstanding by incremental parsing of words in sequence, the framework does not dispense with the concept of grammaticality despite the basic hypothesis that the only representation needed by the grammar is that of semantics.

4. Towards a more dynamic concept of context

Turning now to what the ellipsis data might have to tell us about what is involved in using context to build up interpretation, we will find that we have to see semantics as structural representations of content, syntax as the process of constructing these representations, and context as a store of content, structures, plus actions, the semantic representations and the syntactic process combined. For with these concepts in place, we

[8] Idiosyncratic exceptions such as most can be individually defined, as there is no requirement of exceptionless parallelism of structural and logical type. Indeed there is independent reason to think determiners are not a homogeneous category.

anticipate that hearers/speakers can both retrieve actions stored in context and re-use them to build up interpretation irrespective of which role they had previously been adopting; and a principled account of ellipsis follows immediately.

The types of information that ellipsis can pick up from context are of three basic sorts: the content expressible by some antecedent utterance, the structure made available by some antecedent utterance, and the processes made use of by some previous utterance. We take each in turn.

Context-dependent processing (a): recovery of content The type of ellipsis construal familiar from the linguistics literature is that ellipsis can select terms from (linguistic) context:

(19) Q: Who upset Mary? Ans: John did. (strict readings)

We display this as follows:

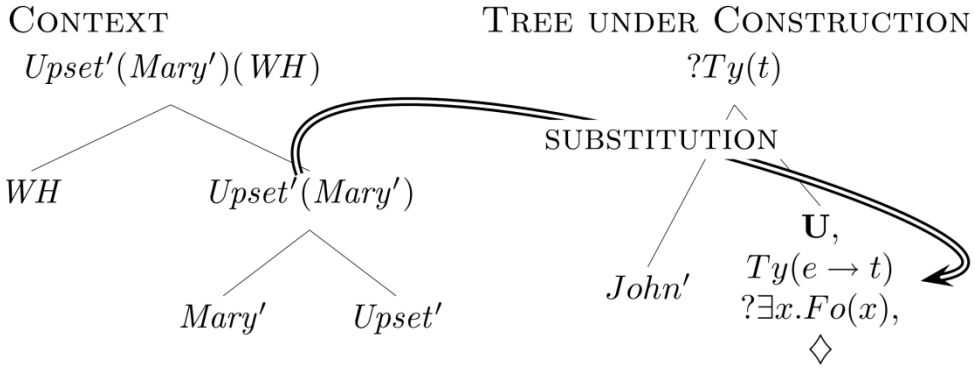

The formulation of this is exactly analogous to pronominal anaphora, though for some predicate type. The encoded expression projects a metavariable of type *(e→t)*, which then acts as trigger to license recovery of some appropriate content from the immediate context.

It is notable that this commits us to the correct observation that ellipsis can be interpreted relative to the non-linguistic, indeed nonverbal context if the predicate is sufficiently salient (contra Hankamer and Sag 1976 and others):

(20) Parent to teenage son with surf-board standing in shallows:
I wouldn't if I were you. The flag's flying, so it'll be dangerous

Context-dependent processing (b): re-use of structure In addition, the very structure made available by the immediate context can also be used, in some sense directly, with the respondent presuming on its availability as the point of departure for their own utterance. This is characteristic of question answer pairs in particular:

(21) Q: Who did John upset ? Ans: Himself.

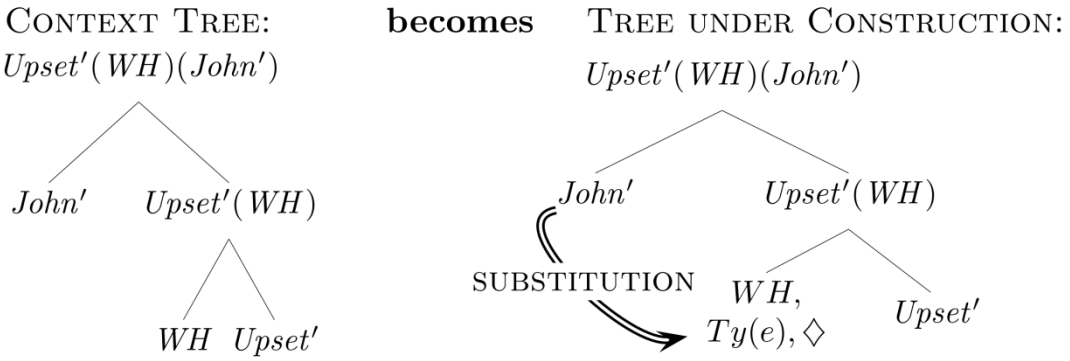

As the above display shows, the structure is in some sense shared between speaker and hearer in these exchanges, and the question is what that amounts to. In this case, this structure contains a specialised variable which the *wh*-expression is defined to provide as a place-holding device for a subsequent answer (see Kempson et al 2001: Chapter 5), and this in the reply is replaced by a reflexive pronoun, which duly has to be identified as the term inhabiting the subject node. Hence the interpretation of the fragment as providing an answer to the question, by update of the very structure which it itself provides. As we shall very shortly see, it is this hand-over of structure which is characteristic of the split-utterance phenomenon.

Context-dependent processing (c): re-using actions from context Before turning to such data however, there is one further important preliminary. It is not merely the output structure or content which participants in language performance can make use of but also the actions which each of them may severally use in building up such structural representations, for it is such re-use which enables patterns of structure-building to be replicated without identity of resulting content:

(22) A: Who hit himself?
B: John did. Bill too.

In (22), as in (19), B's (initial) answer involves the predicate-anaphoric device, the verb *do*, which constitutes a trigger to establish some appropriate predicate from what context provides. But in this case, it is actions recovered from context which yield that value (rather than some given predicate-formula as in the strict interpretation). What is recovered are the actions lexically encoded in the verb *hit* immediately followed by the actions encoded in the reflexive pronoun. As applied in the new environment provided by the ellipsis site – *John did* – this sequence of actions will again involve the projection of the two-place predicate *hit′* (with its past-time temporal indication) and a place holding metavariable at the object node, but the subsequent local identification of that metavariable as subject will, in this new environment, ensure that the object node decoration will be the formula *John′*. The result is the sloppy interpretation 'John hit himself'. A similar pattern of reiterated actions can be applied also in interpreting the following add-on of *Bill too* – all without any repeat of the word *himself* to trigger such distinct local bindings. It is simply the selection of an action sequence, as indicated by the type-specification of the trigger, which are taken over from context and reiterated.

This re-use of actions is what underpins the broad array of sloppy readings that linguists have identified:

(23) John upset his mother. Harry too.
(24) The man [who arrested *John*] failed to read *him his* rights. The man who arrested Tom did too (Wescoat 1993)

In all these cases, it is not the output content which is replicated, or needs manipulation in order to create some novel content: it is simply reiteration of actions stored in the evolving context of the emergent dialogue.

The account has, as a bonus, the prediction of seamless switching between speaker and hearer roles that is diagnostic of conversational dialogue. Unlike other frameworks, for which such split utterances pose

very considerable problems (see Gregoromichelaki et al. 2009), on the DS account, this phenomenon is predicted to be wholly straightforward, indeed their existence is a consequence of the DS account of production (Purver et al 2007). According to the DS account, the very same mechanisms are used in production as in parsing. Tight coordination between the parties is expected. Each party is building up structure relative to their own context, so at any point, making use of that individually constructed representation whether as parser or producer, they can switch roles and take over the other role, the only difference between the two activities being the greater specificity of the goal to be achieved in production. So, even with a role switch and the first and second person pronouns having to be reinterpreted, the mechanism for processing them remains identical.

(25) A: Did you give me back
B: your penknife? It's on the table.
(26) A: I heard a shout. Did you
B: Burn myself? No, luckily.

As these display, the context used by a participant as a producer/hearer is exactly that of the context they use in their shifted role as a hearer/producer. So with the incorporation of the dynamics of structure built in to the grammar itself, a very considerably larger dataset becomes characterisable. The split utterances, so signally ignored in accounts of ellipsis that purport to be a sub part of sentence-based grammar become core data, relative to which competing grammars can be compared.

The significance of this success in extending the ellipsis account to include split utterance data, is that we have an account of ellipsis that reflects directly the folk intuition of ellipsis being a window on context; we can see ellipsis in parallel with anaphora as both being intrinsically context-dependent and we can capture the heterogeneity of the resulting contents that are expressible while nevertheless retaining an integrated account of the phenomenon itself.

5. Grammars for linguistic processing

With this sketch of ellipsis and the attendant concept of context as a promissory note, we now turn finally to the concept of grammar that the DS system provides, and what this tells us about the syntax-pragmatics interface. We have a grammar that is defined in terms of mechanisms for word-by-word incremental parsing and production in context. The core notion is that of representations of content, and more precisely growth of such representations. Syntax constitutes meta-level constraints on tree growth, with lexical specifications inducing procedures for tree growth. Even quantifiers are seen as procedures for name construction, with mechanisms for inducing arbitrarily rich restrictor specifications. This gives rise to a process of name construction which is just as incremental and growth-driven as phenomena that are more familiarly syntactic. The ontology is thus whole-heartedly representational.

These assumptions involve no necessary mapping onto denotations defined by grammar: the account is exclusively in terms of transitions over partial structure. This commits us to some notion of a language of thought but it is not Fodorian. First of all, the notion of 'concept' is quite different -words are capacities for constructing concepts, and these do not stand in one to one correspondence. Of the grammar architecture itself, there is no structural level as such, even at the level of such constructed logical forms: the system is, rather, a set of mechanisms for incrementally building such structures. There is considerable philosophical significance in this shift of perspective. Because the system provides tools for inducing transitions from one partial structure to some enrichment, there is no multi-level representationalism. All generalisations over the input specifications, whether morphological, morpho-syntactic or syntactic are expressed over growth of logical structures, with the resulting structures reflecting thoughts.

In this double move of no incorporation of semantics into the grammar system on the one hand, and no level of structure inhabited by words on the other, it might seem that the framework runs aground because it fails to provide any notion of compositionality, a yardstick against which semantic theories have been evaluated for the last forty years. However, this would not be the right conclusion to draw. There

isn't just one concept of compositionality for articulating the relation of NL sequences to a compositional account of interpretation, but two such concepts, which are linked together. First, there is the restriction of word-by-word incrementality, as articulated in the restriction that all derivation steps involve monotonic growth of partial trees reflecting growth of information content as progressively induced from the sequence of words. Second, there is strict bottom-up compositionality of content for the resultant structure. Spelling this out, the only licensed output representations are those which have no outstanding requirements for further update; and decoration of mother (hence nonterminal) nodes invariably depends on having appropriately assigned type and formula values to those daughter nodes to allow application of functor to argument by a step of type deduction (*beta* reduction), with a requirement on those daughter specifications as input to such a step of type deduction that these input specifications include no outstanding requirements. Apparent instances of non-compositionality are those where the initial processing of an expression only determines some partial set of decorations on the node(s) it introduces, leaving an outstanding requirement on any such node. These meet the constraint that the actions given in lexical specifications must be carried out in the order of the word sequence, but they give the appearance of the system being non-compositional, in so far as bottom-to-top compositionality cannot be defined over the emergent string but only over progressively established updates to an intermediate characterisation of that string which later expressions provide. But the advantage of having a system in which partiality and update is central is that it allows nodes as constructed to be only partially decorated, with license for a return to that node for further update. In addition to the familiar case of long-distance dependency, there is the case of expletive pronouns:

(27) It's likely that I am wrong

In these an incompletely decorated but type-specified node is taken to inhabit the subject node allowing the incremental parse process to proceed, but with an outstanding requirement at that node for some formula decoration that is not yet provided. In virtue of this outstanding

requirement, the pointer must return to this node in order to allow such a value to be determined by further (structural) development of substructure from that node. In the case of the subject specification for (27), it is only once such a subtree is duly constructed and suitably decorated so that all requirements for values are met, that a formula value can be provided at that subject node, and from there to the assignment of value to the propositional formula of type *t* as a whole, finally satisfying the initially imposed ?*Ty*(*t*) requirement. The concept of wellformedness articulated within the system turns on the feeding relation between these two concepts. Wellformedness of a sentence string inheres in the existence of at least one derivation from a given axiom using each word in sequence and appropriate computational lexical and pragmatic actions of the system as applicable – the constraint imposed by the first compositionality requirement – to an end result which is a tree in which no requirements on node decorations of either terminal or nonterminal nodes are left outstanding – the constraint imposed by the second compositionality requirement. In sum, this is a representationalist theory of language, but not because it posits multiple and obscure levels of representation, but because the sole core of the grammar is the projection of structures of thought in a way that directly reflects the word sequences as incrementally parsed.

Finally there is the significance of the lack of domain-specificity or encapsulation. The vocabulary with which the concepts of tree growth are defined is not specific to the domain of NL processing. It is indeed essential to the successful application of the sets of actions defined in any grammar formalism, that their application may take place in a context of structured representations which may themselves have been derived either from previous language processing or from other forms of input, vision, hearing, reasoning over inputs derived from any of these, etc. Moreover, it is essential that such application not be encapsulated in the sense of being insulated from any other such input, if we are to meet the challenge of explaining context-dependency, as set out initially. This is because, as we saw in Section 2, any such insulation would immediately lead to the bifurcation of data into those whose characterisation fits within the sententialist remit and those which do not. And, with the split-utterance data of dialogue amongst those that do

not fit this remit, no structurally identifiable patterns in any language would receive a complete characterisation.

Coming to terms with the lack of domain specificity has the bonus of opening whole new avenues of research into areas such as language acquisition. We no longer need to see first language acquisition as the acquiring of an encapsulated system of an order of complexity that demands assumptions of innateness of a capacity specific to language. This is not however to adopt a "blank sheet" view of language, notoriously ridiculed in Chomsky (1959) in his onslaught against behaviorism. However, the shift into a system which is not domain specific in all aspects other than processing particularities of the input signal would lead one to expect that putative innate capacities would be those in which cognitive activity in general would be grounded. This remains a program for further research; however in anticipation of such research, we note the following. First, if an agent is to function in the world around them, they must be able to individuate information derivable from such input signals, which we might categorise as of primitive type *e*. In order, then, to use this information in activities of reasoning, an agent must also have the ability to construct predicative concepts in order to use the individuated information as itself constituting an input to a system of reasoning (in other words having the ability to construct functions from forms of type *e* onto forms of type *t*). There may of course be different modes of reasoning, but, using tools familiar from the study of logic, the two primitive types and the attendant concept of *predicate*, of type $e \rightarrow t$ thus seem minimal. Finally, to be able to engage in any such constructive act that can lead to inferential effects, the agent must have a means of processing within restrictions determined by processing in real time. For this to be sufficiently flexible to be viable, the agent will need to be able to hang on to recovered information across some restricted period in that time line until such time as other information with which it can be combined is recoverable. This means having the capacity to build "a parsing platform", an insight which is at the heart of the Dynamic Syntax analysis of discontinuity effects. This is arguably a capacity specific to human beings, but on the perspective opened up by these assumptions, this is a consequence of the human-specific capacity to process

information in real time, and not of an encapsulated, language-particular module separated off from such general cognitive capacity.

Acknowledgments

This paper reports joint work with Patrick Healey, Matthew Purver, Arash Eshghi and Wilfried Meyer-Viol, carried out as part of the ESRC-funded The Dynamics of Conversational Dialogue project, RES-062-23-0962 for whose support we are grateful. The third author also thanks the Leverhulme Foundation for a Major Leverhulme Fellowship held during this period. We thank our co-workers for all input, and also the numerous friends, colleagues and audiences who have provided helpful comments at various stages.

References

Cann, Ronnie, Ruth Kempson and Eleni Gregoromichelaki 2009. *Semantics: an Introduction to Meaning in Language*. Cambridge: Cambridge University Press.

Cann, Ronnie, Ruth Kempson and Lutz Marten 2005. *The Dynamics of Language.* Elsevier (now Emerald publications).

Cann, Ronnie, Ruth Kempson and Matthew Purver. 2007. Context-Dependent Wellformedness: the Dynamics of Ellipsis. *Research on Language and Computation* 5, 333-58.

Chomsky, Noam 1959. A review of B. F. Skinner's *Verbal Behavior*. *Language* 35(l), 26-58.

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

Dalrymple, Mary 2001. *Lexical Functional Grammar.* New York: Academic Press.

Dalrymple, Mary, Stuart Shieber and Fernando Pereira 1991. Ellipsis and Higher-order Unification. *Linguistics and Philosophy* 14, 399-452.

Ferreira, Victor and Garry Dell 2000. The effect of ambiguity and lexical avail ability on syntactic and lexical production. *Cognitive Psychology* 40, 296-340.

Ginzburg Jonathan and Robin Cooper 2004. Clarification, ellipsis, and the nature of contextual updates. *Linguistics and Philosophy* 27, 297-366.

Gregoromichelaki, Eleni, Ruth Kempson, Matthew Purver, Gregory Mills, Ronnie Cann, Wilfried Meyer-Viol and Patrick Healey 2011. Incrementality and Intention-Recognition in Utterance Processing. *Dialogue and Discourse* 2, 199-233.

Hankamer, Jorge and Ivan Sag 1976. Deep and Surface Anaphora. *Linguistic Inquiry* 7, 391-428.

Hilbert, David and Paul Bernays 1939. *Grundlagen der Mathematik.* Berlin: Verlag Julius Springer.

Kempson, Ruth, Wilfried Meyer-Viol and Dov Gabbay 2001. *Dynamic Syntax*. Oxford: Blackwell.

Merchant, Jason 2009. Ellipsis. In: Artemis Alexiadou, Tibor Kiss, and Miriam Butt (eds.), *Handbook of Contemporary Syntax* (2nd edition). Walter de Gruyter: Berlin.

Purver, Matthew, Ronnie Cann and Ruth Kempson 2006. Grammars as Parsers: the Dialogue Challenge. *Research on Language and Computation* 4, 289-326.

Purver, Matthew, Christine Howes, Eleni Gregoromichelaki and Patrick Healey 2009. Split Utterances in Dialogue: a Corpus Study. In *Proceedings of SIGDIAL 2009.* London.

Purver, Matthew, Arash Eshghi and Julian Hough 2011. Incremental semantic construction in a dialogue system. In: Bos Johan and Stephen Pulman (eds.), *Proceedings of the Ninth International Conference on Computational Semantics.* Oxford, 365-369.

Stainton, Robert 2006. *Words and Thoughts: Subsentences, Ellipsis, and the Philosophy of Language*. Oxford: Blackwell.