Res Lang Comput (2007) 5:333–358
DOI 10.1007/s11168-007-9033-3

Context and Well-Formedness: The Dynamics of Ellipsis

Ronnie Cann · **Ruth Kempson** · **Matthew Purver**

Received: 2 July 2007 / Accepted: 2 July 2007 / Published online: 21 September 2007

Abstract This paper challenges the tradition of defining grammars and grammaticality independently of the context of utterance. Using dialogue phenomena, in particular elliptical utterances, it argues that the obvious dependence of such utterances on context to recover the intended interpretation should be regarded as an inherent characteristic of natural language grammars and thus applicable to the characterisation of grammaticality for all natural language strings. The paper adopts the framework of Dynamic Syntax which shifts the burden of syntactic explanation away from the definition of de-contextualised syntactic structures defined over strings of words towards the characterisation of syntax as a context-dependent, incremental process whereby interpretations of strings in context are progressively built up as an utterance proceeds. This change in the way syntax is conceived, together with a demonstration that the same processes for building interpretations are used in generation as in parsing, is shown to allow a unitary account of anaphora and a range of elliptical phenomena that is typically precluded in non-dynamic, structure-based theories of syntax. The paper ends by providing formal definitions of well-formedness with respect to context that preserve traditional notions of grammaticality while allowing more fine-grained characterisations of well-formedness to distinguish acceptability from full (un)grammaticality.

R. Cann (✉)
Department of Linguistics, Edinburgh University, Edinburgh, EH8 9LN, UK
e-mail: r.cann@ed.ac.uk

R. Kempson
Department of Philosophy, King's College London, London, WC2R 2LS, UK
e-mail: ruth.kempson@kcl.ac.uk

M. Purver
Center for the Study of Language and Information, Stanford University, Stanford, CA 94305, USA
e-mail: mpurver@stanford.edu

Keywords Dynamic Syntax · Incrementality · Parsing · Context-dependence · Grammaticality · Well-formedness · Anaphora · Ellipsis

1 Preliminaries: Shifting the Ground

Following the pattern of grammars for formal languages, two methodological assumptions have driven the development of natural-language grammars over the past half-century. The first is that a natural-language grammar is a set of principles that determines the set of wellformed sentence-strings by associating with all such strings a structure and an interpretation compositionally derived from the interpretations of the elementary expressions and their mode of combination in that structure. Linguists may disagree over the concept of interpretation to be assigned to strings; but none doubt that syntactic principles induce structure over the strings on the basis of which interpretations for those strings are definable. The second is that these syntactic principles, whatever form they take, are independent of any properties that might be attributable to the dynamics of how language is used in processing in real time, and have to be seen as feeding theories of performance/pragmatics to determine how language is used/processed in context. Accordingly a string is said to be well-formed iff it is licensed by principles internal to the grammar formalism without any reference to context or to mechanisms for processing the string. For any phenomenon which simultaneously displays structural restrictions and dependence on context—anaphora, ellipsis, etc—these assumptions preclude a unitary account (see (Kamp and Reyle 1993) and many others); and it is standard to distinguish grammar-internal anaphoric binding and discourse anaphora. In this paper we challenge these assumptions by arguing that an integrated characterisation of ellipsis becomes possible if we shift to a methodology in which a grammar formalism makes available a set of procedures for progressively building up structure (corresponding to the interpretation of a string) *relative to context*, rather than assigning structure to the string itself independent of context. We adopt the framework of Dynamic Syntax (DS) (Kempson et al. 2001; Cann et al. 2005), in which syntax is defined in terms of procedures for building up interpretation from words in sequence, in a way that reflects the time-linear dynamics of parsing. We define a concept of context which records not only the interpretation assigned to a string but also the process of building it. We show how these assumptions allow an account of ellipsis which captures the diversity of effects while nevertheless sustaining a unitary form of analysis, and we then articulate a range of concepts of context-dependent wellformedness to provide a characterisation of what it means for an elliptical fragment to be well-formed. The result is a more fine-grained characterisation of well-formedness which nevertheless preserves traditional concepts of (un)grammaticality.

2 Context and Well-Formedness

Almost every natural language expression displays some form of context-dependence. This is most obvious with processes of anaphora and tense construal, but the effects go much further than this. An utterance of a sentence like *I bumped into Mary yesterday* *in the park* will convey different information according to who the speaker is, who Mary might be, where the park is (and what park) and when the sentence was uttered. Within linguistic theory, such matters are generally treated as the province of semantics and/or pragmatics, the syntax merely providing some decontextualised compositional analysis of the string of words that inputs into the semantic interpretation. Some justification for this stance can be seen when contradictory information arises, as this gives rise to judgements of anomaly rather than ill-formedness (1.2–3).

(1) 1. I bumped into Mary yesterday.
  2. #I will bump into Mary yesterday.
  3. #I bumped into Mary tomorrow.

Given such an approach, however, a string like *He cried* containing a pronoun is treated by the syntax as well-formed irrespective of whether the context provides an appropriate antecedent for the construal of the pronoun, although a sentence like (2.2) is peculiar, if not ill-formed, with no prior context to provide a male referent. Anaphora thus at least provides an argument for contextual constraints on the well-formedness of discourse, and thus of acceptability, if not grammaticality.

(2) 1. Bill hit his head on the doorframe and he cried.
  2. #Mary hit her head on the doorframe and he cried.

There are, however, well-known syntactic phenomena that more obviously require some reference to context to determine well-formedness, in particular elliptical constructions, where the preceding linguistic context is essential in determining the well-formedness of the string. In these cases, reference to context is essential in determining grammaticality, and responses to contradictory (incompatible) information gives rise to responses of ungrammaticality or stronger problems with acceptability. So (4.1,4) are fully ungrammatical, (4.2) is very odd and (4.3) anomalous because the natural reading of the elliptical fragment is 'some lieder were sick'.

(3) 1. Mary washed her hair and so did Bill.
  2. Bill dislikes something but it's not clear what.
  3. Sue sang a ballad for John and some lieder too.
  4. Sue gave John a book and Bill a CD.

(4) 1. *Mary was tall and so did Bill.
  2. #Bill dislikes coffee but it's not clear what.
  3. #Sue is sick, and some lieder too.
  4. *Sue sings well and Bill a CD.

While the obvious context dependence of elliptical constructions such as those in (3) has received a considerable amount of attention in the literature, (Dalrymple et al. 1991; Hardt 1999; Kempson et al. 1999; Lappin 1996; Reinhart 1991, etc.), it is only within the confines of the sentence that this is definable, since grammars are taken, by definition, to describe only the structure of sentences treated in isolation. Yet such constructions may cross sentential boundaries (5.1–3), may be uttered by other speakers (5.4), and may constitute answers to questions (5.5–6).

(5) 1. A: Mary washed her hair. B: So did Bill.
2. Bill dislikes coffee. I don't know why.
3. Sue sang a ballad for John. Some lieder too.
4. A: Sue gave John a book. B: And Bill a CD.
5. A: Who washed the dishes? B: John (did).
6. A: Who does everyone love? B: Themselves / Their mother.

The licensing context for ellipsis may thus be extra-sentential and so not purely to be determined within a single sentential domain (see e.g. (Pulman 2000) for VP ellipsis, (Ginzburg and Cooper 2004) for NP fragments). This is particularly problematic because exactly the same sort of well-formedness effects hold across sentences as within them. Compare (4) with (6).

(6) 1. Mary was tall. *So did Bill.
2. Bill dislikes coffee. *It's not clear what.
3. Sue is sick. #Some lieder, too.
4. Sue sings well. *And Bill a CD.

Despite the fact that the existence of such data is well known, ellipsis is standardly treated as either a syntactic process involving null productions from a complete structure (Stanley 2000; Merchant 2001), or as a semantic process which involves an abstraction operation on some antecedent content in order to provide something with which the content of the elliptical expression may combine ((Dalrymple et al. 1991) and others following). This is important because it is otherwise unclear what status elliptical fragments have within the grammar. Under the normal Chomskyan definition of competence as knowledge of how *sentences* are constructed independently of anything external to the linguistic system, fragments can be defined with respect only to the context provided by the sentence under construction.

On the other hand, to account for elliptical fragments in discourse such as those in (5) on this view requires an entirely different story: such strings cannot be licensed in the same way, as the licensing context is not part of the same sentence, a problem that is further compounded by dialogue data where interlocutors complete each other's utterances, as in (7):

(7) 1. *Ruth*: What did Alex …
Hugh: give Eliot? A rabbit.
2. *Ruth*: Where have you got to …
Hugh: with your book? Not past the first page.
3. *Ruth*: Did you remember to give …
Hugh: Eliot his present? Of course I did.

We immediately face the problem of what the grammatical status is of fragments in a discourse: are all fragments well-formed? The obvious and immediate answer is 'Of course not'. But if that is so, we must address the issue of context dependence full on, since the responses in (7), at least, have some grammatical status (containing as they do one or more complete constituents leading to complete sentence formation) unlike the initial fragments (which contain an incomplete constituent).

So, it would appear that either we need a theory of ellipsis in dialogue that is independent of sentence internal ellipsis, or we need to assume that any fragment that is licensed within a sentence can appear as a well-formed string in its own right. Neither option is particularly attractive. In the first case, we abandon any pretense at a unitary characterisation of intra- and inter-sentential ellipsis, despite the fact that they show the same syntactic properties. In the second, we lose sight of the fact that elliptical strings have a very restricted distribution, dependent entirely on the immediate linguistic context. All such analyses fail to reflect the informal intuition that ellipsis is a device in which context itself directly provides the way the ellipsis site is to be interpreted. An alternative strategy to account for elliptical phenomena is to abandon the entrenched idea that context is irrelevant to syntax and provide a general characterisation of such processes that is blind to whether the triggering context is internal or external to the sentence. It is this perspective that we propose, against the background of Dynamic Syntax (Kempson et al. 2001). Moreover, as we shall see, the reason why an integrated account of ellipsis is possible within this framework is because syntax itself is defined as a process of structural update. As a bonus, we shall find that the account extends seamlessly to characterise instances of syntactic dependence displayed across speakers in dialogue, while nevertheless retaining a concept of a grammar of a language sui generis.

3 The Flow of Language Understanding

Dynamic Syntax (DS) is a parsing-directed grammar formalism in which a decorated tree structure that represents the semantic interpretation for a string is incrementally built up following the left-right sequence of the words. The concept of process is central, with syntax construed as the process by which semantically transparent structure is incrementally built up. There is no vocabulary articulating syntactic structures independent of these progressively established representations which themselves constitute the basis of syntactic explanation. General syntactic principles and lexical specifications provide actions that update partial structures, with the overall goal of defining a propositional structure representing the content of some string as uttered in a particular context, starting from a universal *requirement* to construct a propositional tree, i.e. one rooted in a type t node which is decorated by some propositional formula.[1] Importantly, this tree is not a model of syntactic structure in the sense of a structure inhabited by the linguistic string, but is a representation of some predicate-argument structure conveyed by an utterance of the sentence. This process is construed solely from a parsing perspective and is defined to update trees on a strictly time-linear and word-by-word basis. It is thus a sequence of labelled partial trees which constitutes the core of the structural characterisation, defined to progress from the initial propositional requirement (shown as $?Ty(t)$) to some complete tree:

[1] Fo is a predicate that takes a logical formula as value, Ty a predicate that takes logical types as values, Tn a predicate that takes tree-node addresses as values.

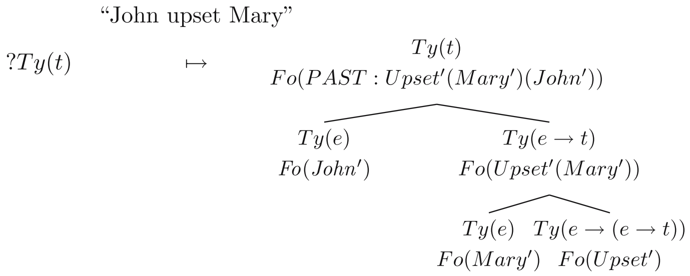

(8)

3.1 Parsing

The central tree-growth process of the model is defined in terms of the procedures whereby such structures are built up; taking the form both of general structure-building principles (*computational actions*) and of specific actions induced by parsing particular lexical items (*lexical actions*). The core of the formal language is the modal tree logic LOFT (Blackburn and Meyer-Viol 1994), which defines modal operators $\langle\downarrow\rangle$, $\langle\uparrow\rangle$, which are interpreted as indicating daughter and mother relations, respectively, with two subcases $\langle\downarrow_0\rangle$, and $\langle\downarrow_1\rangle$ distinguishing daughters decorated with argument or functor formulae, and two additional operators $\langle L\rangle$, $\langle L^{-1}\rangle$ to license paired *linked* trees.[2] The actions defined using this language are transition functions between intermediate states, which monotonically extend tree structures and node decorations. The concept of *requirement* is central to this process, $?X$ representing the imposition of a goal to establish X, for any label X. Requirements may thus take the form of requirements to construct formulae of particular types at the current node, such as $?Ty(t)$, $?Ty(e \rightarrow t)$, to construct formulae of particular types at some other node, e.g. $?\langle\downarrow_1\rangle Ty(e \rightarrow t)$, or to ensure that *some* value is found for a formula or treenode label at a node, e.g. $?\exists \mathbf{x}.Fo(\mathbf{x})$, $?\exists \mathbf{x}.Tn(\mathbf{x})$, etc. All requirements that are introduced have to be satisfied during the construction process. For example, one first action-sequence in parsing a string is the development of the standard initial Axiom state into a partial tree with requirements to find a subject and a predicate (in all such partial tree-structures, the pointer, $\diamond$, indicates the node under development):[3]

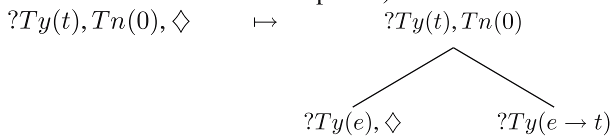

Words are specified in the lexicon to have *lexical actions*, each a sequence (ordered multi-set) of tree-update actions in an $\langle IF..THEN..ELSE\rangle$ format, employing the

[2] From node n, $\langle\downarrow\rangle X$ denotes 'X holds at a daughter of n'; $\langle\downarrow_0\rangle X$ 'X holds at an argument daughter of n', $\langle\downarrow_1\rangle X$ 'X holds at a functor daughter of n', $\langle\uparrow\rangle X$ denotes 'X holds at the mother of n' with subscripts 1, 0 indicating the node (functor, argument) from which the mother relation is described.

[3] One referee queries the non-lexical character of construction rules like axiom. It is possible that a fully lexical account of such rules could be given using the same vocabulary in which lexical actions are couched, but in this paper we continue to follow the theory as it is developed in (Cann et al. 2005), leaving the question open for future research.

explicitly procedural predicates make, go, put. A simple lexical action for a proper name *John* is given as follows:[4]

$$John \quad \left| \begin{array}{ll} \textbf{IF} & ?Ty(e) \\ \textbf{THEN} & \mathtt{put}(Ty(e)); \\ & \mathtt{put}(Fo(John')); \\ & \mathtt{put}([\downarrow]\bot) \\ \textbf{ELSE} & \text{ABORT} \end{array} \right.$$

This entry first checks that there is a requirement $?Ty(e)$ for the correct type at the active node, then adds decorations which specify a semantic formula $Fo(John')$ of this type, and that this is now a terminal node (shown by the modality $[\downarrow]\bot$ "below this node nothing holds"). More complex lexical actions are associated with (e.g.) transitive verbs, like *dislike*, which first make a new node of type $e \rightarrow (e \rightarrow t)$ to which a predicate term is assigned and then an argument node with a requirement for type e (to be fulfilled by parsing the object):

$$dislike \quad \left| \begin{array}{ll} \textbf{IF} & ?Ty(e \rightarrow t) \\ \textbf{THEN} & \mathtt{make}(\langle\downarrow_1\rangle); \mathtt{go}(\langle\downarrow_1\rangle); \\ & \mathtt{put}(Fo(\lambda x \lambda y.Dislike'(x)(y))); \\ & \mathtt{put}(Ty(e \rightarrow (e \rightarrow t))); \mathtt{put}([\downarrow]\bot); \\ & \mathtt{go}(\langle\uparrow_1\rangle); \\ & \mathtt{make}(\langle\downarrow_0\rangle); \mathtt{go}(\langle\downarrow_0\rangle); \mathtt{put}(?Ty(e)) \\ \textbf{ELSE} & \text{ABORT} \end{array} \right.$$

This format of lexical specification is general: all lexical items induce a mapping from one partial tree to another, but they are defined as idiosyncratic macros of update actions, the concept of lexical content being essentially procedural. These obligatory lexical actions, together with optional computational actions, induce a sequence of partial trees in a monotonic growth relation as each word is consumed in turn.

The closing stages of tree decoration, once tree node relations in a tree are fixed and all terminal node decorations fully determined, involve a modal form of type deduction progressively compiling decorations on mother nodes reflecting functional application of formulae on their daughter nodes. Once all requirements are satisfied and all partiality and underspecification is resolved, trees are *complete* (i.e. a topnode formula of type t is derived), parsing is successful, and the input string is said to be grammatical. Provisionally, then, we might say that a string is well-formed just in case it can be parsed using the computational rules of the system and lexical actions of each word in turn to produce a propositional tree that contains no outstanding requirements.

Giving more substance to the parsing process, we define a triple: $\langle T, W, A\rangle$, where T is a (possibly partial) propositional tree, W is the string of words so far parsed and A the sequence of actions (computational and lexical) used to construct T from W.

[4] These actions are a procedure for constructing a logical proper name which has a fixed denotation so that the name constructed from the word *John* is some $John_i$, which we represent here simply as *John'*. This analysis is not in principle disturbed by quantification, though we exclude all details in this paper, since quantified expressions are analysed as a sequence of actions for constructing epsilon calculus terms, all of type e, with scope expressed as constraints on evaluation of such terms.

Since at any point in a parse sequence, there may be more than one such triple, we define a PARSE STATE as a set of triples. The initial parse state P_0 contains only a single triple, in which T consists only of the initial AXIOM and W and A are both empty:

(9) $P_0 = \{\langle\{?Ty(t), \Diamond\}, \emptyset, \emptyset\rangle\}$

As parsing progresses, the cardinality of the parse state set may increase, as multiple hypothesised analyses of the string so far parsed are usually possible. A final, acceptable parse state is one in which there is a *complete* propositional tree T, i.e. one with no requirements outstanding, W is the complete parsed string and A the complete sequence of actions deriving T from W, taking each word in W in order. A string can thus be defined as well-formed just in case there is some sequence of parse states that can lead to a complete propositional tree, i.e. there is some set of actions that can map W in an ordered sequence onto a complete tree. In short grammaticality is defined as parsability: there is no central use-neutral grammar of the kind assumed by most approaches to parsing and/or generation.

3.2 Underspecification and Update

This analysis in terms of progressive update is part of a general pattern. At every non-final step in a sequence of tree transitions, input and output tree may be underspecified; and each parameter for tree decoration (values of the predicates Fo, Ty, Tn) is a possible source of underspecification. An example of an explicitly encoded underspecification of content (i.e. Fo value) is provided by anaphora, which has to be updated during the construal process. In this system, the lexical specification of a pronoun is defined to project a metavariable, together with an accompanying requirement $?\exists\mathbf{x}.Fo(\mathbf{x})$:[5]

he	**IF**	$?Ty(e)$
	THEN	$\texttt{put}(Ty(e))$;
		$\texttt{put}(Fo(\mathbf{U}_{Male'}))$;
		$\texttt{put}(?\exists\mathbf{x}.Fo(\mathbf{x}))$;
		$\texttt{put}(?\langle\uparrow_0\rangle Ty(t))$
	ELSE	ABORT

This requirement must be satisfied by substituting a fully specified Fo value from context as part of the construction process.[6] However, other than an analogue of the Binding Principles (Chomsky 1981) determining the local environment in which a value may *not* be provided, there is no constraint on the process determining what

[5] Though model-theoretic characterisations of anaphora construal have been predominant in the literature, there are also proof-theoretic accounts (Ranta 1994; Fernando 2002; Piwek 1998), to which this account is allied.

[6] The specification of the metavariable as $\mathbf{U}_{Male'}$ here expresses a (presuppositional) constraint restricting potential substituends to the correct gender. The additional final constraint in the lexical action shown above is a case constraint determining relative configurational position in the resulting tree, here $?\langle\uparrow_0\rangle Ty(t)$ (which is equivalent to requiring that this node fill the subject position). Other constraints, e.g. restriction to finite clauses, we ignore here: see (Cann et al. 2005).

does provide this value, and, as we shall see, there is more than one way in which this might be achieved, giving rise to the diversity of effects associated with anaphora construal.

A more radical form of underspecification, following up the concept of tree-growth dynamics, is provided by allowing tree node relations (Tn values) to be only partially specified, with subsequent update fixing that initial weak specification. Long-distance dependency effects are characterised in these terms: a tree-node with decorations provided by that left-peripheral expression being introduced in a partial tree as "unfixed", the relation of the newly introduced node to the node n from which it is introduced specified only as a constraint on some fixed extension (following D-Tree grammar formalisms (Marcus 1987)):[7]

$$\langle \uparrow_* \rangle Tn(n), ?\exists \mathbf{x}.Tn(\mathbf{x})$$

As with other requirements, such underspecification of tree-relation must get resolved within an individual tree constructed as part of the left-to-right construction process.[8]

The parallelism between anaphora and long-distance dependency effects is deliberate: both are identified in terms of underspecification and update, with this update required during the process of tree construction itself. And as we shall see below, the inclusion of explicitly context-dependent update processes allows resolution from context not only for the underspecification of content associated with anaphoric expressions, but also for aspects of structural underspecification which form part of the spectrum of ellipsis data. Quite generally in DS, concepts of underspecification and update are extended from semantics/pragmatics to syntax, and the various forms of underspecification are expressed in similar formal terms. In this way core syntactic phenomena such as long-distance dependency, relative clause binding, expletives, etc are expressed in terms essentially identical to those of anaphora: an immediate bonus is the anticipation of feeding relations between anaphora construal and structural processes, as established in detail elsewhere (Kempson et al. 2001; Cann et al. 2005). This bringing together of semantic and syntactic vocabularies has immediate application in the analysis of ellipsis, as we shall shortly see.

3.3 Preliminaries for Context Definition

As stated above, and as defined in (Kempson et al. 2001), pronouns project *metavariables*, to be replaced by some selected term. This may be found in context through a pragmatic process of SUBSTITUTION, as constrained by conditions on 'binding', Relevance Theoretic principles and any associated 'presupposition' (gender, person, number). Thus, in interpreting the answer in (10), **A**'s question provides the context

---

[7] Recall that $?\exists \mathbf{x}.Tn(\mathbf{x})$ is a requirement to find some specified value for the treenode label for the current node.

[8] In this, the system is like LFG, modelling long-distance dependency in the same terms as the LFG concept of *functional uncertainty* (Kaplan and Zaenen 1989), differing from that concept in the dynamics of update internal to the construction of a single tree, with relative clauses and other strong islands modelled as paired *linked* trees.

that allows *her* to be construed as *Mary* as there are two terms provided but only one satisfies the gender presupposition (and the binding principles):

(10) **A**. Who upset Mary?
B. John upset her.

If we can define the context $\mathfrak{C}$ available to a partial tree, this will allow us to define pragmatic processes like SUBSTITUTION in the same language as standard lexical and computational actions. (11) shows a possible formulation of the simple SUBSTITUTION action required for anaphora resolution: given a current node with a particular type $Ty(X)$ and an unfulfilled requirement for a formula value $?\exists \mathbf{x}.Fo(\mathbf{x})$, we can use a suitably typed and formula-specified node N in the context $\mathfrak{C}$ to provide a Fo value. In example (10), as long as the context $\mathfrak{C}$ contains the tree produced by A's utterance, this will license the identification of B's pronoun with $Mary$:

(11) SUBSTITUTION

IF	$Ty(X), ?\exists \mathbf{x}.Fo(\mathbf{x}),$
	$N \in \mathfrak{C},$
	$N = \{Ty(X), Fo(Y)\}$
THEN	$\mathtt{put}(Fo(Y))$
ELSE	ABORT

The question, of course, is: what is the context $\mathfrak{C}$, and what does it mean for a node to be a member of it? This is a question which is not formally raised in (Kempson et al. 2001) but is central to our considerations here. With respect to simple pronominal anaphora, the concept of context could be identified as a list of terms provided by the discourse so far constructed, much as in Discourse Representation Theory (Kamp and Reyle 1993). However, the range of potential referents for pronouns is wider than simply terms, including at least propositional formulae (12.1), and the sorts of anaphoric elements are not restricted to nominals (12.2–3):

(12) 1. Mary thought the man was a maniac and John believed it, too. *(Proposition)*
2. John screamed and Mary did, too. *(Predicate)*
3. Bill wanted a chocolate ice cream, but Sue wanted a strawberry one. *(Common Noun)*

So, our context must provide (at least) all terms, predicates, and propositions. But this is precisely what our tree representation of content already provides. So instead of adding an extra mechanism (a set of discourse referents) to the framework, a first DS step might be to assume that the trees representing the content of previous utterances themselves constitute the context. The process of SUBSTITUTION then targets a node from some tree in this context, selecting a formula value and writing it to the node decorated by the metavariable.[9] We can then represent the process of interpretation of (10) as in (13), where the double arrow indicates in shorthand the pragmatically

[9] Note that the theory does not currently provide an account of accessibility for available antecedents as does DRT. We assume that the basis of this process is given by Relevance Theory, rather than syntactico-semantic constraints.

constrained SUBSTITUTION operation between two partial trees, the one shown on the right hand side and another where $Fo(Mary')$ substitutes for $Fo(\mathbf{U})$.[10]

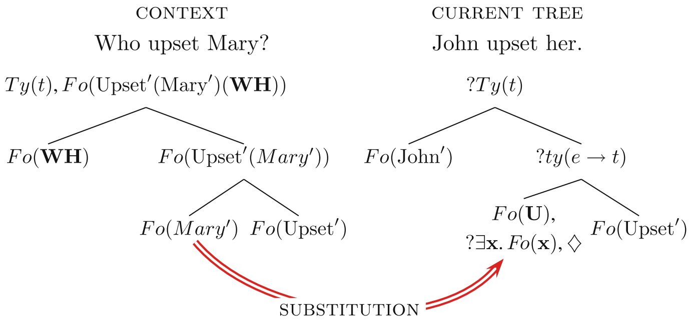

(13)

Note here that SUBSTITUTION must occur in the development of the second tree, updating the metavariable $Fo(\mathbf{U})$ of the object argument node to the value $Fo(Mary')$ in (13), for otherwise there remains an outstanding requirement ($?\exists\mathbf{x}.Fo(\mathbf{x})$) rendering the tree incomplete and the utterance ill-formed, given our initial characterisation of well-formedness. Thus, *John upset her* is treated as **not** well-formed in the absence of any accessible antecedent for the pronoun, a matter to which we return below (Sect. 5).

However, while the major cases of substitution for pronoun construal are restricted to substitution of some formula value established in some previous parse state,[11] we are going to find that context-based phenomena may make reference not merely to some resulting representation, but also to the process of building up such representations. One example is the so-called lazy use of pronouns, in which a pronoun may be interpreted using the actions whereby the antecedent had its construal established, rather than by simply copying the content assigned to the antecedent:

(14) John puts his pay in his bank account but Bill puts it in his post-office savings.

In (14), the pronoun *it* can be construed as referring to Bill's pay, rather than John's pay (as would be obtained by direct term substitution). This can be achieved by taking the *actions* previously used in establishing the interpretation of *his pay*, where the content of *his* was taken to be provided by the subject expression, and re-using them in the new environment, hence again identifying the value of the pronoun in the second conjunct as picking out the subject, Bill in this instance. It is this re-use of actions already employed in some immediate context that we wish to explore further in connection with ellipsis.

So instead of taking just trees and their decorations to be constitutive of the context, we use the richer concept of the triple $\langle T, W, A\rangle$, defined above, consisting of a (possibly partial) tree, a string of words, and a sequence of actions. Furthermore, we assume

[10] **WH** is a specialised metavariable whose value is not provided within the propositional tree that hosts it, that is restricted to co-occurring in a local type t domain typed with a Q feature denoting question-hood (thanks to Jonathan Ginzburg for this reminder).

[11] Or inference over it. We make no claim here that the immediately previously established structures are sufficient for cases of anaphora construal such as *She fainted outside the hospital. They operated on her right away.* Given the general inferential perspective, we do not take such cases to be especially problematic, though no mechanism is provided for them at this stage of DS development.

that a context consists of a sequence of such triples, representing the output provided by parsing all previous strings in the discourse. This characterisation is not quite sufficient, however, as substituends can be found from within the *current* propositional tree:

(15) Janet thinks she is pregnant.

Hence, we further include in the context sequence the current triple under development. A context $\mathfrak{C}$ is thus defined as a sequence of *inactive* triples (where an inactive triple is one whose tree is not currently under development i.e. does not contain the pointer) which we may call the discourse context, $\mathfrak{C}_D$, together with the current, active, triple, $\mathfrak{T}_\phi$ (i.e. $\mathfrak{C} = \mathfrak{C}_D \oplus \mathfrak{T}_\phi$). A consequence of this characterisation, together with our notion of parse state defined above, is that where a final parse state contains more than one triple (reflecting ambiguity): either the intended interpretation is identifiable through pragmatic means, in which case only the relevant triple becomes part of the discourse context; or each triple in the final parse state is added to the previous context to create a set of different contexts, differing at least in the last triple to be added.[12]

We assume that discourse-initially, the initial inactive sequence of triples $\mathfrak{C}_D$ is empty.[13] At the end of each utterance U_n, new contexts are derived by taking each triple $\mathfrak{T}$ in the final parse state, P_n, and added to the previous context sequence ($\mathfrak{C}_{Dn+1} = \mathfrak{C}_{Dn} \oplus \mathfrak{T}_i$) and a new initial Axiom parse state P_0 created for the next utterance.[14] In many cases, the final parse state will be acceptable (containing at least one complete tree) and we assume that all partial (unsuccessful) trees are removed at this stage. In cases where an utterance is incomplete, there must be a choice for the parser either to keep using the current incomplete parse state, expecting this to be

---

12 This analysis promises to give an account of ambiguity in one utterance getting resolved only after processing a later one.

(a) A: Mary's at the cricket ground.
(b) B: Right.
(c) A: I saw her duck.
(d) B: uh huh.
(e) A: It was waddling around on the boundary.
(e′) A: Lucky she did—the ball almost hit her right in the head.
(e″) A: She was very disappointed that she scored no runs.

After processing A's second utterance, B has multiple possible contexts corresponding to the 3 possible interpretations of *I saw her duck*. On processing A's third utterance, B tries to extend each of these possible contexts. In (e), the pronoun *it* plus the meaning of *waddle* ensures that only the 'duck=bird' context can survive; in (e′), the ellipsis signalled by *did* makes sure that only the 'duck=action' context can persist; while in (e″) the ambiguity as to which context is correct is only resolved by the content of the embedded clause which provides the information required to ensure that only the 'duck=score-no-runs-in-cricket' context survives.

13 A fuller account might allow the initial context to include e.g. information given by the current visual situation for deictic anaphora.

14 We have nothing to say here about how hearers select between competing interpretations, simply presuming upon relevance or other pragmatic principles to determine how such choices get made. There is also clearly much to say about how a speaker or hearer decides when an utterance is complete, but here we simply assume that this can be done.

further developed, or add one or more triples to $\mathfrak{C}_D$ to create a set of new contexts and begin a new initial P_0 state.[15]

This definition now allows us to finalise our tentative definition of SUBSTITUTION (11), resulting in (16):

(16) SUBSTITUTION

IF	$Ty(X)$, $?\exists\mathbf{x}.Fo(\mathbf{x})$,
	$\langle T, W, A\rangle \in \mathfrak{C}$,
	$\{Ty(X), Fo(Y)\} \in T$
THEN	$\mathtt{put}(Fo(Y))$
ELSE	ABORT

3.4 Generation

Before turning to ellipsis, it should be noted that this account of context, and indeed the dynamics of the parse process in general, apply equally in the DS account of production (or rather its computational equivalent, generation).[16] Generation, according to (Otsuka and Purver 2003; Purver et al. 2006),[17] uses the same lexical entries and actions as parsing (these being constitutive of the grammar formalism); the difference from parsing comes in the existence of a known intended representation of content, the *goal tree*, against which the emergent parse tree is checked for subsumption at every putative parse step.

Formally, a *generation state* is defined exactly in the same terms as a parse state, except that there is an additional tree, the goal tree. A generation state G is a pair (T_g, X) of a goal tree T_g and a set X of pairs (S, P), where S is a candidate partial string and P is the associated parse state (a set of $\langle T, W, A\rangle$ triples). And the concept of context for generation can then be defined to match that of parsing: the context $\mathfrak{C}$ for any triple $\mathfrak{T}$ in a parse state P_i is $\mathfrak{T}$, plus $\mathfrak{C}_D$, a sequence of inactive triples derived from the final parse states from previous utterances.

Discourse-initially, the set X will contain only one pair, of an empty candidate string and the standard initial parse state, $(\emptyset, P_0)$. As generation progresses, multiple pairs are produced as candidate partial strings S are considered, each with their own associated parse state P. Consider the generation of a string like *John upset Mary*, for which the opening step is the initial parse state in (9) and as goal tree the second tree in (8). Just as in parsing, the parse state needs to be extended to eventually contain the goal tree, but unlike parsing the generator needs to select words from the lexicon to do

[15] Notice that this definition allows there to be triples entered into the context even if the string parsed was incomplete, leaving the parse incomplete. In cases where there has not been selection between competing interpretations, this may give rise to sets of contexts, but it may also instigate clarification procedures in order to reduce such expansion. Arguably this is one function of dialogue interruptions, where utterances may complete partial clauses that do not immediately precede the current utterance or even pick up anaphorically on utterances of incomplete clauses.
(i) A: I am concerned that Mary …B: Is she the one in accounts?

[16] Given the perspective on parsing, this cannot be more than an account of *tactical* generation, associating a tree-structure with a word sequence, rather than *strategic* generation, the determination of the intended tree-structures from underlying goals.

[17] Note that (Otsuka and Purver 2003)'s definition has no explicit context-dependence, for which see (Purver et al. 2006).

this. Having used computational rules to induce subject and predicate requirements, the word 'John' may be selected to yield a partial tree that subsumes the goal tree, as shown in the parse state in (17).

(17) Generating 'John' in *John upset Mary*:

Generating 'John' in *John upset Mary*:

T W A

$\Big\langle$

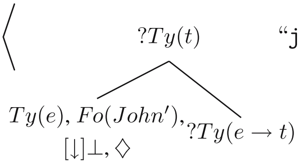

"john",

$\langle$make($\langle\downarrow_1\rangle$);
go($\langle\downarrow_1\rangle$);
put($?Ty(e \to t)$);
go($\langle\uparrow_1\rangle$);
make($\langle\downarrow_0\rangle$);
go($\langle\downarrow_0\rangle$);
put($?Ty(e)$)
put($Ty(e)$);
put($Fo(John')$);
put($[\downarrow]\bot$)$\rangle$ $\Big\rangle$

Thereafter, the generation of the rest of the string occurs in just the same way as parsing, but with each step constrained by the subsumption check on the goal tree. Once generation is complete, the final parse state P_1 paired with the chosen string W_1 in the final generation state is added to the sequence of final states maintained in context (just as with parsing).

Note here the close relationship between the parsing and generation processes. They share the same parsing actions, and the same basic component of their state (a parse state P, a set of tree/word-sequence/action-sequence triples)—a generation state merely adds to such triples the (partial) candidate strings and a goal tree. They must therefore make parallel use of context: the generation of *He smiled* in the context provided by an utterance of *John came in* is licensed not simply because the metavariable provided by *He* allows its partial tree to (trivially) subsume the goal tree, but because, following the parsing dynamics, a value for this metavariable must be identified from context, and the parse of the previously uttered string provides such a value $Fo(John')$ by SUBSTITUTION which (less trivially) allows subsumption.[18]

In addition, as both processes are strictly incremental, there is no requirement that their initial states be empty or contain only complete trees—they can start from *any* parse state or generation state. Switching between the two processes of parsing and generation, even in mid-sentence as in (7), therefore necessarily becomes straightforward. This result is strikingly different from formalisms in which the grammar formalism is use-neutral; there, any inter-process switching must be a result of good parsing/generation strategy design, rather than a necessary result of the grammar formalism itself (Purver et al. 2006).

[18] Use of context in generation also allows the high processing burden of lexical search to be minimized, important with this parsing-driven strategy, also helping explain psycholinguistically observed alignment and parallelism effects—see Purver et al. (2006).

4 Context-Dependence: Ellipsis

Ellipsis poses multiple challenges for grammatical theories, being the *sine qua non* of a phenomenon where strings of the language can only be interpreted with reference to the context; and there has been much dispute as to whether the phenomenon should be analysed syntactically, semantically, or pragmatically. On the one hand it appears to require construal over modes of interpretation and not over linguistic antecedents, in virtue of the so-called 'vehicle-change phenomenon' whereby on the assumption of syntactic reconstruction (or copy plus delete processes), there has to be license to replace some morphological features while retaining others (see (Fiengo and May 1994), among others), a phenomenon which favours a semantically based account:

(18) 1. John doesn't look after himself well, and neither do I
2. I never get around to washing my socks until the weekend. Do you?

On the other hand, at least some cases seem to require a linguistic antecedent, for example gapping (Hankamer and Sag 1976):

(19) 1. John interviewed Mandela, and Mary, Clinton.
2. *John interviewed Mandela and Mary, Clinton a sock.

Problems for both styles of analysis are posed by so-called antecedent-contained ellipsis (Fiengo and May 1994; Lappin 1996; Hornstein 1995), in which the ellipsis site is contained within the VP string, threatening syntactic accounts of VP ellipsis, but nevertheless displaying structural (island) restrictions imposed by the structure within which the ellipsis site occurs, which are problematic for semantic characterisations e.g. (Dalrymple et al. 1991; Hardt 1999), since they provide no basis for defining the fine structure specific to syntax with which to articulate such restrictions:

(20) 1. Joan read every book which Mary believed she had.
2. *Joan read every book which Mary believed the claim she had.

Despite lack of resolution to this debate, the focus of the disagreement has moved on; and protagonists are now arguing over whether all cases of ellipsis can be analysed in structural (syntactic) terms (Stanley 2000; Merchant 2001), or whether at least some cases require a much freer pragmatic basis in the light of (21), where there may be no linguistic antecedent at all (Stainton 2006):

(21) Bacon and eggs, please.

The challenge posed by this sequence of debates is whether ANY unitary basis for ellipsis construal can be provided; and the general consensus is that it remains irresolvably heterogeneous, with the informal intuition that ellipsis is a phenomenon where context provides interpretation having to be set aside. The question from a DS perspective is whether the defined concept of context provides a basis for meeting this challenge, given that it records not merely structure but also actions. We believe it does; but the characterisation involves relinquishing the assumption that grammar articulates structure without any reference to context, as we shall now see.

First, there are cases where the content of the ellipsis site is identified directly by some term from the context, in exactly the manner of anaphora: these are the strict readings of VP ellipsis, as in (22). Since *do* is an anaphoric expression constrained to ranging over event predicates, we take its lexical specification to project a predicate metavariable which is the necessary trigger for ellipsis construal to be licensed. In the conjunction context in (22),[19] the value for this metavariable is taken from the predicate decorating the initial tree:

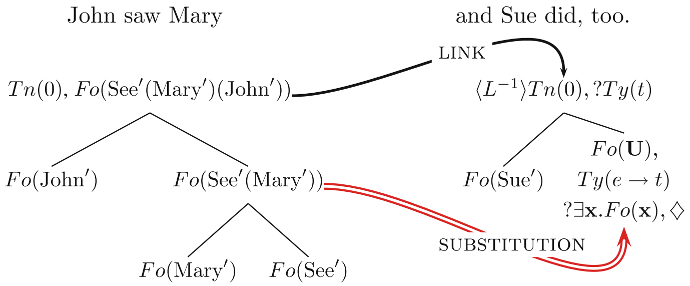

(22)

The result of the substitution operation yields a structure whose output interpretation is, as desired:[20]

$$Fo(\text{See}'(\text{Mary}')(\text{John}') \wedge \text{See}'(\text{Mary}')(\text{Sue}'))$$

Given that context is identically defined for both parsing and generation, both parties having reconstructed the same parse tree, the licensing tree for resolving strict ellipsis need not be part of the interpretation of the current utterance:

(23) **A**: Who upset Mary? **B**: John did.

It is however not just terms decorating some tree in the context that can be used in construal of fragments, but also the structure provided by that tree. Thus, a reflexive may be a perfectly natural response to a question such as in (24):

(24) **A**: Who did John upset? **B**: Himself.

These are the examples which confront purely semantic accounts of ellipsis as problematic, since the acceptability of such fragments is syntactically conditioned.[21] The problem is entirely general. Dependencies of all sorts, scope, negative polarity items, etc., can be distributed across question and answer and other use of elliptical fragments relative to their context:[22]

---

[19] Which we analyse in terms of *linked* tree structures (Cann et al. 2004, 2005). This relation is shown by the solid black arrow in (22) and in subsequent trees. It is to be distinguished from the double-arrow which shows the SUBSTITUTION process and not a tree relation.

[20] In addition to allowing for ambiguity as to the parsing strategy used, we also allow for ambiguity as to the antecedent that is chosen, as in:
(i) John criticised Mary because she wrote sloppy lectures; and Sue did too.
There is obviously a great deal more that needs to be said about this.

[21] See (Ginzburg et al. 2001) for a combined syntactic/semantic approach to bare answers to *wh*-questions.

[22] Again, limitations of space prevent us from giving an account of these examples, but their analysis is straightforwardly definable, given the DS use of epsilon terms and the incremental development of scope statements over the variables bound within them.

(25)

	1. **A**: Who did every student upset?	**B**: A lecturer. (Ambiguous: $\forall < \exists$, $\exists < \forall$)
(25)	2. **A**: What did everyone read?	**B**: A magazine.
	3. **A**: What haven't you understood?	**B**: Anything.
	4. **A**: Which classes did you fail to get to?	**B**: Any of the back-up classes?

This problem is solved if the tree constructed from the parse of the question is itself used as the structure that the fragment updates. Since, on the DS characterisations of parsing (perception) and generation (production), both speaker and hearer share the same discourse context, both have constructed the $\langle T, W, A\rangle$ triple which provides the successful parse of the question. So we have, as part of the context for the processing of the fragment in (24), the tree analysing the question (the first tree in (27)). The set of lexical actions associated with a reflexive targets some local formula value (indicated by the modality $\langle\uparrow_0\rangle\langle\uparrow_1^*\rangle\langle\downarrow_0\rangle$) and copies it onto the current node (26). The effect of this update on the first tree in (27) is shown in the tree in the second tree giving the required output: $Fo(Upset'(John')(John'))$. The shift in roles from speaker to addressee (and vice versa) crucially depends on context (and mechanisms) being the same in both production and perception.

(26) *himself*

IF	$?Ty(e)$		
THEN	IF	$\langle\uparrow_0\rangle?Ty(t)$	
	THEN	`Abort`	
	ELSE	IF	$\langle\uparrow_0\rangle\langle\uparrow_1^*\rangle\langle\downarrow_0\rangle Fo(\alpha)$
		THEN	$\mathtt{put}(Ty(e), Fo(\alpha))$
		ELSE	`Abort`
	ELSE	`Abort`	
ELSE	`Abort`		

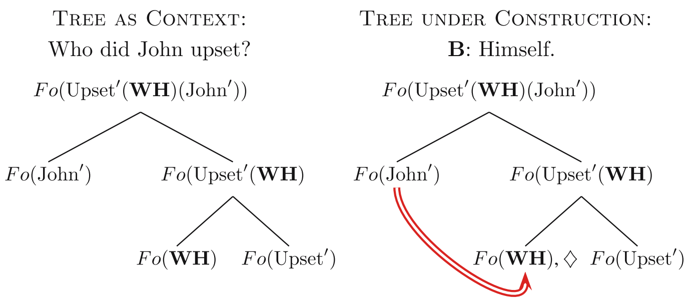

(27)

This is not the end of the account. Because the context is made up of a sequence of triples, it includes not only **trees** but **actions**, and this provides a major saving in the processing task; and it is the re-use of actions which we propose as the basis for sloppy readings of VP ellipsis. As with strict readings, the trigger for such re-use of context-provided constructs is triggered by the metavariable projected by *did*, but in re-running the actions, new interpretations can nevertheless be derived. So in the sloppy construal of the ellipsis site in (28), a new predicate is constructed using the very same procedures as used in setting up interpretation for the question—even to the identification of the antecedent from a particular node in the emergent tree—the one difference being that in the re-run, a different formula decorates the subject node, hence the different resulting interpretation (see (Purver et al. 2006) for details).

(28) **A**: Who upset his mother? **B**: John did. (= John upset John's mother)

The advantage of this form of analysis is that it provides a natural reflection of what may be matching parallelism in the way the ellipsis site is to be construed. Thus in (29), the construal of the fragment [3] as 'I trust Tom' appears to be echoing not merely the interpretation of the antecedent from which it got its interpretation, but the way in which that is built up too, as a detailed look at the actions used in building a construal of (29) [1]–[3] now makes plain.[23]

[1] I'll approach John. [2] Him, I trust. [3] Tom, too.

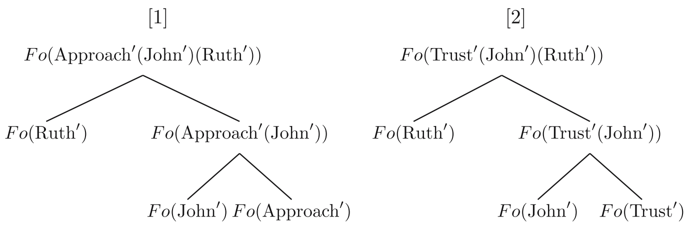

ACTIONS [1] from $?Ty(t)$
$\texttt{make}(\downarrow_0); \texttt{go}(\downarrow_0);$
$\texttt{put}(Fo(\mathbf{U}'_{spkr}), Ty(e)); \texttt{go}(\uparrow_0);$
$\texttt{make}(\downarrow_1); \texttt{go}(\downarrow_1);$
$\texttt{put}(?Ty(e \to t));$
$\texttt{make}(\downarrow_1); \texttt{go}(\downarrow_1);$
$\texttt{put}(Fo(\text{Approach}'),$
$Ty(e \to (e \to t)));$
$\texttt{go}(\uparrow_1); \texttt{make}(\downarrow_0);$
$\texttt{go}(\downarrow_0); \texttt{put}(?Ty(e))$

ACTIONS [2] from $?Ty(t)$
$\texttt{make}(\downarrow_0); \texttt{go}(\downarrow_0);$
$\texttt{put}(Fo(\mathbf{U}'_{spkr}), Ty(e)); \texttt{go}(\uparrow_0);$
$\texttt{make}(\downarrow_1); \texttt{go}(\downarrow_1);$
$\texttt{put}(?Ty(e \to t));$
$\texttt{make}(\downarrow_1); \texttt{go}(\downarrow_1);$
$\texttt{put}(Fo(\text{Trust}'),$
$Ty(e \to (e \to t)));$
$\texttt{go}(\uparrow_1); \texttt{make}(\downarrow_0);$
$\texttt{go}(\downarrow_0); \texttt{put}(?Ty(e));$
MERGE

(29)

The processing of both [1] and [2] in (29) share structural facets, in particular both induce a two-place predicate structure with two arguments. They differ solely in that the processing of [2] involves the initial construction of an unfixed node, and then the merging of this node with the argument node projected by the verb *approach*. The actions used in the processing of [1] and [2] both provide a sequence of actions which the parsing of the fragment *Tom, too* may pick up on. The first step in parsing the fragment [3], neutrally between any selected interpretation, allows the building of an unfixed node, decorating it with $Fo(Tom')$.

TREE UNDER CONSTRUCTION[3]:

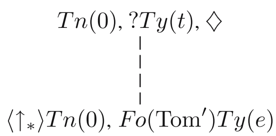

(30)

However, that structure may be completed by either the actions used to construct tree [1] or by the actions used to construct tree [2] up to the point at which *John* is parsed to

[23] We make the assumption here that it is only the actual actions used to construct the tree that are recorded. It may, however, be the case that the conditions on actions also get recorded in order to prevent their re-use in improper contexts. We leave the exploration of such a refinement for future research.

decorate the internal argument node. These two strategies yield two possible structures, both with the unfixed node now decorating the internal argument position, providing the two readings $Fo(Approach'(Tom')(Ruth'))$ and $Fo(Trust'(Tom')(Ruth'))$, respectively.

In either case, the fragment is initially taken to decorate an unfixed node, but then the manipulation of intervening actions culled from either one of the derivations in the context allows the fragment to be assigned distinct interpretations. The availability of more than one putative sequence of available actions from context yields ambiguity, exactly as in anaphora construal.[24] It should be noted here that re-use of actions will only be available if linguistic input has been processed; and this opens up a means of exploring a whole range of cases thought to be irreducibly syntactic such as antecedent-contained ellipsis, gapping etc, while sustaining an analysis of ellipsis as essentially context-dependent (see (Purver et al. 2006) for an outline of an account of antecedent contained ellipsis).[25]

Given that ellipsis has conventionally been analysed as a grammar-internal phenomenon, it might come as a surprise that indexical construal of fragments should even be possible, as observed by (Stainton 2004):

(31) Bacon and eggs

The advantage of the present formalism is that such fragments are not especially problematic, although, like indexical construal of pronouns, an assumption needs to be made that cognitive reasoning about the visually presented environment involves constructing structural representations (a relatively uncontroversial assumption given a representationalist methodology). All that is then required is an assumption that, in default of any other strategy, the hearer can construct an unfixed node and decorate it with some suitable existential term, and then construct out of the presented scenario a suitable relation whose object-argument node the constructed unfixed node can merge with.[26] The particular significance of this example is its buttressing of the DS claim of parallelism between the process of lexical underspecification plus update encoded

[24] This account of course remains to be explored in depth, given that we have not addressed quantification here. However, given that in DS scope constraints are construed as incrementally collected actions for evaluating the structurally complete predicate-argument structure, in principle parallelism effects involving quantification are expected (Kempson et al. 2001).

[25] One referee worries that the current account of ellipsis is not unitary, involving as it does both SUBSTITUTION and re-use of actions. However, a unitary account is not likely to be possible because of the substitution of metavariables by phrasal formulae. In computing the formulae decorating non-terminal nodes, the computational action uses functional application over the formulae of its two daughter nodes. While actions can be used to construct such nodes, some expressions that decorate nodes with metavariables (e.g. pronouns in English) explicitly disallow the node they decorate to dominate any other material, through the so-called 'bottom restriction' ($[\downarrow]\bot$) that signals a necessarily terminal node. If only actions were involved in determining the content of underspecified nodes, this would exclude (at least) pronouns being able to be interpreted as full terms, contrary to fact.

[26] In such commonly recurring scenarios, it is arguable that protagonists have stored ready-made the appropriate incomplete structures as templates to serve as contexts for processing such commonly presented types of input. We assume that such structures come without associated sequences of lexical actions, as they are not produced by parsing or generation, and therefore consist of semantic tree structures alone. As such this approach has parallels to approaches to fragments which assume the existence in context of a semantic *question under discussion* (Ginzburg et al. 2001) or a rhetorical relation (Schlangen 2003).

in anaphoric expressions and structural underspecification plus update, for with this analysis, we complete the picture: structural underspecification parallels anaphora underspecification in every respect. Both can be updated either from context, as this type of example shows, and, more familiarly both can be updated from the construction process, a parallelism between long-distance dependency effects and anaphora resolution which is uniquely expressible in DS (see (Cann et al. 2005) for an account of expletive pronouns in these terms).

Finally, as already noted in Sect. 3.4, this account of ellipsis extends naturally to dialogue phenomena in which utterances are split between interlocutors, with one of (7) repeated here:

(7) Ruth: What did Alex …
Hugh: give Eliot? A rabbit.

The essentially fragmentary follow-on provided by the interrupting interlocutor who completes the utterance is doing no more than using as their starting point the actions and structure which the context has provided (being the record of what they have just parsed), a strategy identical to that used in providing the answer. Conversely, too, for the speaker in becoming the hearer.

So the overall view of ellipsis as expressions whose interpretation is provided by the immediate context promises to be sustainable while nevertheless providing a basis for the diverse ellipsis effects.[27]

5 Well-Formedness in Context

Since the analyses provided depend on having defined structure for strings and the context relative to which they are evaluated, both the limited distribution of fragmentary expressions and their context-dependence are captured directly. Given a grammar formalism which articulates the progressive (time-linear) build up of interpretation over partial structures, we can then express a range of concepts of well-formedness: well-formedness with respect to a given context, well-formedness with respect to at least one context, well-formedness with respect to all/no contexts. This enables us to take into account not only fragments but, equally, continuations in shared utter-

[27] There are forms of ellipsis, where the antecedent form appears to have to be structurally distinct from that provided by the ellipsis site:

(i) Handouts are supposed to be circulated in the first week, and normally I try to.
(ii) John and Mary were dancing together, though Mary's mother had told them not to.
However, given that structure projected from strings is not taken in DS to be inhabited by the string itself, each putative problem-case needs to be taken in turn. For example passive verb forms might be said to induce initial construction of an unfixed node (by requiring this as the trigger for their own update), the resolution of which has to take place within the domain of a single tensed propositional structure, so that construal of *handouts* (i) involves decorating an initially unfixed node and subsequently resolving it as the object argument of the predicate 'circulate', a result which would provide an appropriate input structure for construal of the subsequent elliptical form. Though examples such as (ii) remain, at present, only analysable by stipulation of some essential intermediary inference step, presumption of inference over structures in context is not in any case incompatible with the DS style of analysis.

ances (Pickering and Garrod 2004), while still distinguishing all such expressions from classical context-independent well-formed sentences.

The notion of well-formedness expressed in (Kempson et al. 2001), and mentioned above, presumes on a concept of *complete tree*, where a complete tree contains no node decorated with an outstanding requirement. In other words, we may say that a string, ϕ, is well-formed just in case there is a possible development of the initial parse state defined by AXIOM (i.e. $P_0 = \{\langle\{?Ty(t), \Diamond\}, \emptyset, \emptyset\rangle\}$) to a parse state, P_n which contains at least one triple $\langle T, W, A\rangle$ where T is complete, W is the complete string of words uttered and A is the complete set of computational, lexical and pragmatic actions used to construct T from a strictly time-linear parse of W. We refer to such a triple as a *complete triple*. Our definitions of parsing and generation, together with the definition of *pragmatic* actions such as SUBSTITUTION, mean that such a characterisation requires reference to context to define well-formedness while maintaining the DS insight that well-formedness depends on parsability. We thus take an utterance of a string to be well-formed in a particular context, $\mathfrak{C}$, (a sequence of inactive triples plus the current, active triple) just in case the parsing of it using rules of the system gives rise to at least one complete triple:

(32) A string ϕ uttered with respect to a context, $\mathfrak{C}$, is well-formed iff:[28]

$$\mathfrak{C}_D \oplus \mathfrak{T}_0 \quad \mapsto_{A_\phi} \quad \mathfrak{C}_D \oplus \mathfrak{T}_\phi$$

where $\mathfrak{C}_D$ is the context given by the prior discourse (a sequence of inactive triples); $\mathfrak{T}_0$ is $\langle T_0, \emptyset, \emptyset\rangle$, the standard initial state; A_ϕ is the set of lexical, computational and pragmatic actions used in parsing ϕ on a strictly time-linear basis; and $\mathfrak{T}_\phi$ is complete (i.e. $\mathfrak{T}_\phi = \langle T_\phi, \phi, A_\phi\rangle$ where T_ϕ is a complete tree).

It is no coincidence that this concept of well-formedness with respect to a fixed context is exactly that of felicitous utterance, as put forward by (Heim 1982) (and the very similar concept of proper DRS of (Kamp and Reyle 1993)). The consequence of defining context-dependent concepts of well-formedness is thus that the characterisation of anaphora/ellipsis resolution falls fully within the remit of the formalism.

This shift to a concept of well-formedness explicitly in terms of context update allows us to articulate different concepts of *grammaticality*—well-formedness with respect to *all* contexts, with respect to *some* context and with respect to *no* context. The first essentially provides us with the standard (context-independent) concept of full grammaticality. Sentences which are well-formed in all contexts are either those in which no specific context is required (33a), or those in which the context required for the latter part of the string is provided by the first part (33b–d):[29]

[28] A more liberal definition might allow introduction of a distinct context $\mathfrak{C}'_D$, allowing the extension of $\mathfrak{C}_D$ by the addition of inferences derived from $\mathfrak{C}_D$: see Cann et al. (2005).

[29] Note here that all the examples in (33) still require some context dependence in terms of tense, natural gender and so on. We ignore these details here, but it is possible that there exist no sentences in any natural language that are purely context independent except perhaps for the logically true sentences.

(33) 1. No man is mortal.
2. A woman likes mustard though it makes her hot.
3. If John is a teacher, he will have a degree.
4. Janet thinks she is pregnant.

We may say that say that such strings are *fully grammatical* since they are well-formed in every context or, equivalently, they require no context to support their interpretation:

(34) A string ϕ is fully grammatical iff an utterance of ϕ is well-formed in the null context:

$$\emptyset \oplus \mathfrak{T}_0 \quad \mapsto_{A_\phi} \quad \emptyset \oplus \mathfrak{T}_\phi$$

where $\mathfrak{T}_0$, A_ϕ and $\mathfrak{T}_\phi$ are as defined in (32).

At the other end of the scale we have utterances that are not well-formed in any context. A sentence is *fully ungrammatical* if there is no context relative to which a string is derivable.[30]

(35) A string ϕ is fully ungrammatical iff there is no context, $\mathfrak{C}$, in which an utterance of ϕ is well-formed:

$$\neg\exists\mathfrak{C}[\mathfrak{C} \oplus \mathfrak{T}_0 \quad \mapsto_{A_\phi} \quad \mathfrak{C} \oplus \mathfrak{T}_\phi]$$

where $\mathfrak{T}_0$, A_ϕ and $\mathfrak{T}_\phi$ are as defined in (32).

We can in fact distinguish between two kinds of fully ungrammatical sentence: firstly those in which at some point the pointer cannot proceed, the parse process must abort, and the output state is empty, as in (36.1–2); and secondly those which license full sequences of actions, and thus lead to a non-empty output state, but one in which all trees are incomplete, as in (36.3-6). In these latter cases, some requirement on either computational or lexical action fails to be satisfied; the appropriate update cannot take place; and the result is a failure to derive an output tree with all requirements satisfied.

(36) 1. *The a in run.
2. *Word every no salad sleeps snores.
3. *Which man did you interview the man from London?
4. *The man from London emerged that he is sick.
5. *The man John saw whom is outside.
6. *Who did you see the man who came in with?

The concept of ungrammaticality is thus treated as categorical. However, only a relatively restricted set of sentences can be characterised as fully grammatical: arguably every uttered string displays some element of context-dependence, minimally that of the context of the act of utterance; and the concept of a null context is, in any case, a purely theoretical construct. Accordingly we articulate a concept of general *grammaticality* for strings: in essence, a string is grammatical (but not necessarily

[30] We make no claims here about whether the ungrammaticality of some string is decidable given this definition as the set of possible contexts is infinite.

fully grammatical, as not necessarily being grammatical given any arbitrary context) as long as there is at least one context in which an utterance thereof is well-formed, i.e. parsing it can lead to a complete triple.

(37) A string ϕ is grammatical iff an utterance of ϕ is well-formed in some context:

$$\exists \mathfrak{C}[\mathfrak{C} \oplus \mathfrak{T}_0 \quad \mapsto_{A_\phi} \quad \mathfrak{C} \oplus \mathfrak{T}_\phi]$$

where $\mathfrak{T}_0$, A_ϕ and $\mathfrak{T}_\phi$ are as defined in (32).

These definitions mean that the string *Janet thinks she is pregnant* is fully grammatical, since the transition from the initial state to completed triple is licensed through parsing the string even in the null context; while the string *He was in tears* is grammatical, since it is well-formed in any context that supplies a male antecedent, but there is no guarantee that it will be well-formed in any given context, according to the definition in (32). Strings in this class may thus be referred to in terms of acceptability, since their felicity will change from context to context. Thus, we can express the distinction between examples (39) and (38):

(38) John upset Mary. He ignored her.
(39) A book fell. *He ignored her.

The utterance of the string *He ignored her* with respect to a context obtained from having parsed *John upset Mary* is well-formed. In contrast, although the second string in (39) is potentially well-formed, its utterance with respect to the context provided by having just parsed *A book fell* is not well-formed, at least with no further expansion of the given context. The step of expanding the minimal context which would be necessary to determine the acceptability of *He ignored her* in the context of an utterance of *A book fell* is often taken as trivial, but on the characterisation of well-formedness for context-string pairs provided here, this sidewards move would have the effect of substituting the pair under characterisation with a different pair of objects.

The well-formedness of fragments and elliptical expressions is again subject to the contexts in which they are uttered. Well-formedness is thus defined very liberally to include the strings in (40):

(40)
1. John did, too.
2. John.
3. In the drawer.
4. Drawer.

The difference between these and fully sentential strings (such as *John upset Mary* or *Your keys are in the drawer*) is that the latter will typically have a wider range of contexts that define them as well-formed, while the former will be much more restricted. While liberal with respect to some data, the definition remains strict with respect to strings that cannot lead to well-formed complete propositional outputs, such as the fully ungrammatical examples of (36), and indeed the partial expressions in (41):

(41)
1. Have you read…?
2. Where are…?

3. I'm about to…
4. I'm about to visit the…

While such strings may be completed to provide a well-formed propositional tree, they cannot lead, by themselves, to a well-formed complete propositional parsing state given any context.[31] They are thus not themselves well-formed in any context and so ungrammatical, as desired. However, the same goes for *continuations* of these incomplete utterances, as in (42):

(42) 1. …read your latest chapter?
2. …visit the dentist?

While such continuations do lead to well-formed propositional trees, they do so by extending an active triple produced by the incomplete antecedent, rather than by proceeding from the standard axiom $\mathfrak{T}_0$ as required by (37). Given the definitions so far, there would therefore be no distinction between continuations, which intuitively lead to completeness in some sense, and the incomplete expressions of (41), which are undoubtedly incomplete as they stand. However, we can provide one, by defining a corresponding notion of *potential grammaticality*:

(43) A string ϕ is potentially grammatical iff an utterance of ϕ is well-formed in some context resulting from a previous utterance of a string ψ:

$$\exists \mathfrak{C}[\mathfrak{C} \oplus \mathfrak{T}_\psi \quad \mapsto_{A_\phi} \quad \mathfrak{C} \oplus \mathfrak{T}_\phi]$$

where A_ϕ and $\mathfrak{T}_\phi$ are as defined in (32), but here $\mathfrak{T}_\psi$ is some active triple $\langle T_\psi, \psi, A_\psi \rangle$ resulting from parsing ψ.

Of course, this can be seen as a generalization of (37): in the case where the previous utterance ψ is null, the corresponding active triple $\mathfrak{T}_\psi$ reduces to $\mathfrak{T}_0$, and the definition reduces to (37). We can thereby express the fact that continuations are potentially well-formed (in that they can lead to complete formulae, but only in very specific restricted contexts), while still maintaining a distinction from fully ungrammatical or incomplete strings on the one hand, and generally grammatical complete strings on the other.[32]

[31] Note here that the license for re-use of actions is a metavariable with formula requirement which is lacking in these examples, thus disallowing say (41a) from being well-formed in the context *I've read your chapter*.

[32] One of the interesting consequences of defining linguistic well-formedness in terms of licensing contexts for utterances and distinguishing (full, context-independent) grammaticality from (possible, context-dependent) acceptability, is that it provides a potentially different way to view gradient acceptability. Such phenomena are being increasingly studied, particularly with regard to Optimality Theory (see, in particular, (Keller and Sorace 2005)). The concept of gradience such approaches define, however, is grammar-dependent rather than context-dependent, even though gradient responses vary with respect to contexts. The current proposal points to a different view of the phenomena whereby it is the range of contexts that licenses the well-formedness of an utterance that determines whether it is more, or less, acceptable than another. Although it is unlikely to be possible to determine the gradient acceptability of some string a priori, given the importance of context, it may be possible that a probabilistic view of context will give us some way to account for linguistic data in a more realistic and appropriate fashion. The empirical tests remain to be done, but the implications of the current proposals are clear and leave the way open for the definition of a unitary (if liberal) grammar with context being exploited to explain variant grammaticality patterns.

6 Conclusion

In this paper, we have put forward the hypothesis that fragments of language, as exhibited in ellipsis and dialogue generally, require a structural concept of context for interpretation. With the DS commitment to articulating concepts of structural underspecification and update in both parsing and generation, it was argued that defining a concept of context is essential to defining wellformedness for such fragments. Furthermore, defining well-formedness with respect to context enables more fine-grained concepts of wellformedness. Context dependence and the dynamics of its update, we thus argue is central, not only to semantic interpretation, but also to the syntactic process. The new research perspective that these results impose is to articulate grammar formalisms for natural language whose overall architecture is no longer defined as inducing phonology-syntax-semantics triples independent of any processing dynamics, i.e inducing a set of (interpreted) structured strings. Rather we seek to articulate formalisms for natural languages that express the intrinsically procedural nature of natural language. With this shift in perspective, a grammar formalism also induces a set of strings, but it does so by defining mappings from word sequences onto mappings from representations to representations—mappings, that is, from context structures onto representations of content attributable to the string given that context—hence by definition an update mechanism. Such a view, we contend, leads to new insights into the nature of natural language and provides solutions to linguistic problems long considered recalcitrant to theoretical explanation.

References

Blackburn, P., & Meyer-Viol, W. (1994). Linguistics, logic and finite trees. *Bulletin of the IGPL, 2*, 3–31.

Cann, R., Kempson, R., & Marten, L. (2005). *The dynamics of language, No. 35 in Syntax and Semantics*. Oxford: Elsevier.

Cann, R., Marten, L., Otsuka, M., & Swinburne, D. (2004). On the left and on the right. In D. Adger, C. De Cat, & G. Tsoulas (Eds.), *Peripheries* (pp. 19–47). Dordrecht: Kluwer.

Chomsky, N. (1981). *Lectures on government and binding*. Dordrecht: Foris.

Dalrymple, M., Shieber, S. M., & Pereira, F. C. N. (1991). Ellipsis and higher-order unification. *Linguistics and Philosophy, 14*(4), 399–452.

Fernando, T. (2002). Three processes in natural language interpretation. In W. Sieg, R. Sommer, & C. Talcott (Eds.), *Reflections on the foundations of mathematics: Essays in honor of Solomon Feferman* (pp. 208–227). Natick, Mass: Association for Symbolic Logic.

Fiengo, R., & May, R. (1994). *Indices and identity*. Cambridge, Mass.: MIT Press.

Ginzburg, J., & Cooper, R. (2004). Clarification, ellipsis, and the nature of contextual updates in dialogue. *Linguistics and Philosophy, 27*(3), 297–365.

Ginzburg, J., Gregory, H., & Lappin, S. (2001). SHARDS: Fragment resolution in dialogue. In H. Bunt, I. van der Sluis, & E. Thijsse (Eds.), *Proceedings of the 4th International Workshop on Computational Semantics (IWCS-4)* (pp. 156–172). Tilburg: Tilburg University.

Hankamer, J., & Sag, I. (1976). Deep and surface Anaphora. *Linguistic Inquiry, 7*, 391–428.

Hardt, D. (1999). Dynamic interpretation of verb phrase ellipsis. *Linguistics and Philosophy, 22*, 187–221.

Heim, I. (1982). The semantics of definite and indefinite noun phrases. Ph.D. thesis, University of Massachusetts at Amherst.

Hornstein, N. (1995). *Logical form*. Cambridge: MIT Press.

Kamp, H., & Reyle, U. (1993). *From discourse to logic*. Kluwer Academic Publishers.

Kaplan, R., & Zaenen, A. (1989). Long-distance dependencies, constituent structure, and functional uncertainty. In M. Baltin, & A. Kroch (Eds.), *Alternative conceptions of phrase structure* (pp. 17–42). Chicago: University of Chicago Press.

Keller, F., & Sorace, A. (2005). Gradience in linguistic data. *Lingua, 115*, 1397–1524.

Kempson, R., Meyer-Viol, W., & Gabbay, D. (1999). VP ellipsis: toward a dynamic structural account. In S. Lappin, & E. Benmamoun (Eds.), *Fragments: Studies in Ellipsis and Gapping* (pp. 227–290). Oxford University Press.

Kempson, R., Meyer-Viol, W., & Gabbay, D. (2001). *Dynamic syntax: The flow of language understanding*. Blackwell.

Lappin, S. (1996). The interpretation of ellipsis. In S. Lappin (Ed.), *The handbook of contemporary semantic theory* (pp. 145–176). Oxford: Blackwell.

Marcus, M. (1987). Deterministic parsing and description theory. In P. Whitelock, M. Wood, H. Somers, R. Johnson, & P. Bennett (Eds.), *Linguistic theory and computer applications* (pp. 69–112). New York: Academic Press.

Merchant, J. (2001). *The syntax of silence*. Oxford: Oxford University Press.

Otsuka, M., & Purver, M. (2003). Incremental generation by incremental parsing. In *Proceedings of the 6th CLUK Colloquium* (pp. 93–100). CLUK: Edinburgh.

Pickering, M., & Garrod, S. (2004). Toward a mechanistic psychology of dialogue. *Behavioral and Brain Sciences, 27*, 169–226.

Piwek, P. (1998). Logic, information and conversation. Ph.D. thesis, Tilburg University.

Pulman, S. (2000). Bidirectional contextual resolution. *Computational Linguistics, 26*(4), 497–538.

Purver, M., Cann, R., & Kempson, R. (2006). Grammars as parsers: Meeting the dialogue challenge. *Research on Language and Computation* (forthcoming).

Ranta, A. (1994). *Type-theoretical grammar*. Oxford University Press.

Reinhart, T. (1991). Elliptic conjunctions—Non-quantificational LF. In A. Kasher (Ed.), *The Chomskian turn* (pp. 360–84). Blackwell.

Schlangen, D. (2003). A coherence-based approach to the interpretation of non-sentential utterances in dialogue. Ph.D. thesis, University of Edinburgh.

Stainton, R. (2004). The pragmatics of non-sentences. In L. Horn, & G. Ward (Eds.), *The handbook of pragmatics* (pp. 266–287). Blackwell.

Stainton, R. (2006). *Words and thoughts*. Oxford: Oxford University Press.

Stanley, J. (2000). Context and logical form. *Linguistics and Philosophy, 23*, 391–434.