

Chapter 43

LFG and Tree-Adjoining Grammar

Jamie Y. Findlay

University of Oslo

This chapter gives an introduction to Tree-Adjoining Grammar (TAG) and draws some comparisons with Lexical Functional Grammar (LFG). It is primarily aimed at those familiar with LFG who are looking to learn about TAG and see where the two formalisms differ/overlap, but the comparisons will also be of interest to those coming from a TAG background. After introducing TAG, the chapter considers questions of generative capacity, lexicalisation, and the factoring of redundancies from grammars. It then concludes by illustrating the potential for combining LFG and TAG, and discusses the theoretical implications of doing so.

1 Introduction and roadmap

The purpose of this chapter is to give an introduction to some of the properties of Tree-Adjoining Grammar (TAG: Joshi et al. 1975, Joshi & Schabes 1997, Abeillé & Rambow 2000, Joshi 2005; Kallmeyer 2010: ch. 4) and to draw some comparisons with Lexical Functional Grammar (LFG: Kaplan & Bresnan 1982, Bresnan et al. 2016, Dalrymple et al. 2019; Belyaev 2023b [this volume]). It is primarily aimed at those familiar with LFG who are looking to learn about TAG and see where the two formalisms differ/overlap, but the comparisons will also be of interest to those coming from a TAG background (although details of the LFG formalism will not be covered – the interested reader is directed to the references above and other chapters in this volume).

A TAG is a mathematical formalism for describing a set of trees, just as a context-free grammar (CFG) is a mathematical formalism for describing a set of strings. Unlike a CFG, which generates or recognises a string by repeatedly rewriting symbols until the target string is produced, a TAG does so by combining members from a starting set of ELEMENTARY TREES using two operations



(called *SUBSTITUTION* and *ADJUNCTION*) until a tree whose yield is the target string is produced. This gives TAGs a greater generative capacity than CFGs, which is why they are of particular interest to researchers in natural language syntax: since Shieber (1985), we have known that the complexity of natural language syntax exceeds the context-free space which CFGs are capable of describing. In Section 2, I will briefly discuss this finding and the choice it has forced modern linguistic theories to make regarding their formal foundations. Section 4 delves more deeply into the generative power of TAG and compares it to that of LFG.

When TAG is used in a linguistic capacity, a number of properties are generally added to the basic formalism in order to better align it with certain theoretical assumptions and to enable more natural or transparent analyses of particular grammatical phenomena (e.g. the inclusion of feature structures on nodes to facilitate an analysis of agreement). In Section 3, I introduce some of the details of the TAG formalism along with these linguistically-motivated theoretical assumptions.

One important property generally assumed in linguistic applications of TAG is known as *LEXICALISATION*, the property whereby each of the basic structures of a grammar is associated with a single lexical item. Lexicalised grammars purportedly have a number of desirable traits, and I discuss lexicalisation in more detail in Section 5. This property sets TAG apart from CFG-based formalisms such as LFG, since the latter cannot in general be lexicalised – a perhaps surprising result given the lexical focus of LFG.

Section 6 briefly compares the TAG and LFG approaches to the factoring out of redundancies from grammars. TAG makes use of a so-called *METAGRAMMAR*, a formal system used to produce grammars, which can capture high-level generalisations and make grammar engineering easier. LFG uses *TEMPLATES*, which are part of the grammar proper, and allow pieces of linguistic description to be given names and reused.

Lastly, Section 7 considers the possibility of incorporating a TAG into the LFG architecture, replacing the standard CFG-based description of c-structure. This offers fertile new analytical possibilities, and has pleasing consequences for the descriptive power of LFG more generally, since it allows templates to be extended to the domain of phrase structure, opening the door to a fully constructional LFG.

2 Moving beyond context-free grammars

Context-free grammars have played (and continue to play) a major role in the development of syntactic theory at least since their formal elaboration in the

1950s (Chomsky 1956), with their conceptual roots going back even further (at least to e.g. Harris 1946, Wells 1947). However, there was also always a sneaking suspicion that natural language syntax was formally more complex than CFGs could describe. Nevertheless, by the early '80s there had been no successful proof of this fact (Pullum & Gazdar 1982). Bresnan et al. (1982) demonstrated that the presence of cross-serial dependencies means that the dependency structure of Dutch requires more than context-free power to describe, although owing to the lack of morphological marking of such dependencies, the *string* language of Dutch remains context free. It turns out, however, that Swiss German exhibits the same cross-serial dependencies, but its nouns are case-marked, and since different verbs can assign different cases to their objects, this means that the dependencies show up in the string language as well. Thus, greater-than-context-free power is definitely needed to describe Swiss German (Shieber 1985). Since there is no reason to suspect that speakers of Swiss German are biologically distinct from speakers of other languages, or that some people would be intrinsically unable to learn Swiss German as a first language, this means that the human language faculty generally allows for languages which require greater than context-free power to describe, and so CFGs alone are inadequate as the basis of a grammatical formalism.

Given this fact, there are two different kinds of response for the syntactic theorist. We can either

1. replace the CFG with something more powerful; or
2. beef up the CFG with something extra, so that the *combination* is more powerful.

Chomskyan generative grammar had already taken the second approach from the start: the addition of transformations to a CFG base pushes the formalism well beyond context-freeness, into the space of Type-0, unrestricted grammars (Peters & Ritchie 1973). LFG similarly adds something extra to the CFG component: in this case, a separate level of representation, f(unctional)-structure; the combination of the two takes the formalism as a whole at least into the Type-1, context-sensitive space (Berwick 1982 – although see Section 4).

However, in order to account for cross-serial dependencies, we do not need a full-blown context-sensitive (or even more powerful) grammar. Instead, we only need a more modest MILDLY CONTEXT-SENSITIVE grammar (Joshi 1985). Such grammars have the useful property, shared with context-free grammars, of being parsable in polynomial time, unlike the (worst case) exponential parsing time

of context-sensitive grammars (Joshi & Yokomori 1983, Vijay-Shanker & Joshi 1985), though they still go beyond the expressive power of CFGs in permitting the description of cross-serial dependencies. Responses to the challenge of the non-context-freeness of natural language which take approach number 1 above, and replace the CFG wholesale, tend to do so with a formalism which is explicitly mildly context sensitive, therefore, rather than anything more powerful. One example of this is Combinatory Categorical Grammar (CCG: Steedman 1987, 2000); another is TAG.¹

3 An introduction to TAG

A TAG is a tree rewriting system which consists of a set of ELEMENTARY TREES and the two operations of SUBSTITUTION and ADJUNCTION for combining them. Substitution simply inserts one tree at the frontier of another (at a non-terminal node), while adjunction inserts a tree *inside* another, attaching it at a non-frontier node (more formal definitions of these processes will be given below).²

Most linguistic work in TAG now assumes a LEXICALISED version, LTAG (Schabes et al. 1988), in which each tree is anchored by, i.e. has as its terminal node(s), a single lexical item (which may still consist of several words, as in the case of phrasal verbs, idioms, etc.). Similarly, while trees in a TAG (*qua* mathematical formalism) can be of any size, in linguistic applications the general principle applied is that trees should correspond to the EXTENDED MAXIMAL PROJECTION (Grimshaw 2000, 2005) of a lexical item, i.e. the syntactic projection which includes all functional heads and the full argument structure of the item. Some examples of elementary trees matching these restrictions are given in Table 1. In this chapter, I will use “TAG” to refer specifically to this sub-class of lexicalised,

¹In fact, things are a little more complex than this. In some definitions of mild context sensitivity (e.g. Kallmeyer 2010: 23–24), “permutation-complete” languages like MIX (the language consisting of the subset of $\{a, b, c\}^*$ with an equal number of *as*, *bs*, and *cs*; see Bach 1981) are included (Salvati 2015, Nederhof 2016); in others (e.g. Joshi et al. 1991), they are not (Kanazawa & Salvati 2012). TAG and CCG are in the class which does not contain MIX, and so Steedman (2019: 415) suggests they should be called SLIGHTLY NON-CONTEXT-FREE to distinguish them from the larger class which does contain MIX. Nothing in this chapter hinges on this distinction, so I continue to use the more traditional “mildly context sensitive”, without taking a position on whether or not this refers to the class of languages containing MIX or not.

²In the original formulation (Joshi et al. 1975), TAG only has one combining operation – adjunction. The addition of substitution, however, improves the descriptive capabilities of the framework, making it easier to use for linguistic purposes, while leaving its formal expressive power the same, since adjunction can be used to simulate substitution (Abeillé 1988: 7). In this chapter I will therefore continue to assume that both operations are used.

Table 1: Some elementary trees

Initial trees		Auxiliary trees	
NP	S	VP	S
	/ \	/ \	/ \
N	NP↓ VP	AdvP VP*	NP↓ VP
	/ \		/ \
Benjamin	V NP↓	Adv	V S*
	loves	really	thinks

linguistically-constrained TAG rather than merely to the mathematical formalism, except where otherwise indicated.

Elementary trees come in two types, illustrated in Table 1. An INITIAL TREE is one where all of the frontier nodes are either terminals or non-terminals marked as SUBSTITUTION SITES using the down arrow (\downarrow).³ Substitution sites indicate arguments of a predicate. An AUXILIARY TREE is like an initial tree except that one of the frontier nodes, called the FOOT node, shares the same label as the root, and is marked with an asterisk (*). Auxiliary trees are combined with other trees via adjunction, to be described below. When two elementary trees have been combined, we have a DERIVED TREE, which can then be further manipulated just like an elementary tree.

In the next two subsections, I introduce the two operations used to combine trees in TAG: substitution and adjunction.

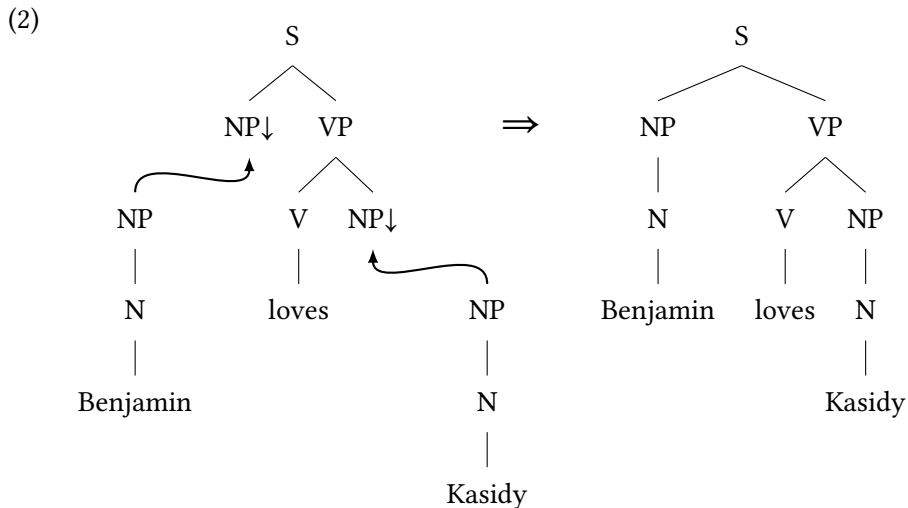

3.1 Substitution

When the root of a tree has a label which matches that of a non-terminal frontier node in another tree, the first tree can be inserted at that frontier node in the second; this is called substitution, and is the process normally used to combine a predicate and its arguments. Example (1) shows this process schematically, while

³In fact, every non-terminal frontier node can serve as a substitution site, so the \downarrow -annotation is formally redundant. Nevertheless, it is often included in expository text (if not in computational implementation) to make it clear at a glance that a tree does not represent a fully completed derivation.

(2) gives a linguistic example involving two cases of substitution: the derivation for *Benjamin loves Kasidy*.

(1) **Substitution** (after Abeillé & Rambow 2000: 5)




A tree rewriting grammar which makes use of substitution alone is called a **TREE SUBSTITUTION GRAMMAR (TSG)**, and is at least weakly equivalent to a CFG – that is, such grammars describe the same set of string languages (weak equivalence), although there are some tree languages which can be described by a TSG for which an equivalent CFG does not exist (so strong equivalence is not guaranteed).⁴ This is easy to see if we imagine converting a CFG into a TSG: all we do


⁴Any CFG can easily be converted into a TSG by simply turning each phrase-structure rule into a tree rooted in the left-hand symbol with the right-hand symbols as daughters, as will be illustrated in the text. But to convert a TSG into a CFG, it may be necessary to relabel some nodes, since the dependency between a mother and its daughters may be tree-specific, and so not hold generally (for example, it might be the case that only trees anchored by transitive verbs have a VP dominating both a V and an NP node) – so the CFG might have to have more non-terminal symbols than the TSG (e.g. VP_{trans} and $VP_{intrans}$ instead of just VP).

is replace each phrase-structure rule with the equivalent tree which it describes (cf. McCawley's 1968 conception of phrase-structure rules as node admissibility conditions, i.e. descriptions of trees). For example, the CFG in (3) corresponds to the TSG in (4):

- (3) S → NP VP
 NP → Miles
 VP → V
 V → sighs



Although TSGs and CFGs are formally very close (at least weakly equivalent), there is an important theoretical difference between them: TSGs have an EXTENDED DOMAIN OF LOCALITY with respect to CFGs. Every rule in a CFG describes a tree of depth 1, but trees in a TSG can be of arbitrarily large size, which means that certain grammatical dependencies, like agreement or extraction, can be described locally in a TSG (i.e. in a single elementary tree) when they cannot be in a CFG (i.e. they cannot be described in a single rule). For example, the TSG in (5) is equivalent to that in (4), except that now the verb and its subject are in the same elementary tree, and so the agreement relationship between the two could be described locally.




Many possibilities exist as to *how* to describe such a dependency – for example, by using complex categories in the style of GPSG (on which see Gazdar et al. 1985), or by using feature structures (on which see Section 3.3); but the point is that however one chooses to represent this relationship, it can be described locally in a TSG when it can only be described indirectly in a CFG, e.g. by percolating features up from V to VP, thus making them visible to the S → NP VP rule which introduces the subject.

3.2 Adjunction


If we add adjunction, a second type of combining operation, to a TSG, we obtain a TAG. While substitution allows a tree to be inserted at a frontier node of another tree, adjunction allows insertion at a *non*-frontier node: the adjoining tree expands the target node around itself. A tree which adjoins into another tree, called an auxiliary tree, must therefore have at least one frontier node with the same label as its root – this is called the foot node. The process of adjunction is represented schematically in (6):

(6) **Adjunction** (after Abeillé & Rambow 2000: 9)

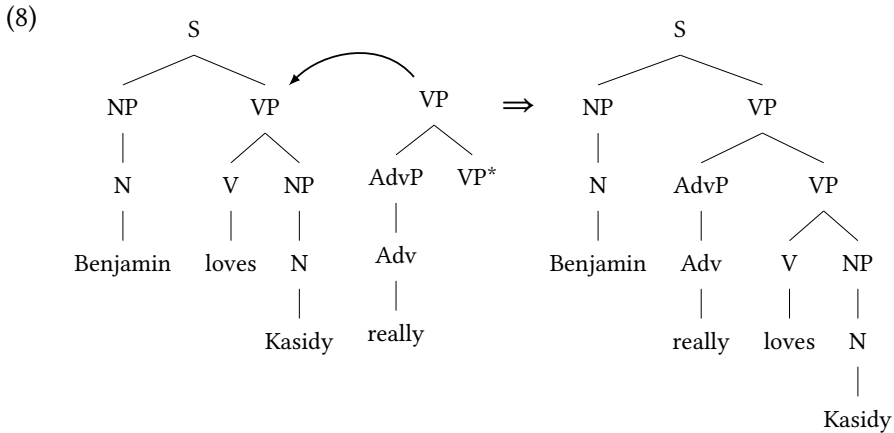


Because of the requirement that the root and foot of an auxiliary tree have the same label, such trees can be seen as factoring recursion out of the grammar. Rather than having a cyclic path through the rewrite rules (as in a CFG), we have a tree which directly encodes such a cycle (in (6), an A contained within an A), which can then be added into a structure via adjunction. For this reason, such auxiliary trees are used to model the recursive aspects of natural language syntax – most notably modification and sentential embedding.


Modifiers such as adjectives and adverbs, but also e.g. relative clauses, are represented as auxiliary trees. For example, *really* is a VP-adverb which appears to the left of the VP it modifies, and so is represented by the tree in (7):



To see this in use, consider the derivation for the sentence *Benjamin really loves Kasidy*: after the substitutions shown above in (2) to generate *Benjamin loves Kassidy*, we can then adjoin the tree from (7) at the VP node in the clause, as in (8):




Of course, such a process can be repeated indefinitely many times, since there is always still a tree-internal VP node available to be adjoined to, as in (9):

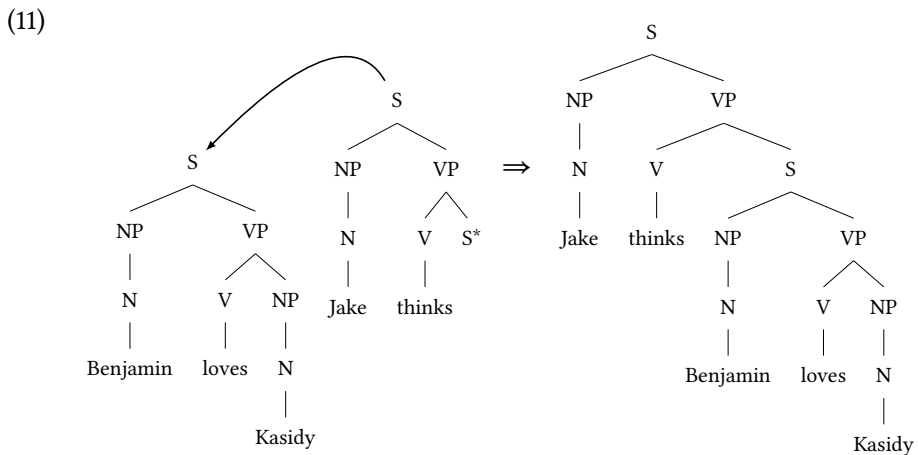


This accounts for the iterability of modifiers like *really*.

Another, perhaps more theoretically interesting, area of recursion in the grammar is in the domain of subordination, i.e. sentential embedding. Verbs which take sentential complements are represented as auxiliary trees, as in (10), for example:




Notice that this means that sentential arguments are treated rather differently from other arguments in TAG: while arguments are normally combined with their governors by means of the former being substituted into the latter, sentential arguments combine with their governors by means of the latter being adjoined into the former – this is shown in (11):



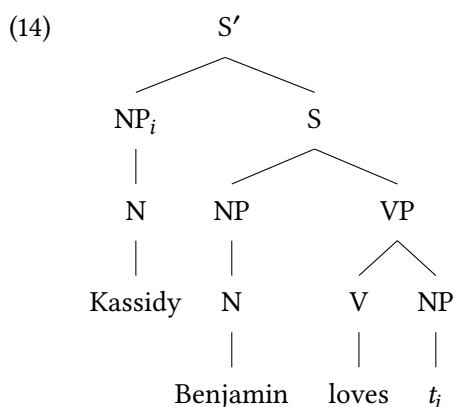
For simple declarative sentences this is a rather unnecessary complication, since the same effect could be achieved by making the foot node of the sentential-embedding verb a substitution site instead. However, the factoring of recursion into auxiliary trees interacts with another TAG principle – the local representation of syntactic dependencies. Owing to their extended domain of locality, it is possible to represent many kinds of syntactic dependencies locally (i.e. in a single elementary structure) in TAG that would require some additional mechanism in other frameworks. This principle extends to filler-gap relations as well, such as that between a fronted focus phrase and its verbal governor, as in (12):

(12) *Kassidy* Benjamin loves (whereas Kira he merely likes).

The tree in (13) represents the appropriate form of the verb *loves*, with its object extracted:⁵



Through substitution alone, this can be used to derive (14):



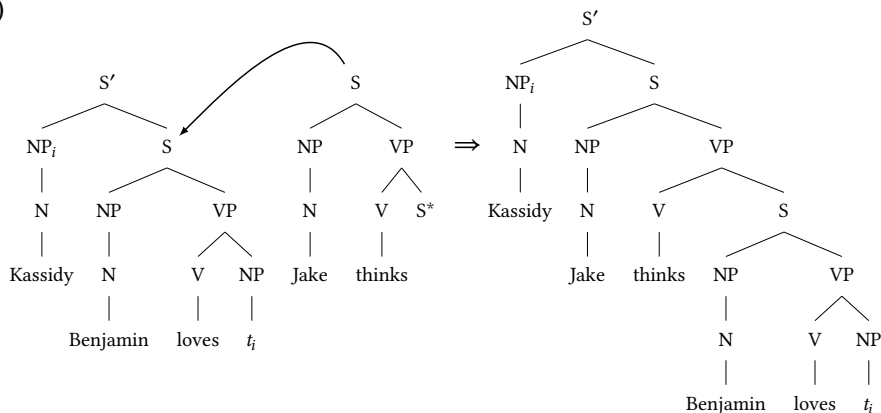
But of course the distance between the fronted phrase and the gap can span multiple clauses, and can be arbitrarily large, as shown in (15):

(15) *Kassidy* [Jadzia knows [Jake thinks ... [Benjamin loves]]].

⁵The use of a trace in object position here is not an essential part of the TAG analysis, though in practice it is common. One reason for this, as argued for by Kroch (1987) and Kroch & Joshi (1985), is that empty elements allow for easier specification of some constraints on extraction in terms of the topology of trees, rather than necessitating additional mechanisms, like functional uncertainty and off-path constraints in LFG. A reviewer points out that traces are also useful in the metagrammar (see Section 6 on this concept), since they allow tree fragments to be reused more easily.

Since sentential embedding verbs are treated as auxiliary trees, this poses no problem – they are adjoined to the internal S node, and thus extend the distance between the gap and the filler:

(16)



What is more, this can clearly be repeated: other trees can be adjoined at the topmost S node, further increasing the distance between the filler and the gap. Thus, a potentially quite radically non-local dependency, between the fronted expression and its governing verb, can be expressed locally in the grammar, in a single elementary tree, because the operation of adjunction allows for the distance between nodes in a tree to grow over the course of a derivation. This same process can be applied to other kinds of filler-gap dependencies, such as *wh*-questions and relative clauses, though for ease of exposition I have chosen not to illustrate these here (since for these we must also account for things like subject-auxiliary inversion and *do*-support).⁶ The TAG approach is in contrast to that of many other syntactic theories which instead derive or infer the relation between filler and gap via some additional syntactic mechanism, be that movement or, in the case of LFG, a functional uncertainty path.

3.3 Expressing constraints

Adjunction of an auxiliary tree, which has a root and foot node with the same label, captures the effects of recursion in other formalisms. However, once an auxiliary tree has been adjoined in, there will be two nodes with the same label where there was previously just one. This means that if we adjoin the same

⁶The interested reader should consult the detailed analyses of these and other constructions in English provided by the XTAG project (XTAG Research Group 2001).

tree again (e.g. as in (9), where we adjoin *really* twice), there are two distinct possible targets (and after that adjunction there will be three, etc.). This has the potential to dramatically complicate parsing, since there will be multiple distinct possible derivations for the same tree (without there also being a genuine ambiguity of interpretation), and so we would like a means of controlling where adjunction takes place. TAG originally did this by using local constraints on adjunction (Joshi 1987: 100ff.), annotations added to the nodes of elementary trees indicating which auxiliary trees can be adjoined there. If the list of adjoiners is empty, we have a NULL ADJUNCTION (NA) constraint, which prohibits adjunction at the node. If the list is non-empty, then we have a SELECTIVE ADJUNCTION (SA) constraint, which limits the trees which can adjoin. There are also OBLIGATORY ADJUNCTION (OA) constraints, which are like SA constraints except that one of the listed trees *must* be adjoined at the annotated node. In classic TAG, this is achieved simply by a diacritic indicating that the constraint is an OA one rather than an SA one. We will see below how this can be achieved in a less stipulative way by making use of feature structures.


Using these constraints, we can avoid having multiple possible parses for sentences by marking the foot node of auxiliary trees with an NA constraint (as is done in the XTAG grammar of English, for example – XTAG Research Group 2001). This means we do not add an extra potential target for adjunction each time such a tree is adjoined in, since only the root of the auxiliary tree is available for further adjunction.

In addition to this practical motivation, adjunction constraints play a crucial theoretical role: they are a vital part of what makes TAG mildly context sensitive. Without adjunction constraints, the formalism is still more powerful than a CFG, but there are several mildly context-sensitive languages which it cannot express, such as the copy language $\{ww \mid w \in \Sigma^*\}$, or the language $\{a^n b^n c^n \mid n \geq 0\}$, also called COUNT-3 (Kallmeyer 2010: 27, 58; we return to COUNT-3 in Section 4).

Let us now consider an example illustrating the linguistic utility of selective and obligatory adjunction constraints. Vijay-Shanker (1987: 134–135) considers non-finite sentential complements such as (17):

(17) John tried [PRO to leave].

Assuming the subordinate clause has the tree in (18), our analysis needs to do two things: ensure that such clauses cannot appear on their own as full sentences – as illustrated in (19) – and ensure that they can be embedded only under verbs that select for infinitival forms – as illustrated in (20).




(19) *To leave.

- (20) a. John tried to leave.
 b. *John imagined to leave.

In other words, *something* must be adjoined into the root node S (an OA constraint), and that something must only be a sentential embedding verb that selects for a non-finite complement clause (an SA constraint).


Originally, these constraints were simply seen as listings of (permitted/required) auxiliary trees, but this is not particularly linguistically illuminating, and also difficult to maintain for a grammar writer. This is remedied in later TAG work through the use of feature structures. It is common to associate nodes with feature structures in CFG-based grammars (e.g. in GPSG – Gazdar et al. 1985) in order to represent grammatical features such as case, number, tense, etc. Indeed, this is what LFG’s f-structures do too (although there multiple structures from different nodes are merged into one). However, in a TAG, we cannot guarantee the integrity of each node in the tree: through adjunction, it may be split up into two nodes, corresponding to the root and foot nodes of an auxiliary tree. For this reason, in feature structure-based TAG (FTAG: Vijay-Shanker 1987: ch. 5; Vijay-Shanker & Joshi 1988), each node is associated with a *pair* of feature structures, called the *TOP* and *BOTTOM* feature structures. The top features refer to the relation of the node to its siblings and its ancestors, i.e. the view from above the node in a tree. The bottom features refer to its relation to its descendants, i.e. the view from below (Vijay-Shanker 1987: 129). Ultimately, the top and bottom features of a node must unify, to give a single description of the properties of that node. However, during the course of a derivation, adjunction may split up the node so that it is now two nodes instead; in that case, its top features will be unified with the top features of the root of the auxiliary tree involved, and its bottom features will be unified with the bottom features of the auxiliary tree’s foot node. This is shown schematically in (21):

(21) Adjunction in FTAG (after Vijay-Shanker 1987: 130)




Substitution is simpler, since the root node of the substituted tree is simply identified with the substitution site, and so both top and bottom feature structures unify, as shown in (22):⁷

(22) Substitution in FTAG





These feature structures can be used to enforce various linguistic constraints. For example, we can enforce subject agreement in initial trees for verbs by specifying number and person features on the subject NP position. More interestingly, we can use features to account for the constraints on adjunction discussed above. Because the features associated with whatever tree is adjoined at a node have to unify appropriately with its top and bottom features, we can control which trees are compatible by giving them (mis)matching features which make unification possible or not. What is more, we can give a more principled account of obligatory adjunction constraints by making the top and bottom features of a particular node incompatible with one another. This means that unless adjunction takes place and the node is split up, unification will be impossible, and the derivation will fail. Returning to our example from above, (23) shows the tree from (18) with two added feature annotations:

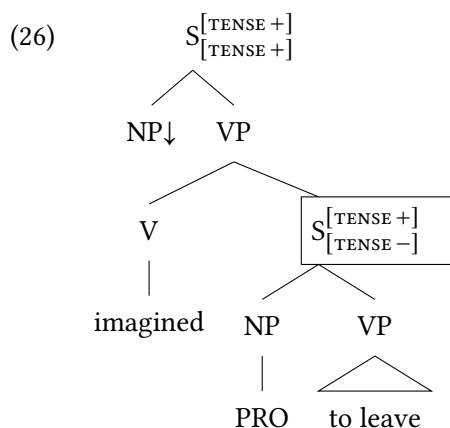
⁷The diagram in (22) follows Vijay-Shanker's (1987) original formulation, where substitution sites also contain bottom features. In much other work using FTAG, this is not the case, so b_n in (22) would be absent, and the final bottom features of A in the derived tree would simply be b_m (XTAG Research Group 2001: 13). This is of course equivalent to (22) with b_n instantiated as the empty feature structure.



Since these features are incompatible and cannot unify, we have achieved the first of our goals, which is to ensure that this tree cannot appear on its own – i.e., to implement an OA constraint. Owing to the feature mismatch, this tree is illicit unless something adjoins to the root node. To achieve the SA constraint, we need to consider the elementary trees of verbs like *tried* and *imagined*. In (24) and (25) we present them with just the relevant features added:



The difference between the two verbs is that *tried* selects for a non-finite, untensed, complement clause, whereas *imagined* selects for a tensed one – this is indicated by the top features on their foot nodes. If we attempt to adjoin *imagined* into the tree for the subordinate clause in (23), then we end up with mismatching features on the foot node, which means they cannot unify, and the tree remains illicit, as shown in (26):




If we adjoin the tree for *tried*, however, then there is no mismatch, and the derivation succeeds.

If we allow a fully-fledged unification-based feature system in FTAG, with recursive feature structures of potentially unbounded size, then FTAG becomes undecidable (Vijay-Shanker 1987: 155f.). This is a very bad result given the emphasis that TAG places on tractable, polynomial parsing. For this reason, the feature structures in FTAG are more restricted, and do not permit recursion/re-entrancy, which makes them quite unlike LFG's f-structures.


3.4 Derivation trees and dependencies

In a CFG as classically conceived, the familiar phrase-structure tree is in fact a representation of the derivation, i.e. of the process by which the output, namely the string, was produced. TAGs also have these DERIVATION TREES, representing the way in which trees were combined during a derivation – but, of course, in a TAG, the output of the derivation is already a tree, called the DERIVED TREE to set it apart. The derived tree represents word order, constituency, and category information, like LFG's c-structure. So what linguistic information does a TAG derivation tree encode? Since each elementary tree in a (lexicalised) TAG corresponds to a lexical item, the derivation tree actually represents relations between lexical items, and so has much in common with a dependency grammar representation of the sort illustrated by Meaning-Text Theory (Mel'čuk 1988), the more contemporary Universal Dependencies project (UD: Nivre et al. 2016), or, indeed, an LFG f-structure (on the relationship between dependency grammars and LFG, see also Haug 2023b [this volume]).

A derivation tree for *Benjamin really loves Kasidy*, the derivation for which was shown in (2) and (8), is given in (27) (cf. Joshi & Schabes 1997: 74ff.):




Here nodes are labelled with the lexeme corresponding to the elementary tree in question. Whenever a tree is substituted or adjoined into another tree, it becomes its daughter in the derivation tree. The derivation tree in (27) shows that three trees, corresponding to *Benjamin*, *Kasidy*, and *really*, were combined with the tree for *love*. Each edge is also labelled, standardly with a node address which indicates where the tree was substituted or adjoined.⁸ However, we can equally well use different labels, such as assigning grammatical function names to argument positions and then treating other positions as ADJ (cf. Rambow & Joshi 1997: 175). This would give us the derivation tree in (28) instead of (27), making the parallel with dependency structures quite explicit:




Rambow & Joshi (1997) discuss the relationship between TAG and dependency grammars in more detail.

Unfortunately, the TAG treatment of sentential embedding somewhat undermines the neat parallel between derivation trees and dependency structures (Rambow et al. 1995, 2001). Recall that arguments are normally substituted into their governors, but that clausal complements have their governors adjoined into them. This reverses the normal dependency relations, and means that “(standard) LTAG derivation trees do *not* provide a direct representation of the dependencies between the words of the sentence, i.e., of the predicate-argument and modification structure” (Rambow et al. 2001: 117, emphasis in original). To see why this is so, consider the derivation tree for (11), *Jake thinks Benjamin loves Kasidy*:

⁸These are so-called GORN ADDRESSES (Gorn 1967). The root has the address 0 (or sometimes ϵ , i.e. the empty string), the k th child of the root (reading left-to-right) has the address k , and for all other nodes, the q th child of a node with address p has the address $p.q$ (Joshi & Schabes 1997: 75).



In (30) *think* is a dependent of *love*, because the *think* tree is adjoined into the *love* tree, when of course in any real dependency grammar the relation would be reversed:



There are technical means of handling this unhappy result (see e.g. Joshi & Vijay-Shanker 2001, Kallmeyer & Kuhlmann 2012), but it nevertheless makes the parallel with dependency structures rather less direct. All the same, we might be tempted to see the division of labour between derived trees and derivation trees in TAG as analogous to that between c-structure and f-structure in LFG, where the former represents constituency, word order, and category information, and the latter encodes a sentence's dependency structure. This is certainly true to a point, but the parallel is imperfect, because f-structure also represents other information beyond the dependency structure of a sentence – syntactic features like person, number, tense, aspect, etc., which in TAG are encoded in the feature structures associated with each node instead. Still, one thing that derivation trees and f-structures have in common is that they are both seen as the appropriate level of representation to serve as input to the semantic component of the grammar.

3.5 Semantics

There have been a variety of different proposals for interfacing TAG with a semantic theory, and space precludes a full presentation here. Nonetheless, this section gives a (superficial) overview of the relevant literature, so that the interested reader can investigate further.

An early proposal for doing semantics with TAG makes use of SYNCHRONOUS TAG (STAG: Shieber & Schabes 1990). In STAG, elementary trees from one grammar are paired with those from another, and links are established between individual nodes in those trees. Then, when adjunction or substitution applies in one grammar, it must also take place in the other, at the linked node, and using the equivalent, paired tree. By pairing a TAG grammar with a tree-based semantic representation, we can therefore implement a “rule-to-rule” approach to semantic derivation (to use Bach’s 1976 terminology).⁹ Nothing requires the paired trees to be isomorphic, so on the syntactic side it is the structure of the derivation tree, not the derived tree, which determines the meaning. Although this approach has largely fallen out of favour in TAG circles, see Nesson & Shieber (2006, 2007, 2008) for a modern revival.

Another approach which uses the derivation tree as the basis for semantic interpretation is that of Joshi & Vijay-Shanker (2001). Here elementary trees are associated with triples of semantic expressions, the first of which specifies the main variable of the predication, the second of which gives the predicate with its arguments, and the third of which specifies which argument variables are associated with which nodes in the tree. When a tree is substituted into another tree, its main variable is identified with the corresponding argument variable in the target tree’s semantics (special consideration has to be made for adjunction, as discussed above: the order of dominance in the derivation tree will be different for sentential vs. non-sentential complements – see Joshi & Vijay-Shanker 2001: 152f.). Since this makes use of a unification-based semantics, the order of combination of the elementary trees is irrelevant, and the derivation tree thus offers an appropriate level of representation, since it abstracts away from order, and simply says how and where trees were combined. This unification-based approach has been developed more recently by Laura Kallmeyer and colleagues, introducing a new focus on underspecification (Gardent & Kallmeyer 2003, Kallmeyer & Joshi 2003, Kallmeyer & Romero 2004, 2008). This has also been integrated with Frame Semantics (Kallmeyer & Osswald 2013).

In LFG, the *de facto* standard approach to the syntax-semantics interface is GLUE SEMANTICS (Asudeh 2023 [this volume]). Observing that, for example, the operation of function application as used in natural language semantics is order insensitive, and that quantifier scope ambiguities show that semantic interpretation does not (always) respect the constituent structure of a sentence (see Asudeh 2012: ch. 5), Glue rejects c-structure as the appropriate level of input to semantics, and uses (a projection of) f-structure instead (where order is irrelevant and

⁹Alternatively, by pairing two TAG grammars from different languages, we can implement a machine translation system – see Abeillé et al. (1990).

many c-structure hierarchies are collapsed). Thus, as in TAG, it is the dependency structure, not the phrasal structure, which is taken as relevant for semantic interpretation. Interestingly, however, one of the only examples of TAG theorists criticising the derivation-tree-based approach to semantic interpretation is when Glue Semantics has been combined with TAG (Frank & van Genabith 2001). Frank & van Genabith argue that the derivation tree is not suitable as the input to semantic interpretation, mostly on the basis that, as discussed above, it provides the wrong dependency structure, and they instead make use of the derived tree in their Glue-based framework.

3.6 The big picture

Linguistic theories based on TAG have two key properties (Joshi & Schabes 1997: 95f.):

1. **EXTENDED DOMAIN OF LOCALITY:** Since TAG elementary trees encompass the whole extended projection of a lexical item, dependencies which in a simple CFG-based grammar would be spread across multiple rules, e.g. agreement, can be expressed “locally” in a TAG (i.e. in the same elementary structure). This is what enables a TAG to lexicalise a CFG (see Section 5).
2. **FACTORING RECURSION FROM THE DOMAIN OF DEPENDENCIES:** Relatedly, the elementary trees are the structures over which the vast majority of dependencies are stated, and that includes filler-gap relations. Such dependencies are therefore local in nature, but can become long distance via the adjunction operation. Recursion is thereby factored out of the domain over which these dependencies are initially stated.

This approach is summed up by Bangalore & Joshi (2010: 2) in the slogan “complicate locally, simplify globally”. That is, local, elementary representations are where almost all linguistic constraints are stated, meaning that they can become quite complex, but the payoff is that the composition of elementary trees can be achieved by just two, very general, operations: substitution and adjunction. This also means that cross-linguistic variation is entirely a matter of what elementary trees a grammar contains, a position very much in keeping with what Baker (2008: 353) calls the **BORER-CHOMSKY CONJECTURE**, after Borer’s (1984) proposal and Chomsky’s (1995) later adoption of it, whereby parametric variation is restricted to the lexicon.

How does this compare with LFG? The second property certainly divides the frameworks: LFG grammars include recursive c-structure rules, and filler-gap

dependencies are expressed syntactically, not lexically. This means, moreover, that the lexicon is not the only source of complexity in LFG grammars; many constructions are analysed as instantiating complex annotated phrase structure rules (see e.g. the analysis of long-distance dependencies in Dalrymple et al. 2019: ch. 17). There is more overlap between the two frameworks when it comes to the first property. Via the parallel projection architecture (see Belyaev 2023b: sec. 5 [this volume]), LFG does obtain an extended domain of locality: for example, agreement can be encoded locally in the agreement controller’s lexical entry via the use of paths through f-structure (see Haug 2023a [this volume]). However, since c-structure is generated by a CFG, any non-local dependencies between c-structure nodes (i.e. those spanning more than one “generation” in the tree) can only be expressed indirectly via other levels of representation. That is, we have no extended domain of locality at c-structure *per se*, only parasitically via other levels. To the extent that phrasal constructions larger than a tree of depth 1 are objects we want to be able to represent in the grammar, this is a shortcoming. We return to this point in Section 7.2.


CONSTRUCTION GRAMMAR (CxG: Fillmore et al. 1988, Goldberg 1995, 2006, Kay & Fillmore 1999, Boas & Sag 2012, Hoffmann & Trousdale 2013, *i.a.*) of course considers such objects as basic to linguistic theorising, and for this reason it has been argued that TAG is a natural means of formalising CxG (Lichte & Kallmeyer 2017). For example, among the properties of constructions listed by Fillmore et al. (1988: 501), one is that they “need not be limited to a mother and her daughters, but may span wider ranges of the sentential tree” – precisely the enlarged definition of locality which a TAG provides, and which LFG denies (at least directly). TAG has a natural means of representing both “formal” and “substantive” idioms, to use Fillmore et al.’s (1988) classification: formal idioms can be included in the set of trees associated with each lexical item of the appropriate class (Lichte & Kallmeyer 2017: 208f.), and substantive idioms can be represented as elementary trees in their own right (Abeillé 1995). While LFG can quite well represent formal idioms at the more schematic end of the scale (see e.g. Asudeh et al. 2013), it struggles with substantive idioms, precisely because it lacks an extended domain of locality at c-structure (Findlay 2023: sec. 4).

4 Generative capacity

TAG was designed specifically as a formalism with only mildly context-sensitive power (in the technical sense of Joshi 1985). This means there are languages out of the reach of context-free grammars that TAGs can describe, but also that there are

languages properly considered context sensitive which TAG cannot. Such a constrained expansion into the context-sensitive space enables parsing algorithms for TAG to preserve the computationally appealing property of a polynomial run time.

As a simple demonstration of the increased generative capacity of a TAG when compared to a CFG, consider the artificial formal language $\{a^n b^n c^n \mid n \geq 0\}$, also known as COUNT-3 – that is, the language which contains all strings consisting of some number of *as* followed by the same number of *bs*, then the same number of *cs*. Partee et al. (1990: 497) demonstrate through application of the pumping lemma for context-free languages that COUNT-3 is not context free. By contrast, there is a quite straightforward TAG grammar for COUNT-3, shown in (31).¹⁰



So, we can see that TAGs are more powerful than CFGs. They are not, however, very much more powerful. There are many kinds of language which they cannot describe, including those which it has been shown can be described by similarly modest extensions to context-free grammars. One example of this is the language MIX (Bach 1981), mentioned in footnote 1, which consists of all permutations of each string in the set $\{a^n b^n c^n \mid n \geq 0\}$, i.e. any number of *as*, *bs*, and *cs*, in any order, provided there is the same number of each. Salvati (2015 – originally circulated as a technical report in 2011) showed that MIX is in the class of multiple context-free languages, where a multiple context-free grammar (MCFG) is itself a mildly context-sensitive grammar formalism, for which the parsing problem is also decidable in polynomial time. However, MIX is *not* a tree-adjoining language, as conjectured by Joshi et al. (1991) and proved by Kanazawa & Salvati (2012): so there are languages which are only slightly within the context-sensitive space and which are still not describable by a TAG. More generally, although COUNT-3 and COUNT-4 (i.e. $\{a^n b^n c^n d^n \mid n \geq 0\}$) are tree-adjoining languages, COUNT-5 is not (Joshi 1985: 223f.).

¹⁰Recall that a node annotated with “NA” bears a null adjunction constraint – see Section 3.3. As mentioned above, without adjunction constraints, TAG becomes less expressive, and cannot describe COUNT-3: see Kallmeyer (2010: 222) for a proof.

The carefully constrained computational complexity of TAG is in marked contrast to the situation in LFG (although see below for attempts to constrain the power of the LFG formalism). Whereas the class of tree-adjoining languages is equivalent to that of the mildly context-sensitive languages (or, perhaps, the slightly non-context-free languages: see fn. 1), the languages described by LFGs are equivalent to the class of recursively enumerable languages (Nakanishi et al. 1992). This has the expected deleterious effect on computational complexity: the parsing problem for LFGs is NP-complete (Berwick 1982), and so, in the worst case scenario, computationally intractable (assuming $P \neq NP$).

There have been attempts to remedy this situation, however. While the LFG formalism as a whole may be computationally very complex, some of the properties responsible for this are not relevant for the description of natural languages – this opens the possibility that the formalism could be constrained to allow tractable parsing (i.e. in polynomial time) while still preserving its usefulness as a tool for describing natural languages. Seki et al. (1993) propose one such restriction, which limits the kinds of functional annotations permitted on c-structure nodes, and the number of nodes which can correspond to a single f-structure. This successfully buys tractability for the resulting formalism, but at a heavy theoretical cost: many staple aspects of LFG analyses are no longer available, including the very common $\uparrow = \downarrow$ head-sharing annotation, or functional control equations like $(\uparrow \text{XCOMP SUBJ}) = (\uparrow \text{SUBJ})$. More recently, Wedekind & Kaplan (2020) have addressed this limitation, describing a more expressive but still tractable version of the LFG formalism, which is provably equivalent to a Linear Context-Free Rewriting System (LCFRS), and therefore in the mildly context-sensitive space. (See also Kaplan & Wedekind 2023 [this volume] and references therein for discussion of the formal and computational properties of LFG.) This approach only covers the original LFG formalism of Kaplan & Bresnan (1982), however, and it remains to be seen whether certain extensions to this basic formalism, such as functional uncertainty (Kaplan et al. 1987, Kaplan & Zaenen 1989), can be accommodated as straightforwardly in this new approach.

One point worth noting is that even in the absence of a tractable version of LFG, this contrast between TAG and LFG should not automatically be viewed as a failing on the part of the latter. In fact, it reflects a rather deep meta-theoretical question: do we want the formalism *itself* to say something interesting about the class of natural languages? The view embodied by TAG is that we should be interested in “finding a grammar formalism that, by itself, gives already a close characterization of the class of natural languages” (Kallmeyer 2010: 7).¹¹ By

¹¹This view is also shared by Combinatory Categorical Grammar (CCG: Steedman 2000) and Multiple Context Free Grammars (MCFGs: Seki et al. 1991), among others.

contrast, the view embodied by LFG is that “it is the theory that imposes the constraints, not the language in which the theory is expressed” (Pollard 1997: 9).¹² In theoretical terms, at least, it does not seem obvious that one approach is *better* than the other – they are merely different perspectives on the problem.¹³

5 Lexicalisation

I mentioned at the start of Section 3 that linguistic applications of TAG assume that the grammar is “lexicalised”. Abeillé & Rambow (2000: 7) give the following definition of this term (emphasis in original):

We will call a grammar *lexicalised* if every elementary structure is associated with exactly one lexical item (which can consist of several words), and if every lexical item of the language is associated with a finite set of elementary structures in the grammar.

In contrast to (L)TAG, LFG grammars are not in general lexicalised, which is perhaps somewhat surprising given what the “L” in “LFG” stands for. Although there is a focus in LFG on the lexicon as a richly structured repository of grammatical information, there is no requirement that this information *cannot* be expressed through non-lexical means. This section begins by sketching the potential for lexicalising CFG-based formalisms, like LFG, and then explores what the potential advantages of lexicalised grammars are.

In general, CFGs are not lexicalised. For example, the toy grammar in (3), repeated below, is not lexicalised, since the first and third rules are not associated with any lexical item – they consist purely of non-terminals.


- (3) S → NP VP
 NP → Miles
 VP → V
 V → sighs

¹²This view is also shared by Head-Driven Phrase Structure Grammar (HPSG: Pollard & Sag 1994) and Minimalism (Chomsky 1995), among others.

¹³Of course, from a more practical point of view, it matters very much whether a formalism is tractable if it is to be used in some natural language processing task. However, there is already a very successful computational implementation of LFG in the form of the Xerox Linguistic Environment (XLE: Kaplan & Newman 1997; Crouch et al. 2011), which employs various “packed computation” (Lev 2007) heuristics to ensure efficient parsing (Maxwell & Kaplan 1989, 1993, 1996). So whatever limitations may exist in principle, they do not necessarily apply in practice.

Since LFG is based on a CFG, via c-structure, LFG grammars standardly make use of many non-lexicalised rules like these, which means that LFG grammars are generally not lexicalised.

It is possible to convert a non-lexicalised grammar into a lexicalised one, but this can require a change to the formalism used. We can speak of one grammar (weakly or strongly) *LEXICALISING* another if the former is (weakly or strongly) equivalent to the latter, except that the former is lexicalised whereas the latter is not.¹⁴ For example, the Tree Substitution Grammar shown above in (5), and repeated below, *STRONGLY LEXICALISES* the grammar in (3), since each elementary object in (5) is associated with a lexical item, and the grammar describes the same string and tree language as (3).



Sometimes it is possible to use a CFG to strongly lexicalise another CFG, but it turns out that this cannot be guaranteed in principle. For, although there is a way of converting any CFG into so-called *GREIBACH NORMAL FORM* (Greibach 1965), where the right-hand side of each rule begins with a terminal symbol – thereby lexicalising the grammar – such grammars do not in general generate the same set of trees as the grammars they normalise, since they will include different (and many more) rules. That is, converting a CFG into Greibach normal form only weakly lexicalises it. The extended domain of locality of a TSG/TAG allows us to avoid this problem, however, and makes tree grammars like this “naturally” lexicalised (Schabes et al. 1988: 579). In fact, to strongly lexicalise an arbitrary CFG, we require a TAG, not simply a TSG (see Kallmeyer 2010: 22f. for a proof). And although a TSG may be sufficient to lexicalise many linguistically relevant CFGs, it places syntactically undesirable restrictions on the resulting grammar,

¹⁴Two (classes of) grammars are weakly equivalent if they describe the same (sets of) string languages (though the corresponding (sets of) tree languages may differ). They are strongly equivalent if they also describe the same (sets of) tree languages.

and so a TAG is preferable here too (Schabes et al. 1988: 579; Schabes 1990: ch. 1). But why should we care whether a grammar is lexicalised or not?

One early advantage touted for lexicalised grammars was based on parsing. In a lexicalised grammar, a given sentence can contain at most as many elementary structures as there are words in the sentence. Since each lexical item is associated with a finite number of elementary structures, this also means that the number of analyses for the sentence is finite, thus guaranteeing that the recognition problem is decidable (Schabes et al. 1988: 581f.). As Kallmeyer (2010: 21; emphasis in original) puts it, “[l]exicalized grammars are *finitely ambiguous*, i.e., no sequence of finite length can be analyzed in an infinite number of ways”. However, in practice, the dangers of non-terminating parses are virtually non-existent in sensibly-written natural-language grammars, and so this advantage is not so great as it may seem.¹⁵

A related claim is that lexicalised grammars assist parsing because “parsing need consider only those trees of the grammar that are associated with the lexical symbols in the input string” (Eisner & Satta 2000: 79f.), rather than searching the whole grammar, and so the specific words used in a sentence “help to restrict the search space during parsing” (Kallmeyer 2010: 20). Once again, however, this argument carries less practical weight than it might seem, since parsing times for TAG grammars are actually rather slow: the best parsing algorithms for TAGs have a time complexity of $\mathcal{O}(n^6)$, as opposed to $\mathcal{O}(n^3)$ in the case of CFGs, for example (Kallmeyer 2010: ch. 5).¹⁶

There are, however, more theoretical reasons to be interested in lexicalised grammars. Firstly, it is by virtue of lexicalisation that the derivation tree of a sentence corresponds to its dependency structure (Kuhlmann 2010: 4ff.), as discussed in Section 3.4. Because each elementary object in a lexicalised grammar corresponds to a lexical item, by tracking the combination of those objects we are in fact tracking the combination of lexical items. Especially given the recent interest in dependency grammars prompted by the Universal Dependencies project (Nivre et al. 2016), it is clearly advantageous if our formalism has a transparent connection to dependency structures (see also Haug 2023b [this volume] on the relationship between LFG and dependency grammars).

Secondly, a lexicalised grammar fits very well with a lexicalist view of syntactic theory. Since the 1970s (at least since the publication of Chomsky 1970), there

¹⁵My thanks to Adam Przepiórkowski and Timm Lichte for discussion of this point.

¹⁶Although this is true of TAGs in general, if our only concern is lexicalising an existing CFG-based grammar, then we could likely devise a parser specialised for TAG grammars that lexicalise CFGs which would have a complexity below $\mathcal{O}(n^6)$. My thanks to a reviewer for this observation.

has been a trend in linguistic theory towards giving lexical analyses of many phenomena which were previously treated as purely syntactic. Indeed, driven by this trend, a plethora of linguistic frameworks have emerged which very deliberately place the lexicon front and centre, treating it as a “richly structured” object, and assuming “an articulated theory of complex lexical structure” (Dalrymple 2001: 3) – this includes LFG, as well as (to a greater or lesser extent) Generalized Phrase Structure Grammar (GPSG: Gazdar et al. 1985), Head-Driven Phrase Structure Grammar (HPSG: Pollard & Sag 1994, Müller & Wechsler 2014), Combinatory Categorical Grammar (CCG: Steedman 2000), Minimalism (Chomsky 1995), and others. Such a focus on the richness of the lexicon is in stark contrast to the historically more prominent view of it as a mere “collection of the lawless”, to use Di Sciullo & Williams’s (1987: 4) term, where it is simply a repository of exceptions, “incredibly boring by its very nature”, about which “there neither can nor should be a theory” (*ibid.*: 3f.). Given that a lexicalist syntactic theory assumes a richly detailed lexicon, in its most parsimonious form this is *all* it would require, the syntactic component being encoded in the lexical entries themselves. In fact, this is just what lexicalisation provides: in TAG, for example, aside from the basic operations of adjunction and substitution, any other grammatical constraints are described in the elementary trees of lexical items; that is, in the lexicon. In a lexicalised grammar, the lexicon essentially *is* the grammar.¹⁷ This means that every language shares the same computational component, and the only differences between languages are in the lexicon (cf. the Borer-Chomsky Conjecture, mentioned above). This is unlike LFG, for example, where languages differ both in their lexica and in the set of c-structure rules they employ.

6 Factoring out redundancies

Natural language grammars involve a large amount of redundancy: for example, the TAG elementary trees for *loves* and *thinks* shown in Table 1 are identical except for their lexical anchors and for the fact that *loves* takes an NP complement where *thinks* takes an S complement. Similarly, all proper nouns will have elementary trees like *Benjamin*, and all VP adverbs will have elementary trees like *really*, except they may follow rather than precede the VP they modify (i.e. the

¹⁷Note that it is possible to collapse the lexicon/grammar distinction without also collapsing the word/phrase (or, equivalently, morphology/syntax) distinction: the processes which build word forms, i.e. the leaf nodes of elementary trees in TAG, need not be the same as those which build derived trees in the syntax. Thus, the formal language theory objections to Construction Grammar presented by Asudeh et al. (2013: 4f.) are only objections to the most radical version of the theory, and need not be taken as objections to constructional approaches generally.

order of the foot node VP* and the AdvP node may be reversed). There is less redundancy when it comes to trees in an LFG grammar, because elementary trees are broken down into smaller-scale phrase-structure rules, but there is plenty of repetition in functional descriptions, where, for example, all 3SG verbs in English will bear the same annotations describing the person and number of their subjects.

Such redundancy or repetition is unavoidable, but it brings with it two undesirable properties: firstly, from a theoretical perspective, it means that certain generalisations may not be expressed; e.g. there are things that *thinks* and *loves* have in common, such as requiring a 3SG subject, and so it is not a mere coincidence that there is overlap in their TAG elementary trees or in their LFG functional descriptions. But nowhere in either grammar is this generalisation expressed *qua* generalisation. Secondly, from a grammar engineering perspective, this kind of redundancy makes updating and extending grammars very difficult: if we change how we analyse a particular phenomenon, we have to make sure we change every instance of it in the grammar (e.g. change every transitive elementary tree); and if we introduce a new feature to deal with some new phenomenon, we have to make sure it is handled correctly in all the existing structures, by manually adapting them one by one. This is clearly likely to lead to inconsistencies and inaccuracies due to human error.

It is therefore desirable to find a means of factoring out redundancies from a grammar, and expressing the generalisations they capture in a single place. Both LFG and TAG have a means of achieving this. In TAG, it is common practice to make use of a METAGRAMMAR, essentially a grammar responsible for generating grammars, where such redundancies can be described just once. Candito (1996, 1999) was one of the first to develop such a metagrammar;¹⁸ her version describes elementary trees along three dimensions: 1) subcategorisation (i.e. how many arguments a verb selects for), including the canonical syntactic functions of the subcategorised arguments; 2) valency alternations/redistribution of syntactic functions; i.e. the actual syntactic function of the arguments; 3) the surface syntactic manifestation of these functions. Each of these dimensions is described by an inheritance hierarchy, and the classes of the metagrammar, corresponding to specific linguistic constructions, such as the English *by*-passive, inherit from

¹⁸Candito's approach was the first to make use of non-destructive inheritance hierarchies, a move which has served as the basis for more modern metagrammar implementations (such as XMG, to be introduced below), but it was not the first to tackle the question of factoring redundancies from TAG grammars. Earlier approaches (Becker 1994 and Srinivas et al. 1994), however, make use of (destructive) lexical rules, which has made them less appealing to researchers who prefer a monotonic approach (I thank a reviewer for bringing this to my attention).

one of the terminal classes in the first dimension, one of the terminal classes in the second dimension, and as many of the terminal classes in the third dimension as there are arguments to realise. Constructions can therefore be described by listing the terminal classes they inherit from each of the three dimensions, a label which Kinyon (2000) calls a *HYPERTAG* (following from the notion of *SUPERTAG* introduced by Bangalore 1997 – see also Bangalore & Joshi 2010).

The most recent implementation of the concept of metagrammar is the *EXTENSIBLE METAGRAMMAR* (XMG) of Crabbé et al. (2013). This does away with Candito’s three explicit dimensions, and instead employs a highly expressive description language that enables linguistic structures to be given a single, complex description, including multiple levels of representation (e.g. syntax and semantics). It is also designed so that it can be extended to cover new phenomena or linguistic formalisms, and so fewer theoretical assumptions are baked into the formalism. XMG makes use of an inheritance hierarchy, but a single hierarchy instead of Candito’s three: rather than taking the approach of describing default syntactic function assignments (dimension 1) and then overriding them with specific valency frames (dimension 2), which might vary, e.g. in the case of diathesis alternations, XMG makes heavy use of disjunctions between alternating descriptions, which enables such alternations to be described fully declaratively, and in just one place. For example, we can express the familiar active-passive diathesis of English as in (32), where each term in italics refers to a class in the metagrammar’s inheritance hierarchy that gives a partial description of a (sub-)tree (Crabbé et al. 2013: 616).

$$(32) \quad \textit{TransitiveDiathesis} \rightarrow \begin{aligned} & (\textit{Subject} \wedge \textit{ActiveVerbForm} \wedge \textit{Object}) \\ & \vee (\textit{Subject} \wedge \textit{PassiveVerbForm} \wedge \textit{ByObject}) \\ & \vee (\textit{Subject} \wedge \textit{PassiveVerbForm}) \end{aligned}$$

Each of the disjuncts in (32) combines these descriptions to give a partial description of a full elementary tree schema (an elementary tree minus its lexical anchor). For now I leave aside the details of how these classes actually describe trees; a simplified version of the logical description language employed will be introduced in the next section. The crucial observation here, and the move which sets XMG apart from earlier approaches to metagrammatical analysis, is that the description in (32) does not privilege one elementary tree/realisation of arguments as basic, but simply describes all possible realisations simultaneously.

The terminal classes of the metagrammar are families of trees which are then associated with lemmas, and represent all the different ways of realising that lemma’s arguments (e.g. active vs. passive, *wh*-extraction, clefting, etc.). The *TransitiveFamily* associated with a lemma like LOVE might just consist of (32),

while the *DitransitiveFamily* of a verb like GIVE might inherit from the *TransitiveFamily* but add an additional object argument:

- (33) *TransitiveFamily* → *TransitiveDiathesis*
DitransitiveFamily → *TransitiveDiathesis* \wedge *IndirectObject*

This modular and structured approach to the metagrammar means that, for instance, if the analysis of a particular phenomenon changes, we just need to modify the relevant class(es): when the grammar is compiled anew, all of the implicated elementary trees will be altered accordingly. The choice of classes can also have theoretical implications, and may shed light on important linguistic generalisations.¹⁹

Although the metagrammatical approach has been used to generate LFG grammars as well as TAGs (Clément & Kinyon 2003a,b), this is not common practice in LFG work. Rather, since redundancies in an LFG grammar are far more abundant in the functional descriptions associated with lexical entries than in phrase structure, the standard solution employed here is to make use of *TEMPLATES*, a type of macro which can be used to abbreviate pieces of functional description that are re-used across lexical entries (Dalrymple et al. 2004, Crouch et al. 2011; see also Belyaev 2023a: sec. 5.1 [this volume]). These templates can take arguments, and can also call other templates, creating a hierarchical organisation – though it should be noted that this is an *inclusion* hierarchy rather than an inheritance hierarchy, since template calls can be negated (Asudeh et al. 2013: 18f.). The semantics of template invocation (represented by prefixing the template name with a '@' symbol) is substitution: the template name is replaced by its contents. This means that a grammar without templates is extensionally equivalent to one with them, but in the latter it will be possible to express generalisations that cannot be expressed in the former.

By way of illustration, (34–35) present some templates which capture some of the same information present in the XMG classes shown above. The *TRANSITIVEDIATHESIS* template takes a predicate name as its argument, and consists of a disjunction of three other templates; it will be called by the lexical entry of any transitive verb which participates in the active/passive alternation in English. Each of the three templates it invokes provides a *PRED* value for the verb in

¹⁹Although metagrammars of this sort can certainly be useful theoretical tools, this is not to say that they are intended as models of how the human language faculty functions. As a reviewer points out, it is perhaps implausible that, in the process of language acquisition, the language learner has to recompile their entire grammar every time they make a change or add a new observation. In this regard, LFG's templates (to be introduced below), which are nothing more than abbreviations, might seem more promising as a model of the learner's competence.

question, associating it with the correct set of grammatical functions, and also provides mapping equations which link the GFs to argument positions at semantic structure, or express the fact that the argument is not syntactically realised, in the case of the short passive (this approach to mapping is described in Asudeh & Giorgolo 2012 and Findlay 2016; see also Findlay et al. 2023: sec. 6.2 [this volume]).

$$(34) \quad \text{TRANSITIVEDIATHESIS}(P) \equiv \\ @\text{ACTIVETRANSITIVE}(P) \vee @\text{BYPASSIVE}(P) \vee @\text{SHORTPASSIVE}(P)$$

$$(35) \quad \text{a. } \text{ACTIVETRANSITIVE}(P) \equiv \begin{aligned} (\uparrow \text{ PRED}) &= 'P\langle \text{SUBJ}, \text{OBJ} \rangle' \\ (\uparrow_{\sigma} \text{ ARG1}) &= (\uparrow \text{ SUBJ})_{\sigma} \\ (\uparrow_{\sigma} \text{ ARG2}) &= (\uparrow \text{ OBJ})_{\sigma} \end{aligned}$$

$$\text{b. } \text{BYPASSIVE}(P) \equiv \begin{aligned} (\uparrow \text{ PRED}) &= 'P\langle \text{SUBJ}, \text{OBL}_{\text{BY}} \rangle' \\ (\uparrow_{\sigma} \text{ ARG1}) &= (\uparrow \text{ OBL}_{\text{BY}})_{\sigma} \\ (\uparrow_{\sigma} \text{ ARG2}) &= (\uparrow \text{ SUBJ})_{\sigma} \end{aligned}$$

$$\text{c. } \text{SHORTPASSIVE}(P) \equiv \begin{aligned} (\uparrow \text{ PRED}) &= 'P\langle \text{SUBJ} \rangle' \\ (\uparrow_{\sigma} \text{ ARG1})_{\sigma^{-1}} &= \emptyset \\ (\uparrow_{\sigma} \text{ ARG2}) &= (\uparrow \text{ SUBJ})_{\sigma} \end{aligned}$$

One noteworthy difference between the use of a metagrammar and the use of templates is that the latter but not the former are part of a grammar itself. A metagrammar, as the name suggests, sits outside the grammar proper: it outputs grammars, where the elementary objects do not (necessarily) contain information about which metagrammar classes they instantiate. Templates, on the other hand, are part of the description language of the grammar, although of course they are merely names for pieces of functional description, and so have no special formal status themselves.

7 Combining LFG and TAG

Now that we have seen some of the key concepts of TAG, along with their motivations and apparent benefits, we might wonder whether LFG could also benefit from some of these boons if we were to combine the two approaches – most naturally, by using a TAG instead of a CFG to describe LFG’s c-structure. Joshi (2005: 496) described this idea as being “of great interest”, and it was previously explored by Kameyama (1986) and Burheim (1996) – but unfortunately only in unpublished work, which has proved impossible to track down. More recently,

the idea has been revived by Findlay (2017a,b, 2019). In this section, I outline two different approaches to achieving the goal of combining LFG and TAG, and discuss some of the consequences of adopting such a merger.²⁰

7.1 Two approaches

The most straightforward way of combining TAG and LFG is simply to take a TAG grammar and add appropriate LFG annotations to the elementary trees. Of course, once we have access to the whole tree, we gain a greater degree of flexibility in how we express functional annotations. Most notably, we can refer to any node in the tree directly, rather than being limited to the current node or its mother – a consequence of TAG’s extended domain of locality. For example, instead of relying on a sequence of $\uparrow = \downarrow$ annotations to pass information from a lexical item to the top of its extended projection, we can refer to the top directly. For the sake of simplicity, let us use node labels as shorthand for the nodes themselves.²¹ Then the $(\uparrow \text{ PRED}) = \text{‘love’}$ annotation on the verb *loves*, for example, could be rewritten as $(S_\phi \text{ PRED}) = \text{‘love’}$, using S_ϕ to refer to the f-structure of the clause directly, rather than indirectly via V_ϕ (the instantiation of \uparrow), which is equated with both VP_ϕ and S_ϕ . Indeed, since we can use absolute rather than relative labels for the nodes in the tree, there is no need to mark annotations actually on the tree at all; instead, we can treat lexical entries as pairs consisting of the tree on the one hand and the annotations on the other, which refer to nodes in the tree. This arguably simplifies the process of determining an f-structure from an annotated c-structure, since many identities which would normally have to be computed are instead already given in the descriptions. Table 2 shows the elementary trees from Table 1 augmented in this fashion.

The trees then combine as usual for a TAG, using the operations of substitution and adjunction, albeit understood in a particular fashion. Substitution involves identifying two nodes, so that, e.g. if the tree for *Benjamin* were substituted into the subject position of the tree for *loves*, NP and NP1 would be identified (and therefore so would their f-structures, requiring the nodes to bear compatible annotations – and thereby accounting for the agreement facts, for example). Ad-

²⁰One concern about replacing the CFG component of LFG with a more powerful TAG might be that it makes the formalism as a whole more computationally complex. However, since TAGs are strictly less powerful than LFGs (see Section 4), such a concern is ultimately baseless.

²¹Of course, in reality nodes and their labels are distinct: several nodes can bear the same label, for example (e.g. there can be more than one NP in a tree). When this happens, I follow the TAG convention of suffixing node labels with numbers (e.g. NP1 and NP2), but it should be borne in mind that this is just a representational choice, and that in reality such nodes have identical labels.

junction involves three steps: first we excise a sub-tree rooted at the adjunction site; next, we *replace* it with the adjoining auxiliary tree; finally, we unify the foot node of the auxiliary tree with the root node of the excised sub-tree it replaced.²² This way, we identify the target of adjunction with the foot of the auxiliary tree, and correctly distribute the annotations between the two “parts” of the expanded node without the need for top and bottom feature structures.^{23,24}

This first approach is much more in the spirit of TAG than of LFG, since the c-structure component is DERIVATIONAL, making use of the combining opera-

²²Note that it is particularly important in this setting that adjunction is only defined where the adjoining tree’s root and foot nodes are of the same category. In some TAG settings this would not need to be stated explicitly, depending on how adjunction is defined, but here the root of the auxiliary does not unify with anything, and so there is nothing which formally requires the root and foot nodes of such a tree to have the same category (I thank a reviewer for this observation). Allowing trees with mismatched root and foot nodes to participate in adjunction would have undesirable consequences: for example, we do not want the tree for *loves* in Table 2 to act as an NP modifier (e.g. **the Benjamin loves boy*).

²³A reviewer asks how obligatory adjunction can be implemented in this setting, since in FTAG it exploits the possibility of mismatching top and bottom features (see Section 3.3). Ultimately, the answer is that the greater expressive power of the LFG projection architecture means that the effects of obligatory adjunction constraints will be captured in different ways in different situations. Constraining equations will frequently be relevant: for example, returning to the example of *to leave* from (18), to implement the SA constraint we might specify that *tried* requires its COMP to contain the feature [FINITE –], whereas *imagined* requires it to contain [FINITE +]; if *to leave* specifies that its f-structure contains [FINITE –], then it will be compatible with the former but incompatible with the latter. If we wish to avoid *to leave* appearing on its own (i.e. we rule out fragments), we might implement a general ban on root f-structures containing [FINITE –], or we might rely on the resource sensitivity of Glue Semantics, since an infinitive alone will not permit a linear logic proof terminating in the goal type of propositions.

²⁴Findlay (2017a: 222, fn. 12) claims that we are forced to adopt the second proposal to be discussed below, using descriptions of trees, because adjunction means that the \uparrow and \downarrow chains in annotations will be disrupted. This would be true if we were forced to refer to f-structures only indirectly, via mother-daughter links, but fails to appreciate the additional freedom afforded by being able to refer to nodes absolutely, as discussed above. There is, however, a small wrinkle when it comes to verbal trees for extraction constructions (e.g. *wh*-questions): if nothing is adjoined to them, we want to unify the f-structures of the root *S'* and the *S* node it immediately dominates; but if a sentential embedding verb is adjoined there, we cannot identify the two f-structures, or else we will end up with a cyclic f-structure which is its own COMP. All this shows us though is that we have to take care when writing the functional annotations. Here, for example, we can solve the problem by actually reintroducing an element of relativity: we identify the root node’s f-structure with the f-structure of its *S* daughter (e.g. by defining a predicate $CATDAUGHTER(n, C)$ which identifies the unique daughter of node *n* which bears label *C*, and is undefined if there is none or more than one), regardless of which node that actually ends up being. I omit the formal details of how this can be achieved, since ultimately we will settle on the second approach to integrating TAG and LFG described below, but it is important to note that this first approach is not unworkable.

Table 2: Some elementary trees with associated LFG annotations

Initial trees	
$\left\langle \begin{array}{c} \text{NP} \\ \\ \text{N} \\ \\ \text{Benjamin} \end{array} \right\rangle$	$\left\langle \begin{array}{l} \text{NP}_\phi = \text{N}_\phi \\ (\text{NP}_\phi \text{ PRED}) = \text{'Benjamin'} \\ (\text{NP}_\phi \text{ NUM}) = \text{SG} \\ (\text{NP}_\phi \text{ PERS}) = 3 \end{array} \right\rangle$
$\begin{array}{c} \text{S} \\ / \quad \backslash \\ \text{NP1}\downarrow \quad \text{VP} \\ \quad \quad \quad / \quad \backslash \\ \quad \quad \quad \text{V} \quad \text{NP2}\downarrow \\ \quad \quad \quad \\ \quad \quad \quad \text{loves} \end{array}$	$\left\langle \begin{array}{l} \text{S}_\phi = \text{VP}_\phi = \text{V}_\phi \\ (\text{S}_\phi \text{ PRED}) = \text{'love'} \\ (\text{S}_\phi \text{ TENSE}) = \text{PRES} \\ (\text{S}_\phi \text{ SUBJ}) = \text{NP1}_\phi \\ (\text{S}_\phi \text{ OBJ}) = \text{NP2}_\phi \\ (\text{NP1}_\phi \text{ NUM}) = \text{SG} \\ (\text{NP1}_\phi \text{ PERS}) = 3 \end{array} \right\rangle$
Auxiliary trees	
$\begin{array}{c} \text{VP1} \\ / \quad \backslash \\ \text{AdvP} \quad \text{VP2}^* \\ \\ \text{Adv} \\ \\ \text{really} \end{array}$	$\left\langle \begin{array}{l} \text{VP1}_\phi = \text{VP2}_\phi \\ \text{AdvP}_\phi = \text{Adv}_\phi \\ (\text{AdvP}_\phi \text{ PRED}) = \text{'really'} \\ \text{AdvP}_\phi \in (\text{VP1}_\phi \text{ ADJ}) \end{array} \right\rangle$
$\begin{array}{c} \text{S1} \\ / \quad \backslash \\ \text{NP}\downarrow \quad \text{VP} \\ \quad \quad \quad / \quad \backslash \\ \quad \quad \quad \text{V} \quad \text{S2}^* \\ \quad \quad \quad \\ \quad \quad \quad \text{thinks} \end{array}$	$\left\langle \begin{array}{l} \text{S1}_\phi = \text{VP}_\phi = \text{V}_\phi \\ (\text{S1}_\phi \text{ PRED}) = \text{'think'} \\ (\text{S1}_\phi \text{ TENSE}) = \text{PRES} \\ (\text{S1}_\phi \text{ SUBJ}) = \text{NP}_\phi \\ (\text{S1}_\phi \text{ COMP}) = \text{S2}_\phi \\ (\text{NP}_\phi \text{ NUM}) = \text{SG} \\ (\text{NP}_\phi \text{ PERS}) = 3 \end{array} \right\rangle$


tions of substitution and adjunction. Let us therefore call it LFG-TAG. But as Kaplan (1995: 11) points out, this PROCEDURAL, or CONSTRUCTIVE, approach to grammatical analysis is in contrast to the DESCRIPTIVE (a.k.a. DECLARATIVE or MODEL-BASED) approach which is the “hallmark of LFG” (*ibid.*). Findlay (2019: ch. 5) therefore explores another way of combining the two frameworks which is more in keeping with the LFG spirit. In brief, we associate lexical entries with *descriptions* of trees, rather than with the trees directly, as is standard practice in metagrammars, for instance.²⁵ In the simplest cases there is a one-to-one correspondence between a description and the (minimal) tree it describes, and so we could straightforwardly translate LFG-TAG into a more LFG-like format. However, descriptions can also make use of negation, disjunction, or other operations that go beyond simple conjunction of propositions, and in this case the relation between descriptions and trees is no longer isomorphic (Kaplan 1995: 14).

²⁵The use of tree descriptions has been discussed extensively in the context of TAG – see, for instance, Vijay-Shanker (1992), Rogers & Vijay-Shanker (1994), Rambow et al. (1995, 2001), Kallmeyer (2001).

In order to add descriptions of trees to LFG lexical entries, we need a suitable language to write the descriptions in. There are a variety of different possibilities, but here we will assume a fairly simple language based on that used in XMG (Crabbé et al. 2013: 599), which will consist of the following:²⁶

- (36)
1. a set N of node variables
 2. a set P of unary labelling predicates, including all terminal and non-terminal labels
 3. the following binary predicates:
 - \rightarrow , immediate dominance (the *mother-of* relation)
 - \rightarrow^* , dominance (the transitive, reflexive closure of \rightarrow)
 - \prec , linear precedence²⁷

The tree in (37) can then be described by the set of constraints in (38):²⁸



- (38)
- | | | |
|--------------|-----------------------|-----------------|
| $S(n_1)$ | $n_1 \rightarrow n_2$ | $n_2 \prec n_3$ |
| $NP(n_2)$ | $n_1 \rightarrow n_3$ | $n_4 \prec n_5$ |
| $VP(n_3)$ | $n_3 \rightarrow n_4$ | |
| $V(n_4)$ | $n_3 \rightarrow n_5$ | |
| $NP(n_5)$ | $n_4 \rightarrow n_6$ | |
| $loves(n_6)$ | | |

However, as it stands, the constraints in (38) are too rigid. Specifically, they will not allow adjunction at the VP node, since then at least one statement in the


²⁶We will also assume that sufficient axioms are in place to ensure the usual well-formedness conditions on trees, e.g. that they are singularly rooted, that branches cannot cross, etc. Rogers (1998: 15f.) gives one such set of axioms.

²⁷Here this is to be understood as the transitive closure of immediate linear precedence, i.e. what Crabbé et al. (2013: 599) represent as \prec^+ . In other words, a node linearly precedes everything to its right, but does not linearly precede itself.

²⁸In descriptions, we will assume that all node variables are ultimately existentially bound.

description will no longer be true: if we identify the target n_3 with the root of the adjoining tree, then $n_3 \rightarrow n_4$ will no longer hold (the foot node of the auxiliary tree will dominate n_4 instead), and if we identify it with the foot node of the adjoining tree, then $n_1 \rightarrow n_3$ will not be true instead. The basic problem is that “[t]he composition operation of adjoining creates a new structure that does not maintain all of the properties that held in the original (fully specified) structures of which it is composed” (Vijay-Shanker 1992: 486). What this means is that we cannot operate with fully specified descriptions, but must make use of partial descriptions instead.

For each node where adjunction can apply, we instead describe a pair of QUASI-NODES which stand in the dominance relation (Vijay-Shanker 1992: 486ff.). That is, instead of (38), we have (40), which is represented schematically in (39) (where a dashed line represents dominance rather than immediate dominance):



- (40)
- | | | |
|--------------|-------------------------|-------------|
| $S(n_1)$ | $n_1 \rightarrow n_2$ | $n_2 < n_3$ |
| $NP(n_2)$ | $n_1 \rightarrow n_3$ | $n_5 < n_6$ |
| $VP(n_3)$ | $n_3 \rightarrow^* n_4$ | |
| $VP(n_4)$ | $n_4 \rightarrow n_5$ | |
| $V(n_5)$ | $n_4 \rightarrow n_6$ | |
| $NP(n_6)$ | $n_5 \rightarrow n_7$ | |
| $loves(n_7)$ | | |

As elsewhere in LFG, we take the solution to a set of constraints to be the *minimal* structure (or structures) which satisfies all the constraints. Since the dominance relation is reflexive, the minimal tree which satisfies (40) remains (37), i.e. one where we equate n_3 and n_4 . But, crucially, if something is adjoined here, the nodes

can come apart, with the result that n_3 is identified with the root of the auxiliary tree and n_4 with its foot node.

Now that we have a description of this tree, we can combine it with functional descriptions to form a full LFG lexical entry:²⁹

(41)	$S(n_1)$	$n_1 \rightarrow n_2$	$n_2 < n_3$	$n_{1\phi} = n_{3\phi}$
	$NP(n_2)$	$n_1 \rightarrow n_3$	$n_5 < n_6$	$n_{4\phi} = n_{5\phi}$
	$VP(n_3)$	$n_3 \rightarrow^* n_4$		$(n_{5\phi} \text{ PRED}) = \text{'love'}$
	$VP(n_4)$	$n_4 \rightarrow n_5$		$(n_{5\phi} \text{ TENSE}) = \text{PRES}$
	$V(n_5)$	$n_4 \rightarrow n_6$		$(n_{1\phi} \text{ SUBJ}) = n_{2\phi}$
	$NP(n_6)$	$n_5 \rightarrow n_7$		$(n_{4\phi} \text{ OBJ}) = n_{6\phi}$
	$\text{loves}(n_7)$			$(n_{2\phi} \text{ NUM}) = \text{SG}$
				$(n_{2\phi} \text{ PERS}) = 3$

To parse a sentence, we just collect up all of the constraints associated with each lexical item and find the minimal structures – both c-structure and f-structure – which satisfy them.

Of course, (41) is not particularly readable, so we might prefer to collect some parts of the description in various templates. For example, the tree for any transitive verb will share most of the description in (41), so we can factor out this part of the description, parametrising the only variable, namely the lexical anchor:

(42)	$\text{TRANSITIVETREE}(s, np_1, vp_1, vp_2, v, np_2, a, \text{ANCHOR}) \equiv$			
	$S(s)$	$s \rightarrow np_1$	$np_1 < vp_1$	$s_\phi = vp_{1\phi}$
	$NP(np_1)$	$s \rightarrow vp_1$	$v < np_2$	$vp_{2\phi} = v_\phi$
	$VP(vp_1)$	$vp_1 \rightarrow^* vp_2$		$(s_\phi \text{ SUBJ}) = np_{1\phi}$
	$VP(vp_2)$	$vp_2 \rightarrow v$		$(vp_{2\phi} \text{ OBJ}) = np_{2\phi}$
	$V(v)$	$vp_2 \rightarrow np_2$		
	$NP(np_2)$	$v \rightarrow a$		
	$\text{ANCHOR}(a)$			

²⁹Here I have kept to a more conservative annotation scheme than above, whereby e.g. lexical contributions are associated with the f-structure of n_5 , i.e. the V node, rather than with that of the root S. This is because adjunction may in principle alter the structure of the tree so that it is no longer the case that the f-structure of the S node is the same as the f-structure of the V node. In fact, with verbal trees like this, that will not be the case, because the only auxiliary trees which target VPs in a TAG grammar will be auxiliary verbs or adverbial modifiers, neither of which will break the link between V and S in terms of f-structure-identity. But it will, for example, be relevant for verbal trees containing extraction sites, which can be targetted by sentential embedding verbs, thereby separating the root's f-structure from the head verb's.

We have to “expose” all nodes as parameters of the template, so that they can be referred to by other constraints in the same lexical entry, thereby taking advantage of the extended domain of locality afforded by having the description of the whole tree in one place. However, since all of the parameters in (42) except the lexical anchor will simply be node variables, I propose a shorthand: when calling the template, all but the last parameter will be omitted (though, to repeat, when defining it all the parameters must be specified); if we wish to refer to any of the other parameters, we can do so by using the template name and suffixing it with the appropriate parameter.³⁰ For example, TRANSITIVETREE.*s* refers to the first parameter, a node variable which corresponds to the root node *s* in (42). With these conventions in place, we can write a more readable lexical entry for *loves* as in (43), using a LOCAL NAME, %UP, to refer to the verb’s *f*-structure – see Belyaev 2023a: sec. 3.2.5 [this volume] for the details on local names):³¹

- (43) @TRANSITIVETREE(*loves*)
 %UP = TRANSITIVETREE.*v*_φ
 (%UP PRED) = ‘love’
 (%UP TENSE) = PRES
 (%UP SUBJ NUM) = SG
 (%UP SUBJ PERS) = 3

One theoretical advantage of this approach is that we can build up trees from smaller parts by making use of nested template calls. This allows us to capture connections between phrasal configurations in a way which CFG rules do not. For example, there is no relationship between the two rules in (44), even though the latter is obviously partially described by the former:³²

³⁰This is based on the conventions of XMG for exported variables (Crabbé et al. 2013: 602–604).

³¹Here I have reverted to describing the agreement constraints on the subject via the verb’s *f*-structure rather than via the NP’s, to make this lexical entry closer to the LFG standard. But of course the option is still open to us to describe it via the tree directly, by associating TRANSITIVETREE.*np*_{1φ} with a name, e.g. %SUBJ-NP, and then declaring that (%SUBJ-NP NUM) = SG. Although these options are extensionally equivalent here, there can be theoretical/descriptive reasons to prefer one over the other. Cross-linguistically, for example, we might want to treat subject agreement as the same kind of phenomenon both in languages where phrase-structure position is a clear guide to grammatical function (like English) and in languages where it is not (like Warlpiri); so it would make sense to retain the standard LFG approach of describing agreement via *f*-structure. But in other cases it might make more sense to refer to a particular phrase-structure position, and the integration of an extended tree description into the lexical entry means we now have that choice.

³²Of course, we can use the convention of surrounding optional nodes in parentheses, and then we can express the relationship between the two rules within a single phrase-structure rule as follows:

- (44) a. VP \rightarrow V NP
 $\uparrow=\downarrow$ (\uparrow OBJ) = \downarrow
- b. VP \rightarrow V NP NP
 $\uparrow=\downarrow$ (\uparrow OBJ) = \downarrow (\uparrow OBJ $_{\theta}$) = \downarrow

On the other hand, if we have a template DITRANSITIVE TREE which calls the TRANSITIVE TREE template (as well as another template which adds a secondary object), then this containment relationship is made explicit, as shown in (45).

- (45) DITRANSITIVE TREE(ANCHOR) \equiv @TRANSITIVE TREE(ANCHOR)
 @SECONDARY OBJECT

Of course the TRANSITIVE TREE template can also be decomposed into a call of an INTRANSITIVE TREE template plus a PRIMARY OBJECT one, and so on. By continuing along these lines, we can capture all the various generalities across trees in a template inclusion hierarchy, recreating the class hierarchies of a TAG meta-grammar inside an LFG grammar itself.

7.2 Implications

Having now seen how LFG and TAG can be combined, let us consider the consequences of such a merger. There are several potential gains which such a move could bring, along with several unanswered questions which require further research.

Firstly, the second approach described above offers a pleasing harmonisation of LFG lexical entries. Standard LFG lexical entries contain descriptions of all levels of the projection architecture, but since such lexical entries are really just context-free phrase-structure rules, the description of c-structure is limited to information about the word itself and its mother. In contrast, descriptions of all other levels of structure can refer to arbitrarily distant elements (via functional uncertainty). The inclusion of tree descriptions removes this irregularity from the grammar, since now non-local elements of c-structure can also be included.

$$(i) \text{ VP } \rightarrow \text{ V NP } \left(\begin{array}{l} \text{NP} \\ (\uparrow \text{ OBJ}_{\theta}) = \downarrow \end{array} \right)$$

$$\uparrow = \downarrow \quad (\uparrow \text{ OBJ}) = \downarrow$$

However, once we move beyond simple examples like this, such an approach becomes unwieldy, with multiply nested parentheses and very complex disjunctions. Unlike the templatic approach, which provides a readable front-end to the formal complexity, and allows us to represent the relationship(s) between sub-trees in an inclusion hierarchy, this approach forces us to create fewer but more complex rules, which does nothing to aid human-readability.

Secondly, we now have the opportunity to lexicalise an LFG grammar (indeed, Findlay 2019: ch. 5 calls the description-based approach described above “Lexicalised LFG”). As outlined above, the extended domain of locality of a TAG means that all dependencies, including long-distance ones, can be encoded locally in a lexical entry. Lexicalisation seems a natural goal for a lexicalist theory like LFG, and it is perhaps lamentable that it was not possible before.

Thirdly, we can now straightforwardly account for idioms (Findlay 2019: ch. 6). These are problematic for the current leading account of constructions in LFG (Asudeh et al. 2013), since they do not simply add additional constructional meaning to existing lexical meaning, but rather *replace* the lexical meaning with another, different meaning (that is, *shooting the breeze* involves neither shooting nor a uniquely contextually salient breeze). This forces lexicalist theories like LFG to adopt an approach which treats idioms as conspiracies of independent lexical items that select for one another (see Findlay 2023: sec. 4.3). Such approaches face a host of problems, not least of which is that they singularly fail to capture our intuitions about idioms – viz. that they are “things” (as Williams 2007 puts it), and not mere epiphenomena of the grammar (see Findlay 2019: 58ff. for discussion of various other problems). But now that we can have lexical entries containing multiple, separable word forms, something which is not possible in vanilla LFG, there is no obstacle to encoding multiword expressions in a single place, thus enabling a much more satisfying analysis. Findlay (2019: ch. 3) provides detailed discussion of the need for this kind of constructional approach to idioms, as well as arguments against other types of analysis.

Fourthly, in addition to idioms, we have a straightforward account of constructional phenomena more broadly. Similar arguments can be made here about the need for constructions to have some ontological status in the theory – to be “things”.³³ Admittedly, Asudeh et al. (2013) demonstrate that we do not need to admit constructions as first-class entities in our theory in order to explain *some* kinds of constructional effects, but they only consider constructions which can be described by a single lexical entry or a single context-free phrase-structure rule, and so the constructions in question can be described in a single place. Other constructions require reference to wider spans of phrase structure, or require

³³There is suggestive psycho- and neurolinguistic evidence that the way we process language makes heavy use of prefabricated chunks (“prefabs”) (Pawley & Syder 1983, Wray 2002) and of constructions more generally (Bencini & Goldberg 2000, Kaschak & Glenberg 2000, Goldwater & Markman 2009, Pulvermüller 2010, Allen et al. 2012, Johnson & Goldberg 2012). Obviously grammatical theory need not have anything to say about how language is processed in the mind, but it might still be seen as an advantage if it at least makes available the kinds of objects the mind seems to work with – e.g. constructions.

the presence of multiple specific words, potentially in quite distant parts of the phrase, and here the approach will once again have to rely on multiple interacting lexical entries and phrase-structure rules which conspire to give the correct constructional effects. Even if this gives the right results, one might, again, object that it does so for the wrong reasons, since it fails to account for the unitary nature of constructions as grammatical objects. By contrast, in the description-based approach to merging TAG and LFG, although constructions are still not added as new objects in the ontology of the theory, they nevertheless have a kind of first-class status, since they can either be entire (complex) lexical entries or be a part of a lexical entry in the form of a tree template which can be called by all the different words which can fill the empty slots in the construction. See Findlay (2023) for a broader discussion of the connection between LFG and Construction Grammar, and for arguments that vanilla LFG is inadequate to give a satisfactory analysis of certain multiword (substantive) constructions.

Alongside these advantages of combining LFG and TAG, there remain some unexplored implications which are ripe for future work. Firstly, one of the parade examples of LFG's utility is in describing languages with highly flexible word orders, such as Warlpiri (see e.g. Bresnan et al. 2016: ch. 1). With TAG's focus on configurational properties, we need to ensure that incorporating a TAG into LFG does not undo its ability to describe these non-configurational languages. Given the flatter tree structures generally assumed for such languages (Simpson 1991, Austin & Bresnan 1996), a first pass solution in the present framework would be to simply make use of looser tree descriptions, which, for example, lack linear precedence relations between a verb and its arguments, so that the entry for a verb does not describe a unique minimal tree, but rather several minimal trees which represent the different orderings of arguments. Of course, these different orderings are not just random, and actually correspond to different information structures, so simply allowing free choice between them is inadequate. Instead, we should once again make use of disjunction, this time between the descriptions corresponding to the different orderings of verb and arguments, where each of the different word orders is also accompanied by the correct information-structural annotations.

Such languages also often permit discontinuous constituents, and these will require their own solution. For example, some adjuncts might be represented not as auxiliary trees that induce a more articulated structure, but rather as simpler trees whose root merely unifies with another node, such as the clausal root *S*, adding the adjunct as a sister to the existing daughters. Obviously this sketch needs to be developed into a fully fleshed-out proposal before we can be confident that no analytical clout has been lost.

Another open question arises from the fact that using a TAG as the basis of the c-structure component means that we can employ adjunction to account for long-distance dependencies. This then removes a foundational motivation for functional uncertainty (Kaplan & Zaenen 1989), one of the major sources of formal complexity in LFG. Unfortunately, this does not mean we can simply remove functional uncertainty from the formalism, since it has been employed by researchers in various other domains beyond filler-gap dependencies – most notably in LFG’s binding theory (e.g. Dalrymple 1993, Dalrymple et al. 2018). Determining whether these analyses can be reformulated so that functional uncertainty could be done away with altogether remains a task for future work, perhaps drawing on existing TAG analyses of binding (e.g. Ryant & Scheffler 2006, Champollion 2008, Storoshenko et al. 2008, Storoshenko & Han 2013).

Lastly, including a description of a tree which incorporates the full extended projection of a predicate in its lexical entry means that we can take a rather different view of argument structure. A predicate’s arguments and the possibilities for their realisation can be encoded directly in its lexical entry, rather than relying on a separate level of representation like a-structure (on which see Findlay et al. 2023 [this volume]). And alternative argument realisations, e.g. diathesis alternations, can be expressed through disjunctive templates, as discussed above in parallel with XMG, rather than through a separate mechanism of mapping between a-structure and f-structure. Work on developing templatic approaches to argument structure include Asudeh & Giorgolo (2012), Findlay (2016, 2020), and Przepiórkowski (2017), but these do not take c-structure into account: with the new TAG perspective, the phrase-structural effects of argument structure/mapping phenomena can also be directly expressed.

8 Conclusion

Tree-Adjoining Grammar offers a rather different perspective on some grammatical phenomena from that of CFG-based formalisms.³⁴ For instance, it allows us to describe constraints on filler-gap relationships via the structure of the elementary trees in the grammar rather than via an independent principle like Subjacency. It also provides a natural account of the fact that many “lexical” items in fact incorporate several distinct word forms (e.g. phrasal verbs, compounds, idioms), and of constructional meaning, by virtue of its expanded concept of locality. And, computationally speaking, it possesses just the right degree

³⁴TAG has also played an important role outside of theoretical linguistics – specifically in both computational linguistics (see e.g. Kallmeyer et al. 2008, Kasai et al. 2017, Koller 2017) and psycholinguistics (see e.g. Ferreira 2000, Ferreira et al. 2004).

of context-sensitivity to account for natural languages while remaining parsable in polynomial time. Nonetheless, its representation of dependency structures is imperfect, and its focus on the primacy of phrase structure leaves it somewhat impoverished when compared to the richly expressive parallel projection architecture of LFG, which facilitates a much fuller view of the grammar as a whole. Combining the two approaches might therefore offer a tempting opportunity to acquire the best of both worlds. In Section 7 we saw how this could be achieved, and the possibilities this affords for LFG, both in terms of descriptive power and in terms of potentially further-reaching formal or architectural changes.

Acknowledgements

Much of this work stems from my D.Phil. thesis (Findlay 2019), which was funded by a UK Arts and Humanities Research Council studentship (grant reference AH/L503885/1), and for which I am indebted to my fantastic supervisors, Ash Asudeh and Mary Dalrymple. The writing of this chapter was completed while I was employed under a Norwegian Research Council grant (number 300495, “Universal Natural Language Understanding”), which I gratefully acknowledge.

References

- Abeillé, Anne. 1988. Parsing French with Tree Adjoining Grammar: Some linguistic accounts. In *COLING '88: Proceedings of the 12th Conference on Computational Linguistics*, 7–12. Budapest. DOI: 10.3115/991635.991637.
- Abeillé, Anne. 1995. The flexibility of French idioms: A representation with Lexicalized Tree Adjoining Grammar. In Martin Everaert, Erik-Jan van der Linden, André Schenk & Rob Schreuder (eds.), *Idioms: Structural and psychological perspectives*, 15–42. Hove: Lawrence Erlbaum Associates.
- Abeillé, Anne & Owen Rambow. 2000. Tree Adjoining Grammar: An overview. In Anne Abeillé & Owen Rambow (eds.), *Tree Adjoining Grammars: Formalisms, linguistic analysis and processing*, 1–68. Stanford: CSLI Publications.
- Abeillé, Anne, Yves Schabes & Aravind K. Joshi. 1990. Using Lexicalized TAGs for machine translation. In *COLING '90: Proceedings of the 13th Conference on Computational Linguistics*, vol. 3, 1–6. DOI: 10.3115/991146.991147.
- Allen, Kachina, Francisco Pereira, Matthew Botvinick & Adele E. Goldberg. 2012. Distinguishing grammatical constructions with fMRI pattern analysis. *Brain and Language* 123(3). 174–182. DOI: 10.1016/j.bandl.2012.08.005.

- Asudeh, Ash. 2012. *The logic of pronominal resumption* (Oxford Studies in Theoretical Linguistics). Oxford: Oxford University Press. DOI: 10.1093/acprof:oso/9780199206421.001.0001.
- Asudeh, Ash. 2023. Glue semantics. In Mary Dalrymple (ed.), *Handbook of Lexical Functional Grammar*, 651–697. Berlin: Language Science Press. DOI: 10.5281/zenodo.10185964.
- Asudeh, Ash, Mary Dalrymple & Ida Toivonen. 2013. Constructions with Lexical Integrity. *Journal of Language Modelling* 1(1). 1–54. DOI: 10.15398/jlm.v1i1.56.
- Asudeh, Ash & Gianluca Giorgolo. 2012. Flexible composition for optional and derived arguments. In Miriam Butt & Tracy Holloway King (eds.), *Proceedings of the LFG '12 conference*, 64–84. Stanford: CSLI Publications.
- Austin, Peter K. & Joan Bresnan. 1996. Non-configurationality in Australian aboriginal languages. *Natural Language & Linguistic Theory* 14(2). 215–268. DOI: 10.1007/bf00133684.
- Bach, Emmon W. 1976. An extension of classical transformational grammar. In Jerrold M. Sadock, David J. Dwyer, Seok C. Song & Emmon Bach (eds.), *Problems in linguistic metatheory: Proceedings of the 1976 conference*, 183–224. East Lansing, MI: Michigan State University.
- Bach, Emmon W. 1981. Discontinuous constituents in generalized categorial grammars. In V. A. Burke & James Pustejovsky (eds.), *Proceedings of the 11th annual meeting of the North Eastern Linguistics Society*, 1–12. Amherst: Graduate Linguistics Student Association of the University of Massachusetts, Amherst.
- Baker, Mark C. 2008. The macroparameter in a microparametric world. In Theresa Biberauer (ed.), *The limits of syntactic variation*, 351–374. Philadelphia: John Benjamins. DOI: 10.1075/la.132.16bak.
- Bangalore, Srinivas. 1997. *Complexity of lexical descriptions and its relevance for partial parsing*. University of Pennsylvania. (Doctoral dissertation).
- Bangalore, Srinivas & Aravind K. Joshi. 2010. Introduction. In Srinivas Bangalore & Aravind K. Joshi (eds.), *Supertagging: Using complex lexical descriptions in natural language processing*, 1–31. Cambridge, MA: The MIT Press. DOI: 10.7551/mitpress/8370.003.0004.
- Becker, Tilman. 1994. *HyTAG: A new type of Tree Adjoining Grammars for hybrid syntactic representations of free word order languages*. Saarbrücken: Universität des Saarlandes. (Doctoral dissertation). <http://www.dfki.de/~becker/becker.diss.ps.gz>.
- Belyaev, Oleg. 2023a. Core concepts of LFG. In Mary Dalrymple (ed.), *Handbook of Lexical Functional Grammar*, 23–96. Berlin: Language Science Press. DOI: 10.5281/zenodo.10185936.

- Belyaev, Oleg. 2023b. Introduction to LFG. In Mary Dalrymple (ed.), *Handbook of Lexical Functional Grammar*, 3–22. Berlin: Language Science Press. DOI: 10.5281/zenodo.10185934.
- Bencini, Giulia M. L. & Adele E. Goldberg. 2000. The contribution of argument structure constructions to sentence meaning. *Journal of Memory and Language* 43(4). 640–651. DOI: 10.1006/jmla.2000.2757.
- Berwick, Robert C. 1982. Computational complexity and Lexical-Functional Grammar. *American Journal of Computational Linguistics* 8(3–4). 97–109. DOI: 10.3115/981923.981926.
- Boas, Hans C. & Ivan A. Sag (eds.). 2012. *Sign-Based Construction Grammar*. Stanford: CSLI Publications.
- Borer, Hagit. 1984. *Parametric syntax: Case studies in Semitic and Romance languages* (Studies in Generative Grammar 13). Dordrecht: Foris Publications. DOI: 10.1515/9783110808506.
- Bresnan, Joan, Ash Asudeh, Ida Toivonen & Stephen Wechsler. 2016. *Lexical-Functional Syntax*. 2nd edn. (Blackwell Textbooks in Linguistics 16). Malden, MA: Wiley-Blackwell.
- Bresnan, Joan, Ronald M. Kaplan, Stanley Peters & Annie Zaenen. 1982. Cross-serial dependencies in Dutch. *Linguistic Inquiry* 13(4). 613–635. <https://www.jstor.org/stable/4178298>. Reprinted in Savitch, Bach, Marsh & Safran-Naveh (1987: 286–319).
- Burheim, Tore. 1996. Aspects of merging Lexical-Functional Grammar with Lexicalized Tree-Adjoining Grammar. Unpublished manuscript, University of Bergen.
- Candito, Marie-Hélène. 1996. A principle-based hierarchical representation of LTAGs. In *COLING '96: Proceedings of the 16th Conference on Computational Linguistics*, 194–199. DOI: 10.3115/992628.992664.
- Candito, Marie-Hélène. 1999. *Représentation modulaire et paramétrable de grammaires électroniques lexicalisées : application au français et à l'italien*. Université Paris 7. (Doctoral dissertation).
- Champollion, Lucas. 2008. Binding theory in LTAG. In *Proceedings of the Ninth International Workshop on Tree Adjoining Grammar and Related Frameworks (TAG+9)*, 1–8. Tübingen: Association for Computational Linguistics. <https://aclanthology.org/W08-2301>.
- Chomsky, Noam. 1956. Three models for the description of language. *IRE Transactions on Information Theory* 2(3). 113–124. DOI: 10.1109/tit.1956.1056813.
- Chomsky, Noam. 1970. Remarks on nominalization. In Roderick A. Jacobs & Peter S. Rosenbaum (eds.), *Readings in English transformational grammar*, 184–221. Waltham, MA: Ginn.

- Chomsky, Noam. 1995. *The Minimalist Program*. Cambridge, MA: The MIT Press. DOI: 10.7551/mitpress/9780262527347.001.0001.
- Clément, Lionel & Alexandra Kinyon. 2003a. Generating LFGs with a MetaGrammar. In Miriam Butt & Tracy Holloway King (eds.), *Proceedings of the LFG '03 conference*, 105–125. Stanford: CSLI Publications.
- Clément, Lionel & Alexandra Kinyon. 2003b. Generating parallel multilingual LFG-TAG grammars from a metagrammar. In *Proceedings of the 41st annual meeting of the Association for Computational Linguistics*, 184–191. DOI: 10.3115/1075096.1075120.
- Crabbé, Benoît, Denys Duchier, Claire Gardent, Joseph Le Roux & Yannick Parmentier. 2013. XMG: eXtensible MetaGrammar. *Computational Linguistics* 39(3). 591–629. DOI: 10.1162/COLI_a_00144.
- Crouch, Richard, Mary Dalrymple, Ronald M. Kaplan, Tracy Holloway King, John T. III Maxwell & Paula S. Newman. 2011. *XLE Documentation*. Xerox Palo Alto Research Center. Palo Alto, CA. https://ling.sprachwiss.uni-konstanz.de/pages/xle/doc/xle_toc.html.
- Dalrymple, Mary. 1993. *The syntax of anaphoric binding*. Stanford: CSLI Publications.
- Dalrymple, Mary. 2001. *Lexical Functional Grammar* (Syntax and Semantics 34). New York: Academic Press. DOI: 10.1163/9781849500104.
- Dalrymple, Mary, Dag T. Haug & John J. Lowe. 2018. Integrating LFG's binding theory with PCDRT. *Journal of Language Modelling* 6(1). 87–129. DOI: 10.15398/jlm.v6i1.204.
- Dalrymple, Mary, Ronald M. Kaplan & Tracy Holloway King. 2004. Linguistic generalizations over descriptions. In Miriam Butt & Tracy Holloway King (eds.), *Proceedings of the LFG '04 conference*, 199–208. Stanford: CSLI Publications.
- Dalrymple, Mary, Ronald M. Kaplan, John T. III Maxwell & Annie Zaenen (eds.). 1995. *Formal issues in Lexical-Functional Grammar*. Stanford: CSLI Publications.
- Dalrymple, Mary, John J. Lowe & Louise Mycock. 2019. *The Oxford reference guide to Lexical Functional Grammar*. Oxford: Oxford University Press. DOI: 10.1093/oso/9780198733300.001.0001.
- Di Sciullo, Anna Maria & Edwin Williams. 1987. *On the definition of word* (Linguistic Inquiry Monographs 14). Cambridge, MA: The MIT Press.
- Eisner, Jason & Giorgio Satta. 2000. A faster parsing algorithm for Lexicalized Tree-Adjoining Grammars. In *Proceedings of the Fifth International Workshop on Tree Adjoining Grammar and Related Frameworks (TAG+5)*, 79–84. Université Paris 7. <http://aclweb.org/anthology/W00-2011>.

- Ferreira, Fernanda. 2000. Syntax in language production: An approach using Tree-Adjoining Grammars. In Linda Wheeldon (ed.), *Aspects of language production*, 291–330. Hove: Psychology Press.
- Ferreira, Fernanda, Ellen F. Lau & Karl G. D. Bailey. 2004. Disfluencies, language comprehension, and Tree Adjoining Grammars. *Cognitive Science* 28(5). 721–749. DOI: 10.1207/s15516709cog2805_5.
- Fillmore, Charles J., Paul Kay & Mary Catherine O'Connor. 1988. Regularity and idiomaticity in grammatical constructions: The case of *let alone*. *Language* 64. 501–538.
- Findlay, Jamie Y. 2016. Mapping theory without argument structure. *Journal of Language Modelling* 4(2). 293–338. DOI: 10.15398/jlm.v4i2.171.
- Findlay, Jamie Y. 2017a. Multiword expressions and lexicalism. In Miriam Butt & Tracy Holloway King (eds.), *Proceedings of the LFG '17 conference*, 200–229. Stanford: CSLI Publications.
- Findlay, Jamie Y. 2017b. Multiword expressions and lexicalism: The view from LFG. In *Proceedings of the 13th Workshop on Multiword Expressions (MWE 2017)*, 73–79. Valencia, Spain: Association for Computational Linguistics. <http://aclweb.org/anthology/W17-1709>.
- Findlay, Jamie Y. 2019. *Multiword expressions and the lexicon*. Oxford: University of Oxford. (D.Phil. Thesis).
- Findlay, Jamie Y. 2020. Mapping Theory and the anatomy of a verbal lexical entry. In Miriam Butt & Ida Toivonen (eds.), *Proceedings of the LFG '20 conference*, 127–147. Stanford: CSLI Publications.
- Findlay, Jamie Y. 2023. Lexical Functional Grammar as a Construction Grammar. To appear in *Journal of Language Modelling* 11(2).
- Findlay, Jamie Y., Roxanne Taylor & Anna Kibort. 2023. Argument structure and mapping theory. In Mary Dalrymple (ed.), *Handbook of Lexical Functional Grammar*, 699–778. Berlin: Language Science Press. DOI: 10.5281/zenodo.10185966.
- Frank, Anette & Josef van Genabith. 2001. GlueTag: Linear logic based semantics for LTAG – and what it teaches us about LFG and LTAG. In Miriam Butt & Tracy Holloway King (eds.), *Proceedings of the LFG '01 conference*, 104–126. Stanford: CSLI Publications.
- Gardent, Claire & Laura Kallmeyer. 2003. Semantic construction in feature-based TAG. In *Proceedings of the 10th conference of the European chapter of the ACL (EACL '03)*, vol. 1, 123–130. DOI: 10.3115/1067807.1067825.
- Gazdar, Gerald, Ewan Klein, Geoffrey K. Pullum & Ivan A. Sag. 1985. *Generalized Phrase Structure Grammar*. Cambridge, MA: Harvard University Press.

- Goldberg, Adele E. 1995. *Constructions: A Construction Grammar approach to argument structure*. Chicago: University of Chicago Press.
- Goldberg, Adele E. 2006. *Constructions at work: The nature of generalization in language*. Oxford: Oxford University Press.
- Goldwater, Micah B. & Arthur B. Markman. 2009. Constructional sources of implicit agents in sentence comprehension. *Cognitive Linguistics* 20(4). 675–702. DOI: 10.1515/COGL.2009.029.
- Gorn, Saul. 1967. Explicit definitions and linguistic dominoes. In John Francis Hart & Satoru Takasu (eds.), *Systems and computer science*, 77–115. University of Toronto Press. DOI: 10.3138/9781487592769-008.
- Greibach, Sheila A. 1965. A new normal-form theorem for context-free phrase structure grammars. *Journal of the Association for Computing Machinery* 12(1). 42–52. DOI: 10.1145/321250.321254.
- Grimshaw, Jane. 2000. Locality and extended projection. In Peter Coopmans, Martin Everaert & Jane Grimshaw (eds.), *Lexical specification and insertion*, 115–133. Amsterdam/Philadelphia: John Benjamins. DOI: 10.1075/cilt.197.07gri.
- Grimshaw, Jane. 2005. Extended projection. In *Words and structure*, 1–74. Stanford: CSLI Publications. Published version of 1991 manuscript by the same name.
- Harris, Zellig. 1946. From morpheme to utterance. *Language* 22(3). 161–183. DOI: 10.2307/410205. Reprinted in Joos (ed) (1957), *Readings In Structural Linguistics*, University of Chicago Press, 142–153.
- Haug, Dag. 2023a. Agreement. In Mary Dalrymple (ed.), *Handbook of Lexical Functional Grammar*, 193–218. Berlin: Language Science Press. DOI: 10.5281/zenodo.10185942.
- Haug, Dag. 2023b. LFG and Dependency Grammar. In Mary Dalrymple (ed.), *Handbook of Lexical Functional Grammar*, 1829–1859. Berlin: Language Science Press. DOI: 10.5281/zenodo.10186040.
- Hoffmann, Thomas & Graeme Trousdale (eds.). 2013. *The Oxford handbook of Construction Grammar*. Oxford: Oxford University Press. DOI: 10.1093/oxfordhb/9780195396683.001.0001.
- Johnson, Matthew A. & Adele E. Goldberg. 2012. Evidence for automatic accessing of constructional meaning: Jabberwocky sentences prime associated verbs. *Language and Cognitive Processes* 28(10). 1–14. DOI: 10.1080/01690965.2012.717632.
- Joshi, Aravind K. 1985. Tree adjoining grammars: How much context-sensitivity is required to provide reasonable structural descriptions? In David R. Dowty, Lauri Karttunen & Arnold M. Zwicky (eds.), *Natural language parsing: Psycho-*

- logical, computational, and theoretical perspectives*, 206–250. Cambridge, UK: Cambridge University Press.
- Joshi, Aravind K. 1987. An introduction to Tree Adjoining Grammars. In Alexis Manaster-Ramer (ed.), *Mathematics of language: Proceedings of a conference held at the University of Michigan, Ann Arbor, October 1984*, 87–114. Amsterdam: John Benjamins. DOI: 10.1075/z.35.07jos.
- Joshi, Aravind K. 2005. Tree-adjoining grammars. In Ruslan Mitkov (ed.), *The Oxford handbook of computational linguistics*, 1st edn., 483–498. Oxford: Oxford University Press. DOI: 10.1093/oxfordhb/9780199276349.013.0026.
- Joshi, Aravind K., Leon S. Levy & Masako Takahashi. 1975. Tree adjunct grammars. *Journal of Computer and System Sciences* 10(1). 136–163. DOI: 10.1016/S0022-0000(75)80019-5.
- Joshi, Aravind K. & Yves Schabes. 1997. Tree-Adjoining Grammars. In Grzegorz Rozenberg & Arto Salomaa (eds.), *Handbook of formal languages, vol. 3: Beyond words*, 69–123. Berlin: Springer.
- Joshi, Aravind K. & K. Vijay-Shanker. 2001. Compositional semantics with Lexicalized Tree-Adjoining Grammar (LTAG): How much underspecification is necessary? In Harry Bunt, Reinhard Muskens & Elias Thijsse (eds.), *Computing meaning*, vol. 2 (Studies in Linguistics and Philosophy 77), 147–163. Dordrecht: Springer. DOI: 10.1007/978-94-010-0572-2_9.
- Joshi, Aravind K., K. Vijay-Shanker & David Weir. 1991. The convergence of mildly context-sensitive formalisms. In Peter Sells, Stuart M. Shieber & Thomas Wasow (eds.), *Foundational issues in natural language processing*. Cambridge, MA: The MIT Press.
- Joshi, Aravind K. & Takashi Yokomori. 1983. *Parsing of tree adjoining grammars*. Tech. rep. Department of Computer & Information Science, University of Pennsylvania.
- Kallmeyer, Laura. 2001. Local tree description grammars. *Grammars* 4. 85–137.
- Kallmeyer, Laura. 2010. *Parsing beyond context-free grammars*. Berlin: Springer. DOI: 10.1007/978-3-642-14846-0.
- Kallmeyer, Laura & Aravind K. Joshi. 2003. Factoring predicate argument and scope semantics: underspecified semantics with LTAG. *Research on Language and Computation* 1(1–2). 3–58. DOI: 10.1023/A:1024564228892.
- Kallmeyer, Laura & Marco Kuhlmann. 2012. A formal model for plausible dependencies in lexicalized tree adjoining grammar. In *Proceedings of the 11th International Workshop on Tree Adjoining Grammars and Related Formalisms (TAG+11)*, 108–116. Paris. <https://aclanthology.org/W12-4613>.

- Kallmeyer, Laura, Timm Lichte, Wolfgang Maier, Yannick Parmentier, Johannes Dellert & Kilian Evang. 2008. TuLiPA: Towards a multi-formalism parsing environment for grammar engineering. In Stephen Clark & Tracy Holloway King (eds.), *Proceedings of the workshop on Grammar Engineering Across Frameworks (GEAF08)*, 1–8. Association for Computational Linguistics. <https://www.aclweb.org/anthology/W08-1701.pdf>.
- Kallmeyer, Laura & Rainer Osswald. 2013. Syntax-driven semantic frame composition in Lexicalized Tree Adjoining Grammars. *Journal of Language Modelling* 1(2). 267–330. DOI: 10.15398/jlm.v1i2.61.
- Kallmeyer, Laura & Maribel Romero. 2004. LTAG semantics with semantic unification. In *TAG+7: Seventh International Workshop on Tree Adjoining Grammar and Related Formalisms*, 155–162. Vancouver. <https://aclanthology.org/W04-3321>.
- Kallmeyer, Laura & Maribel Romero. 2008. Scope and situation binding in LTAG using semantic unification. *Research on Language and Computation* 6(1). 3–52. DOI: 10.1007/s11168-008-9046-6.
- Kameyama, Megumi. 1986. Characterising Lexical Functional Grammar (LFG) in terms of Tree Adjoining Grammar (TAG). Unpublished manuscript, Department of Computer and Information Science, University of Pennsylvania.
- Kanazawa, Makoto & Sylvain Salvati. 2012. MIX is not a tree-adjoining language. In *Proceedings of the 50th annual meeting of the Association for Computational Linguistics*, 666–674. Association for Computational Linguistics. <https://www.aclweb.org/anthology/P12-1070>.
- Kaplan, Ronald M. 1989. The formal architecture of Lexical-Functional Grammar. *Journal of Information Science and Engineering* 5. 305–322. Revised version published as Kaplan (1995).
- Kaplan, Ronald M. 1995. The formal architecture of Lexical-Functional Grammar. In Mary Dalrymple, Ronald M. Kaplan, John T. III Maxwell & Annie Zaenen (eds.), *Formal issues in Lexical-Functional Grammar*, 7–27. Stanford: CSLI Publications. Earlier version published as Kaplan (1989).
- Kaplan, Ronald M. & Joan Bresnan. 1982. Lexical-Functional Grammar: A formal system for grammatical representation. In Joan Bresnan (ed.), *The mental representation of grammatical relations*, 173–281. Cambridge, MA: The MIT Press. Reprinted in Dalrymple, Kaplan, Maxwell & Zaenen (1995: 29–130).
- Kaplan, Ronald M., John T. III Maxwell & Annie Zaenen. 1987. Functional uncertainty. In *CSLI Publications Monthly Newsletter*. Stanford: Stanford University.
- Kaplan, Ronald M. & Paula S. Newman. 1997. Lexical resource reconciliation in the Xerox Linguistic Environment. In *Proceedings of the ACL workshop on Com-*

- putational Environments for Grammar Development and Engineering*. Association for Computational Linguistics.
- Kaplan, Ronald M. & Jürgen Wedekind. 2023. Formal and computational properties of LFG. In Mary Dalrymple (ed.), *Handbook of Lexical Functional Grammar*, 1035–1082. Berlin: Language Science Press. DOI: 10.5281/zenodo.10185982.
- Kaplan, Ronald M. & Annie Zaenen. 1989. Long-distance dependencies, constituent structure, and functional uncertainty. In Mark Baltin & Anthony Kroch (eds.), *Alternative conceptions of phrase structure*, 17–42. Chicago: University of Chicago Press. Reprinted in Dalrymple, Kaplan, Maxwell & Zaenen (1995: 137–165).
- Kasai, Jungo, Bob Frank, Tom McCoy, Owen Rambow & Alexis Nasr. 2017. TAG parsing with neural networks and vector representations of supertags. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, 1712–1722. Copenhagen: Association for Computational Linguistics. DOI: 10.18653/v1/D17-1180. <http://aclweb.org/anthology/D17-1180>.
- Kaschak, Michael P. & Arthur M. Glenberg. 2000. Constructing meaning: The role of affordances and grammatical constructions in sentence comprehension. *Journal of Memory and Language* 43(3). 508–529. DOI: 10.1006/jmla.2000.2705.
- Kay, Paul & Charles J. Fillmore. 1999. Grammatical constructions and linguistic generalizations: The *What's X doing Y?* construction. *Language* 75. 1–33.
- Kinyon, Alexandra. 2000. Hypertags. In *COLING 2000: The 18th International Conference on Computational Linguistics*, vol. 1. DOI: 10.3115/990820.990885.
- Koller, Alexander. 2017. A feature structure algebra for FTAG. In *Proceedings of the 13th International Workshop on Tree Adjoining Grammars and Related Formalisms (TAG+13)*. <https://www.aclweb.org/anthology/W17-6201>.
- Kroch, Anthony. 1987. Unbounded dependencies and subadjacency in a Tree Adjoining Grammar. In Alexis Manaster-Ramer (ed.), *Mathematics of language: Proceedings of a conference held at the University of Michigan, Ann Arbor, October 1984*, 143–172. Amsterdam: John Benjamins. DOI: 10.1075/z.35.09kro.
- Kroch, Anthony & Aravind K. Joshi. 1985. *The linguistic relevance of Tree Adjoining Grammar*. Tech. rep. MC-CS-85-16. Department of Computer & Information Sciences, University of Pennsylvania.
- Kuhlmann, Marco. 2010. *Dependency structures and lexicalized grammars: An algebraic approach*. Berlin: Springer. DOI: 10.1007/978-3-642-14568-1.
- Lev, Iddo. 2007. *Packed computation of exact meaning representations*. Stanford: Stanford University. (Doctoral dissertation).
- Lichte, Timm & Laura Kallmeyer. 2017. Tree-Adjoining Grammar: A tree-based constructionist grammar framework for natural language understanding. In *Proceedings of the AAAI 2017 Spring Symposium on Computational Construction*

- Grammar and Natural Language Understanding* (Technical Report SS-17-02), 205–212. Association for the Advancement of Artificial Intelligence. <https://www.aaai.org/ocs/index.php/SSS/SSS17/paper/viewFile/15330/14536>.
- Maxwell, John T. III & Ronald M. Kaplan. 1989. An overview of disjunctive constraint satisfaction. In *Proceedings of the 4th International Workshop on Parsing Technologies (IWPT 1995)*, 18–27. Also published in Tomita (1991) as ‘A Method for Disjunctive Constraint Satisfaction’, and reprinted in Dalrymple, Kaplan, Maxwell & Zaenen (1995: 381–402).
- Maxwell, John T. III & Ronald M. Kaplan. 1993. The interface between phrasal and functional constraints. *Computational Linguistics* 19. 571–590.
- Maxwell, John T. III & Ronald M. Kaplan. 1996. Unification-based parsers that automatically take advantage of context freeness. In Miriam Butt & Tracy Holloway King (eds.), *Proceedings of the LFG '96 conference*, 1–31. Stanford: CSLI Publications.
- McCawley, James D. 1968. Concerning the base component of a transformational grammar. *Foundations of Language* 4(3). 243–269.
- Mel'čuk, Igor A. 1988. *Dependency syntax: Theory and practice*. Albany, NY: State University of New York Press.
- Müller, Stefan & Stephen Wechsler. 2014. Lexical approaches to argument structure. *Theoretical Linguistics* 40(1–2). 1–76. DOI: 10.1515/tl-2014-0001.
- Nakanishi, Ryuichi, Hiroyuki Seki & Tadao Kasami. 1992. On the generative capacity of Lexical-Functional Grammars. *IEICE Transactions on Information and Systems* E75-D(4). 509–516. https://search.ieice.org/bin/summary.php?id=e75-d_4_509.
- Nederhof, Mark-Jan. 2016. A short proof that O_2 is an MCFL. In *Proceedings of the 54th annual meeting of the Association for Computational Linguistics*, 1117–1126. Association for Computational Linguistics. <https://www.aclweb.org/anthology/P16-1106>.
- Nesson, Rebecca & Stuart M. Shieber. 2006. Simpler TAG semantics through synchronization. In Shuly Wintner (ed.), *Proceedings of the 11th Conference on Formal Grammar (FG 2006)*, 129–142. Stanford: CSLI Publications. <https://csli-publications.stanford.edu/FG/2006/nesson.pdf>.
- Nesson, Rebecca & Stuart M. Shieber. 2007. Extraction phenomena in synchronous TAG syntax and semantics. In Dekai Wu & David Chiang (eds.), *Proceedings of the Workshop on Syntax and Structure in Statistical Translation*. Rochester, New York, 26 April 2007.
- Nesson, Rebecca & Stuart M. Shieber. 2008. Synchronous vector-TAG for natural language syntax and semantics. In *Proceedings of the Ninth International Work-*

- shop on Tree Adjoining Grammars and Related Formalisms (TAG+ 9)*. Tübingen. <http://tagplus9.cs.sfu.ca/papers/NessonShieber.pdf>.
- Nivre, Joakim, Marie-Catherine de Marneffe, Filip Ginter, Yoav Goldberg, Jan Hajič, Christopher D. Manning, Ryan McDonald, Slav Petrov, Sampo Pyysalo, Natalia Silveira, Reut Tsarfaty & Daniel Zeman. 2016. Universal Dependencies v1: A multilingual treebank collection. In *Proceedings of the 10th International Conference on Language Resources and Evaluation (LREC'16)*, 1659–1666. Portorož: European Language Resources Association (ELRA). <https://aclanthology.org/L16-1262>.
- Partee, Barbara H., Alice ter Meulen & Robert E. Wall. 1990. *Mathematical methods in linguistics*. Dordrecht: Kluwer Academic Publishers. DOI: 10.1007/978-94-009-2213-6.
- Pawley, Andrew & Frances Hodgetts Syder. 1983. Two puzzles for linguistic theory: Nativelike selection and nativelike fluency. In Jack C. Richards & Richard W. Schmidt (eds.), *Language and communication*, 191–226. London: Longman.
- Peters, Stanley & R. W. Ritchie. 1973. On the generative power of transformational grammars. *Information Sciences* 6. 49–83. DOI: 10.1016/0020-0255(73)90027-3.
- Pollard, Carl. 1997. The nature of constraint-based grammar. *Linguistic Research* 15. 1–18. <http://isli.khu.ac.kr/journal/content/data/15/1.pdf>.
- Pollard, Carl & Ivan A. Sag. 1994. *Head-Driven Phrase Structure Grammar*. Chicago: University of Chicago Press & CSLI Publications.
- Przepiórkowski, Adam. 2017. A full-fledged hierarchical lexicon in LFG: The FrameNet approach. In Victoria Rosén & Koenraad De Smedt (eds.), *The very model of a modern linguist – In honor of Helge Dyvik*, 202–219. Bergen: Bergen Language & Linguistics Studies (BeLLS). DOI: 10.15845/bells.v8i1.1336.
- Pullum, Geoffrey K. & Gerald Gazdar. 1982. Natural languages and context-free languages. *Linguistics and Philosophy* 4. 471–504. DOI: 10.1007/bf00360802.
- Pulvermüller, Friedmann. 2010. Brain embodiment of syntax and grammar: Discrete combinatorial mechanisms spelt out in neuronal circuits. *Brain and Language* 112(3). 167–179. DOI: 10.1016/j.bandl.2009.08.002.
- Rambow, Owen & Aravind K. Joshi. 1997. A formal look at dependency grammars and phrase-structure grammars, with special consideration of word-order phenomena. In Leo Wanner (ed.), *Recent trends in Meaning-Text Theory* (Studies in Language Companion Series 39), 167–190. Amsterdam: John Benjamins.
- Rambow, Owen, K. Vijay-Shanker & David Weir. 1995. D-Tree Grammars. In *Proceedings of the 33rd annual meeting of the Association for Computational Linguistics (ACL '95)*, 151–158. DOI: 10.3115/981658.981679.

- Rambow, Owen, K. Vijay-Shanker & David Weir. 2001. D-tree substitution grammars. *Computational Linguistics* 27(1). 87–121. DOI: 10.1162/089120101300346813.
- Rogers, James. 1998. *A descriptive approach to language-theoretic complexity*. Stanford: CSLI Publications.
- Rogers, James & K. Vijay-Shanker. 1994. Obtaining trees from their descriptions: An application to Tree-adjoining Grammars. *Computational Intelligence* 10(4). 401–421. DOI: 10.1111/j.1467-8640.1994.tb00005.x.
- Ryant, Neville & Tatjana Scheffler. 2006. Binding of anaphors in LTAG. In *Proceedings of the Eighth International Workshop on Tree Adjoining Grammar and Related Formalisms*, 65–72. Sydney. DOI: 10.3115/1654690.1654699.
- Salvati, Sylvain. 2015. MIX is a 2-MCFL and the word problem in \mathbb{Z}^2 is captured by the IO and the OI hierarchies. *Journal of Computer and System Sciences* 18(7). 1252–1277. DOI: 10.1016/j.jcss.2015.03.004.
- Savitch, Walter J., Emmon W. Bach, William Marsh & Gila Safran-Naveh (eds.). 1987. *The formal complexity of natural language* (Studies in Linguistics and Philosophy). Dordrecht: Springer. DOI: 10.1007/978-94-009-3401-6.
- Schabes, Yves. 1990. *Mathematical and computational aspects of lexicalized grammars*. University of Pennsylvania. (Doctoral dissertation).
- Schabes, Yves, Anne Abeillé & Aravind K. Joshi. 1988. Parsing strategies with ‘lexicalized’ grammars: application to Tree Adjoining Grammars. In *COLING ’88: Proceedings of the 12th Conference on Computational Linguistics*, 578–583. Stroudsburg, PA: Association for Computational Linguistics. DOI: 10.3115/991719.991757.
- Seki, Hiroyuki, Takahashi Matsumura, Mamoru Fujii & Tadao Kasami. 1991. On multiple context-free grammars. *Theoretical Computer Science* 88(2). 191–229. DOI: 10.1016/0304-3975(91)90374-b.
- Seki, Hiroyuki, Ryuichi Nakanishi, Yuichi Kaji, Sachiko Ando & Tadao Kasami. 1993. Parallel multiple context-free grammars, finite-state translation systems, and polynomial-time recognizable subclasses of Lexical-Functional Grammars. In *Proceedings of the 31st annual meeting of the Association for Computational Linguistics*, 130–139. Columbus, OH: Association for Computational Linguistics. DOI: 10.3115/981574.981592.
- Shieber, Stuart M. 1985. Evidence against the context-freeness of natural language. *Linguistics and Philosophy* 8(3). 333–343. DOI: 10.1007/978-94-009-3401-6_12.
- Shieber, Stuart M. & Yves Schabes. 1990. Synchronous Tree Adjoining Grammars. In *Proceedings of the 13th International Conference on Computational Linguistics (COLING ’90)*, 253–258. Helsinki.

- Simpson, Jane. 1991. *Warlpiri morpho-syntax: A lexicalist approach*. Dordrecht: Kluwer Academic Publishers.
- Srinivas, Bangalore, Dania Egedi, Christy Doran & Tilman Becker. 1994. Lexicalization and grammar development. In *Proceedings of KONVENS 94, Vienna, Austria, September 1994*, 310–319. DOI: 10.48550/ARXIV.CMP-LG/9410015.
- Steedman, Mark. 1987. Combinatory grammars and parasitic gaps. *Natural Language & Linguistic Theory* 5(3). 403–440. DOI: 10.1007/bf00134555.
- Steedman, Mark. 2000. *The syntactic process*. Cambridge, MA: The MIT Press. DOI: 10.7551/mitpress/6591.001.0001.
- Steedman, Mark. 2019. Combinatory Categorical Grammar. In András Kertész, Edith Moravcsik & Csilla Rákosi (eds.), *Current approaches to syntax: A comparative handbook*, 389–420. Berlin: De Gruyter Mouton.
- Storoshenko, Dennis R. & Chung-hye Han. 2013. Using synchronous tree adjoining grammar to model the typology of bound variable pronouns. *Journal of Logic and Computation* 25(2). 371–403. DOI: 10.1093/logcom/exs064.
- Storoshenko, Dennis R., Chung-hye Han & David Potter. 2008. Reflexivity in English: An STAG analysis. In *Proceedings of the Ninth International Workshop on Tree Adjoining Grammar and Related Frameworks (TAG+9)*, 149–156. Tübingen. <http://www.aclweb.org/anthology/W08-2320>.
- Tomita, Masaru (ed.). 1991. *Current issues in parsing technology*. Dordrecht: Kluwer Academic Publishers. DOI: 10.1007/978-1-4615-3986-5.
- Vijay-Shanker, K. 1987. *A study of Tree Adjoining Grammars*. University of Pennsylvania. (Doctoral dissertation).
- Vijay-Shanker, K. 1992. Using descriptions of trees in a Tree Adjoining Grammar. *Computational Linguistics* 18(4). 481–517. <http://dl.acm.org/citation.cfm?id=176313.176317>.
- Vijay-Shanker, K. & Aravind K. Joshi. 1985. Some computational properties of Tree Adjoining Grammars. In *Proceedings of the 23rd annual meeting of the Association for Computational Linguistics*, 82–93. Association for Computational Linguistics. DOI: 10.3115/981210.981221.
- Vijay-Shanker, K. & Aravind K. Joshi. 1988. Feature structure based Tree Adjoining Grammars. In *Proceedings of the 12th Conference on Computational Linguistics (COLING '88)*, 714–719. Budapest: Association for Computational Linguistics. DOI: 10.3115/991719.991783.
- Wedekind, Jürgen & Ronald M. Kaplan. 2020. Tractable Lexical-Functional Grammar. *Computational Linguistics* 46(2). 515–569. DOI: 10.1162/coli_a_00384.
- Wells, Rulon S. 1947. Immediate constituents. *Language* 23(2). 81–117. DOI: 10.2307/410382.

- Williams, Edwin. 2007. Dumping lexicalism. In Gillian Ramchand & Charles Reiss (eds.), *The Oxford handbook of linguistic interfaces*, 353–382. Oxford: Oxford University Press. DOI: 10.1093/oxfordhb/9780199247455.013.0012.
- Wray, Alison. 2002. *Formulaic language and the lexicon*. Cambridge, UK: Cambridge University Press. DOI: 10.1017/cbo9780511519772.
- XTAG Research Group. 2001. *A Lexicalized Tree Adjoining Grammar for English*. Tech. rep. IRCS-01-03. Philadelphia: Institute for Research in Cognitive Science, University of Pennsylvania. <https://www.cis.upenn.edu/~xtag/gramrelease.html>.

