

Glue logic vs. Spreading Architecture in LFG*

EVERY D. ANDREWS

The Australian National University

Avery.Andrews@anu.edu.au

1 Introduction

Although the f-structures of LFG look like an intuitively clear representation of many aspects of meaning, it has proved somewhat more difficult than one might have expected to connect them to conventional logical forms, over which entailment and other semantic properties and relations can be defined by standard methods. In this paper I will present what I believe to be a relatively easy to use formulation of the ‘glue logic’ approach (Dalrymple 1999, 2000), and then show how to eliminate the ‘restriction projections’ used by Andrews and Manning (1993, 1999) in their analyses of scoping modifiers and complex predicates, allowing the analyses to work with the conventional ‘locational projections’ of Kaplan (1995).

The presentation is based on recent work by de Groote (1999), and Perrier (1999), which allows glue-logic deductions, which linguists sometimes find quite difficult to understand, to be replaced by a proof-net-based technique which I believe to be easier to follow, because it employs totally obviously monotonic processes of link addition and assignment of values to (sub-)formulas, instead of deductive combinations and transformations of premises, which are easy to lose one’s way in. I suspect that the uptake of glue-logic by working linguists is significantly impeded by the difficulties of understanding it, so anything that offers a chance of making it easier is worth trying. This proof-net technique involves no substantive change in the theory, only the methods for calculating its results.

The proposed elimination of restriction projections will however require a substantive change, the elimination of the ‘semantic projection’. Rather than associate semantic information with a single semantic projection, it will be associated both with the f-structure and an ‘a-structure’ along the lines of Andrews and Manning (1999), and the capacity of glue-logic to integrate information present at different locations will be seen to make Andrews and Manning’s use of restriction projections unnecessary.

Glue logic is a theory of ‘semantic assembly’, that is, the way in which information about the meaning that is provided by lexical items and grammatical constructions is put together to get a meaning for the whole utterance. As such, it is somewhat agnostic about the exact nature of the meaning-elements that are being assembled (although the assembly techniques employed presumably impose some restrictions on what the meaning-elements can be). In introductory presentations, there is a tradition of using Montague’s logical forms without intensions, which I will follow here. First I will discuss basic predicate-argument structure, then quantifiers, at the same time developing the ‘content flow’ presentation. Then I turn to scoping modifiers and the elimination of restriction projections. I will close with some more general remarks.

*I am indebted to Ash Asudeh and Mary Dalrymple for discussion of some of the issues considered here, all errors are of course due to me.

2 Predicates and Arguments

In the semantic format we'll be using, the simplest kind of meaning-constructor is one for a proper name. All meaning-constructors consist of two components, on the left, the 'meaning-side', a specification of a meaning, on the right the 'glue-side', a specification of the information needed to control assembly. This will consist of a specification of (Montagovian) semantic type (based on e and t , in the introductory version), combined with f-structural 'locational' information (where the input or output will be collected from/delivered to). So if a proper name such as *Art* is inserted under an N with f-structure g , the resulting constructor will look like this, where the semantic type is subscripted to the location (but is however often omitted, although not in this paper):

$$(1) \text{ Art} : g_e$$


In the presentation we will be developing, this says that the content *Art*, of type e (entity), is delivered to the f-structure g (perhaps the subject of a sentence).

A bit more complex is the meaning-constructor of an intransitive verb, which would collect content of type e from its subject('s f-structure), and deliver content of type t (truth-value) to its own f-structure, which is also that of the clause. This input-output connection is represented with the \multimap symbol, popularly read as 'lollipop', representing implication in linear logic. So if the f-structure of the whole sentence is f , the constructor for *sleep* will be:

$$(2) \text{ sleep} : g_e \multimap f_t$$

It's an important principle of glue logic that the semantic type of the meaning-side is fully determined by the structure of the glue side. Here the semantic type has to be a function from type e to type t ; this is usually represented partially by writing the meaning side as a lambda-expression, such as $\lambda x.sleep(x)$. To reduce clutter, I will usually not do this. One could also subscript the meaning-sides with their type: $sleep_{e \rightarrow t}$.


The next thing we are going to do is connect the location-type expressions, called 'literals', of the meaning-constructors with links according to some rules; we'll call this 'hookup', its result is called a 'proof structure'. For these particular constructors, all we have to do is connect *Art*'s e (output) to *sleep*'s e , (input), with the final output coming out at *sleep*'s t . We can diagram the idea of content being delivered to and collected from f-structure locations with a highly explicit picture like this (positions of the meaning and glue sides of the meaning-constructors swapped):




We see *Art*'s constructor delivering its content to g (output port represented by a solid dot), and *sleep*'s collecting from g (input represented by circle) and delivering to f . These constructors can then be assembled by adding an arrow to connect the output to the input at g . An important fact which (3) doesn't represent is that there can be multiple input and output

ports of the same type at an f-structure; if there are multiple ways of hooking them up, this may result in ambiguity, although it is also possible that constraints to be discussed later will block some of the hookups.

To see a case where the syntactic structure makes a difference, we look at a transitive sentence such as *Art likes Gwen*. We'll follow Montague Grammar and Marantz (1984) in treating a transitive verb as something that takes an e (from the direct object) as input, and yields an intransitive predicate as output, i.e., something of type $e \multimap (e \multimap t)$. So if we have this full syntactic structure:




the f-structure of the object will be h , and the constructors we want will be:




Now the f-structure locational information will control which name's output gets hooked up to which of the verb's inputs, so that the syntactic structure is seen to contribute to the semantic interpretation in the right way.

To move from these ideas to a system that actually does something we need to specify a number of more things. First, how is the f-structure information in the glue-sides produced? The answer is from the lexicon, via the LFG processes of instantiation and functional resolution. In the lexicon, meaning-constructors contain expressions using function application and the \uparrow -metavariable, which in a lexical item means 'the f-structure correspondent of the c-structure node over me'. So that the lexical form of the constructors above will be:



So when the lexical items are used in structure (4), these uninstantiated constructors instantiate to:



And the last step, functional resolution, replaces the functional designator expressions with the f-(sub)structures they designate to produce (5).

Next, we need a formal scheme to officially identify the inputs and the outputs, which we do by defining a notion of ‘polarity’, as follows:

- (8) Polarity Rule 1: An entire glue side has positive polarity.
- Polarity Rule 2: If an implication has positive polarity, then its antecedent has negative polarity and its consequent positive.

So by (8), the polarities for the constructors will be:

- (9) *Art* : g_e
 +

- Gwen* : h_e
 +

- like* : $h_e \multimap (g_e \multimap f_t)$
 - + - + +

where the polarity of an implication will be written under its implication symbol. Finally, inputs and outputs are characterized as follows:¹

- (10) Inputs and Outputs: a positive literal is an output, a negative literal an input.

Given polarity, we can formulate the ‘Hookup Rule’:

- (11) Hookup Rule: Every negative literal must be connected to one and only one positive literal, and every positive to one and only one negative, except for one positive literal, the final output, located at the f-structure of the sentence.

A structure connected in accordance with this rule, but not necessarily obeying a further constraint to be introduced later, is called a ‘Proof Structure’.

Leaving out the polarities to reduce clutter, Hookup in (9) gives us:

- (12) *Art* : g_e
 ↓
 Gwen : h_e
 ↓
 likes : $h_e \multimap g_e \multimap f_t$


¹This is the polarity convention used in the LFG literature; unfortunately de Groote (1999) and Perrier (1999) use the opposite.

We've also begun leaving out rightmost parentheses (that is, linear implication is taken to be right-associative).

Now we move on to the content-assignment rules. We begin by assigning the meaning side of each constructor as content to its (entire) glue-side:

- (13) Content Flow Rule 1: The content of a glue side is its meaning side,


The result of applying this to (12) is (writing the content of an implication under its implication symbol):



Next we propagate content from positive literals (output ports) to negatives that they are linked to (input ports):

- (15) Content Flow Rule 2: The content of a positive literal is propagated to the negative that it is linked to.


Two applications of this produce:



And finally a somewhat more interesting rule, for calculating the content of the consequent of an implication from that of the implication and its antecedent:

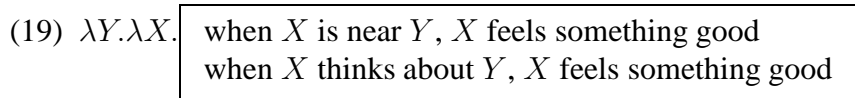
- (17) Content Flow Rule 3: If an implication has a as its own content, and b as its antecedent's content, then the content of its consequent is $a(b)$ (the function a applied to the argument b).

This is where the strict correlation between the semantic type of the meaning-side and the structure of the glue-side does its work: if for example the type of the transitive verb constructor here didn't take two type e arguments, the application of (17) wouldn't make sense. But since it is of type $e \rightarrow e \rightarrow t$, everything is fine, and two applications of CFR 3 produce:



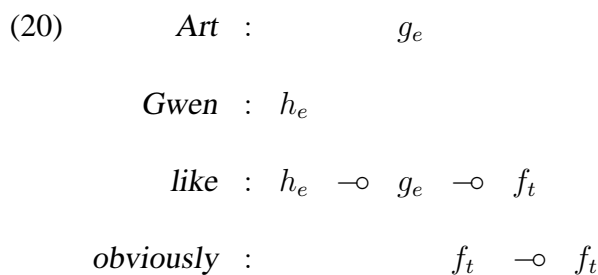
So we've got logical forms being produced by the f-structure.

From some points of view this may not look very illuminating; it simply says that the meaning is what you get by applying a function named *like* to the arguments *Gwen* and *Art* in succession. But this function might be something with some useful content, such as perhaps an explication along the lines of:




Then the resultant logical form will be equivalent to (19) with the lambdas removed and (the meanings of) *Gwen* and *Art* substituted for *Y* and *X* respectively, and since the f-structure can be connected to overt phrase structure in a wide variety of ways, as discussed in the LFG literature, we will have managed to express the connection between semantic roles and overt positions across a typologically diverse range of languages (Bresnan 2001, Falk 2001, Dalrymple 2000).

Unfortunately, the Hookup Rule is not enough to guarantee the delivery of sensible content, as can be seen by considering the possibilities for this structure:



The problem is that alongside of the sensible hookup, where *like* feeds into *obviously* which then provides the final output, there is a silly one where the final output comes from *like*, and the output of *obviously* is fed back to its input.

The following principle will eliminate the unwanted hookup:



²C.f. Theorem 1.8 (Perrier 1999).

You can see that if *obviously* in (20) has its output connected to its input, neither will get content, since the content of the whole implication needs content at its antecedent in order to produce content for its consequent.

This system can handle a significant range of constructions, but additional rules are needed to deal with quantifiers, which are involved in the scoping modifier constructions we will be analysing later.

3 Quantifiers

We will treat quantifiers as designating relations between sets, such as overlap, non-overlap, inclusion etc., so that *some knights sleep* is true iff the set of knights overlaps with the set of sleepers, etc. (this is ‘generalized quantifiers’ (Barwise and Cooper 1981)). Since the semantic type of a set-designator is $e \rightarrow t$, that of a relation between sets will be:

$$(22) (e \rightarrow t) \rightarrow (e \rightarrow t) \rightarrow t$$

So what about the locational information? There is a convention to the effect that the first argument of a quantifier is taken to be the nominal content of the NP it appears in (its ‘restriction’), the second the predication provided by its environment (its ‘(nuclear) scope’). The nominal content is of semantic type $e \rightarrow t$, but where will its inputs and outputs be located? There is a useful convention, borrowed from HPSG, that the e input is located at the value of an attribute VAR, the t output at the value of an attribute RESTR, which will here be located in the NP’s f-structure.³

So if the f-structure is (23), then the instantiated constructor for *knight*, and the first part of the *every* constructor, will be as in (24), with polarities indicated to the extent that we have principles to determine them:

$$(23) \left[\begin{array}{l} \text{SUBJ} \\ \text{PRED} \end{array} \right] f: \left[\begin{array}{l} \text{QUANT} \quad \text{Every} \\ \text{PRED} \quad \text{‘Knight’} \\ \text{VAR} \quad gv:[\] \\ \text{RESTR} \quad gr:[\] \end{array} \right] g: \left[\begin{array}{l} \text{PRED} \quad \text{‘Sleep(SUBJ)’} \end{array} \right]$$

$$(24) \textit{knight} : \begin{array}{ccc} gv_e & -\circ & gr_t \\ - & + & + \end{array}$$

$$\textit{every} : \begin{array}{cccc} (gv_e & -\circ & gr_t) & -\circ \dots \\ ? & - & ? & + + \end{array}$$

The reason we don’t know the polarities of the first two literals is that they are antecedent and consequent of a negative conditional, and the rules so far specify polarity only for those of a positive. But it is obvious what the answer has to be: to connect things up we will

³They are standardly put in the ‘semantic projection’, which will be argued against later in the paper.

And finally we apply CFR 3 again to produce the final result $every(knight)(sleep)$.

An important question is on what basis we chose f as the location of the second argument's consequent. It turns out that for a first approximation theory of quantifier scope, one can do quite well by letting this location be specified as a variable, which can be identified with any location in the f-structure, as long as the final result is delivered back to the same place, that is, specified with the same variable. Using upper case italics for variables over f-structure, the lexical meaning-constructor for the *every* will now look like this:

$$(34) \text{ every} : ((\uparrow\text{VAR})_e \multimap (\uparrow\text{RESTR})_t) \multimap (\uparrow_e \multimap H_t) \multimap H_t$$

We get (33) by identifying H with f . It turns out that the restriction on hookups imposed by CFR 5 (the content of the consequent must contain the content of the antecedent free) excludes various wrong scope possibilities that tend to be generated by naive schemes of 'Quantifier Raising', a result discussed at length in Dalrymple et al. (1999).

Here is an example where variable choice of scope-consequent and final result produces an ambiguity, the wide and narrow scope readings of *everyone seems to sleep*; to keep the structure simple the restriction argument of the quantifier is not expressed:

$$(35) \text{ a. } \left[\begin{array}{l} \text{SUBJ} \quad h: [\text{PRED} \text{ 'Everyone'}] \\ \text{PRED} \quad \text{'Seem(XCOMP)'} \\ \text{XCOMP} \quad g: \left[\begin{array}{l} \text{SUBJ} \quad h: [\quad] \\ \text{PRED} \quad \text{'Sleep(XCOMP)'} \end{array} \right] \end{array} \right]$$

f: g:

b. $(h_e \multimap H_t) \multimap H_t : \text{everyone}$ $(h_e \multimap H_t) \multimap H_t : \text{everyone}$
 $h_e \multimap g_t : \text{sleep}$ $h_e \multimap g_t : \text{sleep}$
 $g_t \multimap f_t : \text{seem}$ $g_t \multimap f_t : \text{seem}$

$seem(everyone(\lambda x.sleep(x)))$ $everyone(\lambda x.seem(sleep(x)))$
 $H = g$ $H = f$

Getting the right logical forms out of *seem*-structures with functional control is one of the trickier aspects of semantic interpretation in LFG; glue logic does it easily. Although an extended discussion of quantifier scope is beyond the scope of this paper, the basic mechanism is important, because it ultimately solves the 'integration problem' that provided Andrews and Manning's motivation for restriction projections.

We have now presented the framework, which is so far substantively almost the same as standard glue logic, with two differences aside from our different technique for deriving the meanings. A minor notational difference is that in the standard formulations, there are universal quantifiers on the glue side binding any f-structure variables. Here we've followed Fry (1999:94) in omitting them, a move that is legitimate because universal quantifiers can be floated upward through implications, and then understood as implicitly applying to free

variables. A more substantial difference is we have dispensed with the ‘semantic projection’ that the meaning-constructors normally associate meanings with, instead associating them directly with the f-structure. The semantic projection doesn’t have any clear empirical justification (no analyses are broken when it is removed), and is inconsistent with the analysis in the next section, which depends on the possibility of controlling semantic assembly with multiple projections.

4 Scoping Modifiers

The problem that ‘scoping modifiers’ present for LFG is that the rather flat f-structures that they get under the theory don’t seem to provide the right guidance for semantic interpretation, which does however seem to be tightly linked to the c-structure in an obvious way (so that various other well-formalized linguistic theories have no issues with these constructions). Complex predicates pose similar issues, as discussed by Alsina (1997) and Andrews and Manning (1993, 1999). The case of scoping modifiers is simpler, so it is what I will discuss here.

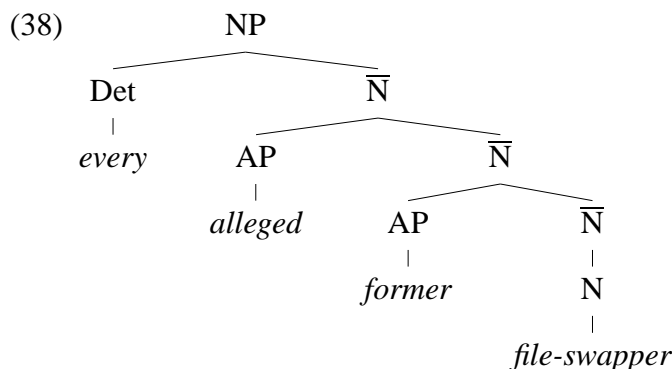
For an example, the pair of NPs:

- (36) a. a former alleged file-swapper
- b. an alleged former file-swapper

both get an f-structure like this, where the two adjectives are members of an intrinsically unordered adjunct set:

$$(37) \left[\begin{array}{ll} \text{QUANT} & \text{EVERY} \\ \text{ADJUNCT} & \left\{ \begin{array}{l} \left[\text{PRED} \text{ 'Former'} \right] \\ \left[\text{PRED} \text{ 'Alleged'} \right] \end{array} \right\} \\ \text{PRED} & \text{'File-swapper' } \end{array} \right]$$

This structure generates the expectation that both of the NPs will be ambiguous, with the same meanings, instead of differing in meaning in the way they clearly do, and which is obviously suggested by the c-structure trees they would plausibly get (Andrews 1983):

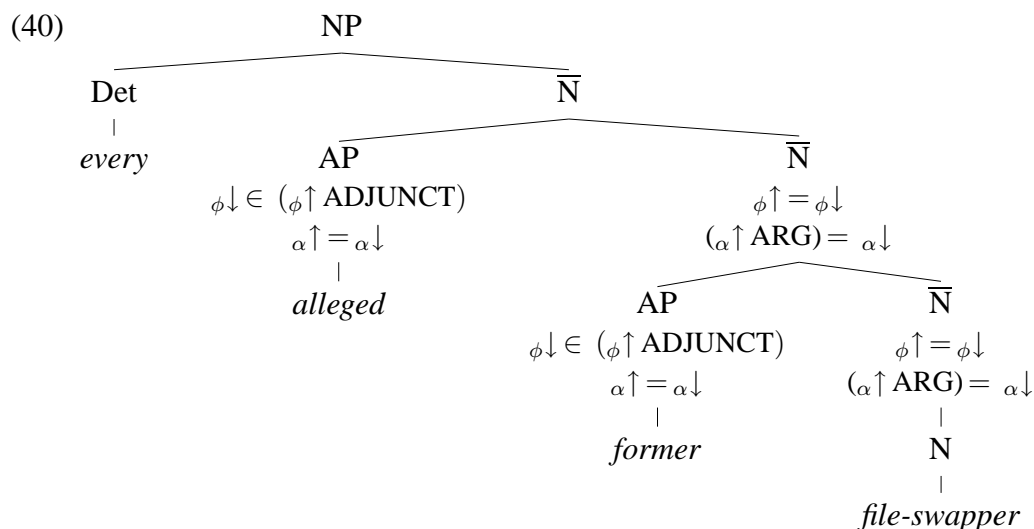


The central component of Andrews and Manning’s solution to this sort of problem was an additional projection that reflected the tree-structure more closely, which we called ‘a-structure’ because it seems suitable to perform at least some of the functions that are often attributed to the somewhat variably defined notion of ‘argument structure’. I will now show how glue logic supports an analysis whereby a-structure is retained, but using Kaplan’s (1995) original ‘locational’ notion of projection rather than Andrews and Manning’s second innovation, restriction projections.

The basic idea is to have a-structure provide different headship relations than f-structure, so that the scoping modifiers are a-structurally heads, with their \bar{N} -sisters being ‘a-structural complements’, which we construed as values of an a-structure attribute ARG. The form of a-structure proposed for (38) is (39) (PRED-attributes included to make it easier to interpret the structure, although I don’t believe they actually exist in a-structure, or perhaps anywhere else):

$$(39) \left[\begin{array}{l} \text{PRED} \quad \text{'Former(ARG)'} \\ \text{ARG} \quad b: \left[\begin{array}{l} \text{PRED} \quad \text{'Alleged(ARG)'} \\ \text{ARG} \quad c: \left[\begin{array}{l} \text{PRED} \quad \text{'File-swapper'} \end{array} \right] \end{array} \right] \end{array} \right]$$

We can get the f-structure (37) and the a-structure (39) from the tree (38) by annotating the tree as follows, where there are annotations specifying both f- and a- structural information, the arrows pre-subscripted with ϕ or α to indicate whether they are providing information about f- or a- structure, respectively (unannotated nodes assumed to share both f- and a-structure):



Because they are differently related to the c-structure than the f-structure is, these a-structures will ‘respect the tree’ (Alsina 1997) as required, but how will we combine their information with that from the f-structure, where grammatical relations control the relationship between NPs and semantic roles? Andrews and Manning showed how restriction projections could be used to solve this problem, and that the alternatives in the literature were not satisfactory;

here I will show that glue-logic can also do it, essentially by virtue of its capacity to shift content from one location in a structure to another.

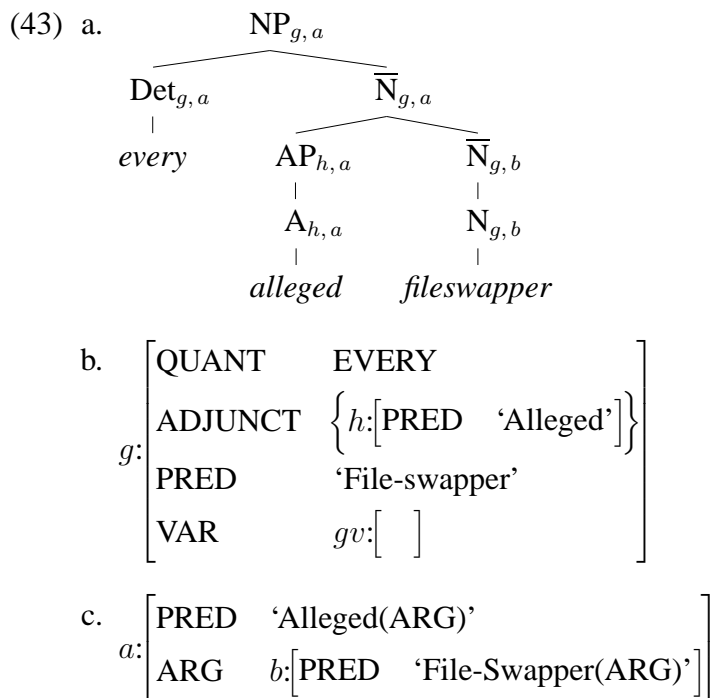
The basic idea will be to locate the *t*-content of NPs at a-structure, but the *e*-content, which is shared through the levels, at f-structure. Therefore the uninstantiated meaning-constructor for *file-swapper* will be:

$$(41) \text{ file-swapper} : (\phi \uparrow \text{VAR})_e \multimap \alpha \uparrow_t$$

The RESTR attribute can now be abolished,⁴ and the constructor of the quantifier adjusted accordingly (I so far see no reason why the quantifier’s constructor would have to say what projection the scope *t*-information is located on, so the *H* variable doesn’t specify this):

$$(42) \text{ every} : ((\phi \uparrow \text{VAR})_e \multimap \alpha \uparrow_t) \multimap (\phi \uparrow_e \multimap H_t) \multimap H_t$$

For simple NPs such as *every fileswapper*, the analysis will work just as before, except that the noun’s consequent output and quantifier’s restriction’s consequent input passes through an a-structure rather than an f-structure RESTR attribute. But when an operator adjective is present, things will be different. Consider a simple case with one scoping modifier, with c-structure, f-structure and a-structure as indicated below, with the c-structure nodes subscripted first with their f-structure and second with their a-structure correspondent:




Now the noun’s meaning-constructor will collect as input the variable provided as output by the quantifier, because this is located at f-structure, which is the same for both the quantifier

⁴As can the VAR attribute, although the explanation of why one can get away with this is somewhat complex.

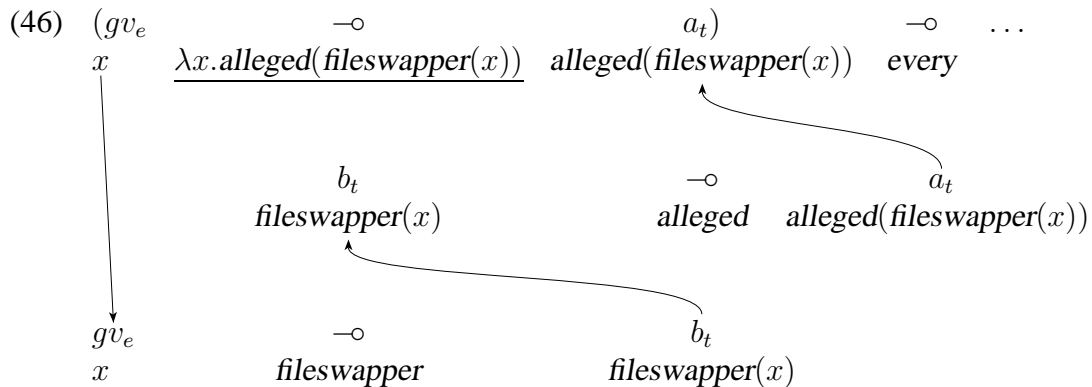
and the noun, and return its output to its a-structure, which is b . But the quantifier collects input back (for its restriction argument) from its own a-structure a , which is different from b . So without some sort of middleman, the flow of content will be blocked. The scoping adjective is of course the middleman: we can construe it as collecting t -type input from the ARG-value b of its own a-structure correspondent a , and outputting the processed result back at a ; this effect will be produced by an uninstantiated meaning-constructor like this:

$$(44) \textit{alleged} : (\alpha \uparrow \textit{ARG})_t \multimap \alpha \uparrow_t$$

Hooking up the literals produced by lexical insertion and instantiation, we get:



The noun's output is hooked up to the scoping modifier's input, and the latter's output to the quantifier's first argument's input (a negative consequent), and the content then flows as indicated here, with the resulting value for the quantifier's restriction argument underlined:



So we see that the modifier gets applied to the noun-meaning in the appropriate manner.

Furthermore, it should be clear that if there is more than one scoping modifier, as in (40), the analysis will respect the tree as required, because each modifier will pass up the processed content of its \bar{N} -sister's a-structure to its own a-structure, and the order of application of the modifiers will be determined by the c-structural relations in the tree. Finally, the treatment of the second (scope) argument of the quantifier will solve the integration problem for which restriction projections were proposed, since the content of the quantifier's scope can be determined (at least mostly) by f-structure as discussed at the end of the previous section, and this is not affected by the change introduced here.

5 Final Remarks

I have shown how glue-logic can be presented with a content-flow conceptualization, and thus presented, can be seen to assist with solving some descriptive problems by virtue of its ability to integrate information located on different levels, or projections, of linguistic structure, a result which might have wider applicability in other theories employing parallel architectures.

I will finish by mentioning two further issues. First, Andrews (2003) extends the analysis to deal with some more complex types of scoping modifiers, such as *self-proclaimed* and *confessed*, and also regular intersective modifiers, such as *greedy* and *unscrupulous*. Some remarks are also made about developing the same kind of analysis for complex predicates. A downloadable implementation, ‘Baby Glue’ covers some of these issues as well.⁵

Second, I’ve here presented content-flow only for linear implication, whereas some glue logic analyses also use linear conjunction (\otimes , read ‘tensor’); as discussed in Andrews (2003), content flow can be extended to \otimes , although it remains to be seen whether \otimes is truly required for natural language semantic assembly.

Bibliography

Alsina, A. 1997. A theory of complex predicates: Evidence from causatives in Bantu and Romance. In A. Alsina, J. Bresnan, and P. Sells (Eds.), 203–246.

Alsina, A., J. Bresnan, and P. Sells (Eds.). 1997. *Complex Predicates*. Stanford, CA: CSLI Publications.

Andrews, A. D. 1983. A note on the constituent structure of modifiers. *Linguistic Inquiry* 14:695–7.

Andrews, A. D. 2003. Glue logic, projections, and modifiers. URL: <http://arts.anu.edu.au/linguistics/People/AveryAndrews/Papers>.

Andrews, A. D., and C. D. Manning. 1993. Information-spreading and levels of representation in LFG. Technical Report CSLI-93-176, CSLI, Stanford University, Stanford, CA. URL: <http://www-nlp.stanford.edu/manning/papers/>.

Andrews, A. D., and C. D. Manning. 1999. *Complex Predicates and Information Spreading in LFG*. Stanford, CA: CSLI Publications.

Barwise, J., and R. Cooper. 1981. Generalized quantifiers and natural language. *Linguistics and Philosophy* 159–219.

Bresnan, J. W. 2001. *Lexical-Functional Syntax*. Oxford: Blackwell.

⁵Available from <http://arts.anu.edu.au/linguistics/people/averyandrews/software/>

Crouch, R., and J. van Genabith. 2000. Linear logic for linguists.

URL: <http://www2.parc.com/istl/members/crouch/>.

Dalrymple, M. (Ed.). 1999. *Syntax and Semantics in Lexical Functional Grammar: The Resource-Logic Approach*. MIT Press.

Dalrymple, M. 2000. *Lexical Functional Grammar*. Academic Press.

Dalrymple, M., R. M. Kaplan, J. T. Maxwell, and A. Zaenen (Eds.). 1995. *Formal Issues in Lexical-Functional Grammar*. Stanford, CA: CSLI Publications.

Dalrymple, M., J. Lamping, F. Pereira, and V. Saraswat. 1999. Quantification, anaphora and intensionality. In Dalrymple (Ed.), 39–90.

de Groote, P. 1999. An algebraic correctness criterion for intuitionistic multiplicative proof-nets. *TCS* 115–134.

URL: <http://www.loria.fr/~degroote/bibliography.html>.

Falk, Y. N. 2001. *Lexical-Functional Grammar: An Introduction to Parallel Constraint-Based Syntax*. Stanford University: CSLI Publications.

Fry, J. 1999. Proof nets and negative polarity licensing. In M. Dalrymple (Ed.), 91–116.

Kaplan, R. M. 1995. The formal architecture of LFG. In M. Dalrymple, R. M. Kaplan, J. T. Maxwell, and A. Zaenen (Eds.), 7–27. CSLI Publications.

Marantz, A. 1984. *On the Nature of Grammatical Relations*. Cambridge MA: MIT Press.

Perrier, G. 1999. Labelled proof-nets for the syntax and semantics of natural languages. *L.G. of the IGPL* 629–655.

URL: <http://www.loria.fr/~perrier/papers.html>.