

Proof Nets and the Complexity of Processing Center-Embedded Constructions

Mark Johnson *
Brown University

(Received ; Accepted in final form)

Abstract. This paper shows how proof nets can be used to formalize the notion of “incomplete dependency” used in psycholinguistic theories of the unacceptability of center-embedded constructions. Such theories of human language processing can usually be restated in terms of geometrical constraints on proof nets. The paper ends with a discussion of the relationship between these constraints and incremental semantic interpretation.

1. Introduction

The distinction between competence and performance popularized by Chomsky1965 (Chomsky1965) is fundamental to modern linguistics. *Competence* is the “knowledge” of a human language possessed by its users (e.g., as formalized in a grammar), while *performance* refers to how this knowledge is used to produce and comprehend utterances. To date, most applications of logic to natural language have concentrated on competence grammar. This paper suggests that techniques from logic such as proof nets (Girard1995) might also be useful for describing the difficulty of processing center-embedded constructions, which is an aspect of human linguistic performance.

It is well-known that center embedded constructions such as (1b) are more difficult to comprehend than corresponding right embedded constructions (1a). (The prefix ‘#’ indicates the example is difficult or impossible to comprehend).

- (1) a. The patient was cured by the drug [that was administered by the intern [that was supervised by the nurse]].
- b. #The drug [that the intern [that the nurse supervised] administered] cured the patient.

* This research was performed while on a sabbatical visit to the Rank Xerox Research Centre in Grenoble. I would like to thank Brown and RXRC for making this possible, my colleagues who encouraged me to learn about Lambek Categorical Grammar and Linear Logic, and Ted Gibson for helpful information about human performance on center-embedded constructions. In addition, I would like to thank the JoLLI reviewers for their helpful and insightful comments on the first draft of this paper.

As Chomsky1965 (Chomsky1965) notes, the difficulty of comprehending (1b) does not seem to be due to any single construction in (1b), but rather to the way those structures are configured. He argues that (1b) should be treated as grammatical (i.e., possessing a well-formed syntactic analysis), but that its incomprehensibility arises from an inability of the human sentence processing mechanism to find that analysis.

Chomsky and Miller1963 (Chomsky and Miller1963) propose that short-term memory overload causes this incompleteness. The strictly right embedded constructions of the kind (1a) can be generated by a right-linear grammar, and hence can be recognized by a finite state machine. In contrast, a simple application of the pumping lemma for regular sets shows that center embedded constructions of the kind (1b) cannot be recognized by a finite state machine.

However, it seems that all such comprehension difficulties cannot completely be explained in terms of the memory requirements of a parsing automaton. As Chomsky1965 (Chomsky1965) and others note, the particular constructions and their arrangement affect comprehensibility; thus the processing difficulty of (1b) relative to (1a) may be due in part to the fact that in (1b) the relative clauses involve object extraction, while in (1a) the relative clauses involve subject extraction.

The following example, from Gibson and Thomas1996 (Gibson and Thomas1996), shows that depth of center embedding alone cannot account for all such processing difficulty. (2a) and (2b) are both doubly embedded structures consisting of a relative clause (RC) and a sentential complement (SC), differing only in the order of embedding, and most phrase-structure parsing automata require the same amount of working memory to accept them both. However (2a) is easier to comprehend than (2b).

- (2) a. The chance [_{SC} that the nurse [_{RC} who the doctor supervised] lost the reports] bothered the intern.
- b. #The intern [_{RC} who the chance [_{SC} that the doctor lost the reports] bothered] supervised the nurse.

The literature contains a number of proposals accounting for a variety of comprehensibility differences (Gibson and Thomas1996; Lewis1996; Stabler1994). These theories of processing overload differ in the predictions that they make, this paper does not attempt to choose between them, or to develop an alternative account. (Indeed, there seems to be considerable uncertainty among theorists as to exactly which sentences cause processing overload).

However, all these proposals share the common feature that they posit that the human sentence processing mechanism is sensitive to

the number and kind of “incomplete dependencies” encountered as the sentence is processed from left to right. The precise configurations of dependencies which purportedly lead to processing overload differ from theory to theory. Unfortunately, the central notion of incomplete dependency, and its relationship both to grammar and to parsing mechanisms, is usually only informally explained in these proposals.

This paper proposes that resource logics provide the appropriate tools to formalize and investigate such incomplete dependencies. The next section introduces resource logics and proof nets, and the following section uses these to express an account of the processing complexity of the examples just discussed. The paper ends with a discussion of the relationship between that account and incremental semantic processing.

2. Resources and Dependencies

Most, if not all, contemporary linguistic theories are based on a view of language in which linguistic entities produce and consume resources. For example, a transitive verb is standardly viewed as something which requires two noun phrases, its subject and object, to fill its “thematic grid,” i.e., fill its argument requirements in order to produce a complete clause, which can in turn serve as an argument of a higher verb.

Linguistic theories differ as to exactly what types of resources are involved in particular sentences, what structural configurations are required for a resource requirement to be filled, exactly what the rules for resource accounting are, and what other mechanisms play a role in language. For example, versions of Categorical Grammar such as Morrill1994 (Morrill1994)’s Type-Logical Grammar are based solely on (directionally sensitive) resource accounting. On the other hand, modern versions of Transformational Grammar specify a set of configurations in which resource consumption can occur (e.g., feature cancellation occurs in a Spec-Head configuration) and provides structural operations (Merge and Move) to move linguistic entities into such configurations.

The notion of “incomplete dependency” adopted in most processing models derives from a fundamentally resource based view of language. A dependency is a producer-consumer pair, and it is “incomplete” at a point in a left to right analysis if exactly one member of the pair has been encountered in the portion of the input string seen so far. For example, at the point marked ‘ \triangle ’ in the simple transitive sentence (3) the verb-object dependency is incomplete, since the consumer resource (i.e., the verb) has been encountered but the producer resource (the object noun phrase) has not.

(3) Kim kissed \triangle Sandy.

Recent work in logic has established a general framework in which various kinds of resource sensitivity can be formalized and studied. This section informally sketches some of the key ideas, but the reader should see Girard, Lafont, and Taylor1989 (Girard, Lafont, and Taylor1989) and van Benthem1995 (van Benthem1995) for a more thorough introduction to this material.

Linguistic entities are classified into *types*, which encode their combinatory capability. For example, the noun phrase *Kim* might be assigned the type NP, and a verb phrase *snores* might be assigned the type $\text{NP} \rightarrow \text{S}$, indicating that it consumes a resource of type NP to produce an object of type S. These types also indicate the *type* of semantic interpretation associated with these phrases: e.g., the verb phrase is a function from NP meanings to S meanings.

Resource logics differ in terms of the kinds of *structural sensitivity* they enforce. For example, a grammar consisting solely of a resource logic (e.g., Lambek Categorical Grammar) needs to be sensitive to the positions of linguistic entities in the sentence. In such a grammar the location of the argument with respect to the functor needs to be specified, so the undirected implication ‘ \rightarrow ’ is specialized into a leftward looking version ‘ \backslash ’ and a rightward looking version ‘ $/$ ’. For example, a transitive verb such as *touches* might be of type $(\text{NP}\backslash\text{S})/\text{NP}$, i.e., an entity which consumes an NP to its right (the object noun phrase) and then consumes an NP to its left (its subject) to produce a saturated sentence.

Directional sensitivity is one dimension of structural sensitivity. In general, the more linguistic detail a resource logic is called upon to account for, the more refined its type system needs to be. Because the domain of locality of linguistic relationships may vary (e.g., head-complement dependencies are more local than WH dependencies), different types may need to be associated with different structural sensitivities. A general framework of multimodal substructural logics has been developed for formalizing and investigating these interactions (Moortgat1997).

A proof in a resource logic specifies a “plugging”, i.e., it identifies which objects fill the requirements of which other objects (Girard1995). These dependencies or pluggings determine the semantic interpretation of the utterance; e.g., via the Curry-Howard correspondence between (intuitionistic) proofs and λ -terms (Girard, Lafont, and Taylor1989; van Benthem1995). As the reader will observe, the approach developed here is not tied to any one particular theory of grammar, but can

be applied to any grammar that can be described in terms of such pluggings.

The job of a grammar is to specify just which pluggings can occur in a particular human language. However, for the purposes of this paper it is not necessary to exactly identify the constraints on grammatical pluggings (indeed, one of its strengths is that it does not depend on the precise details of the particular linguistic theory involved). Rather, it suffices that we can determine the pluggings which have actually occurred in the particular sentences being studied. These pluggings can usually be deduced from fairly general linguistic assumptions (e.g., about the valence of verbs) and the interpretation of the sentences themselves (which must reflect their pluggings, via the Curry-Howard correspondence).

To keep things simple, all the examples in this paper use Lambek Categorical Grammar. This grammatical framework, while sufficient to describe the examples presented here, is incapable of describing many other important constructions that appear in natural languages. More sophisticated grammatical frameworks, such as the multi-modal systems described Morrill1994 (Morrill1994) and Moortgat1997 (Moortgat1997), can account for a much wider set of natural language examples. However, as the interested reader can confirm, the observations presented here about the complexity of particular examples also hold in these more complex systems (as they must, since these more sophisticated grammars assume the same pluggings as presented here).

2.1. PROOF NETS

Proof nets are a graphic depiction of a proof, i.e., of the dependencies or plugging relationships between entities. In the application described here the logic is required to systematically distinguish the inputs to the combinatory process (the lexical items) from the single output produced from these inputs (the completed sentence), i.e., the logic must be intuitionistic and correspondingly the pluggings themselves can be regarded as directed.

A proof net is a directed graph composed of proof net connectives. The edges of the proof nets used here are directed (as the logic is intuitionistic) and each edge is labelled with a type. A proof net must satisfy structural conditions which depend on the kind of structural sensitivity imposed by the logic. Proof nets for Lambek Categorical Grammars must satisfy a planarity condition (Roorda1991; Roorda1992), and Moortgat1997 (Moortgat1997) describes in detail the structural conditions that correspond to the more general class of multimodal resource logics. However, as explained above the work




Figure 1. The first two columns list the proof net connectives used here, and the third column provides the corresponding linear logic proof net connective.

presented here requires us only to identify the proof nets associated with particular examples, and does not depend directly on the general structural constraints themselves, so they are not presented here.

The first two columns of Figure 2.1 lists the proof net connectives used in the examples here, and the third column lists the corresponding linear logic proof net connective for readers familiar with linear logic.

We permit two kinds of axiom links. In addition to the identity axiom (the wire permitting an X phrase to plug an X requirement) we have a non-logical axiom permitting a WH phrase to plug an NP requirement, and will indicate its use with a dashed line. (This distinction is one that any realistic grammar will have to make, as WH and NP dependencies exhibit different structural constraints).

The proof net connectives decompose the types associated with lexical entries into the resources that they produce and consume, so each lexical entry can be regarded as a partial proof net with incoming and outgoing edges labelled with atomic categories. These unpacked lexical entries are connected with other words and phrases by axiom links.

The axiom links identify the *dependencies* between the lexical item and its surrounding material, so the unsatisfied dependencies at any

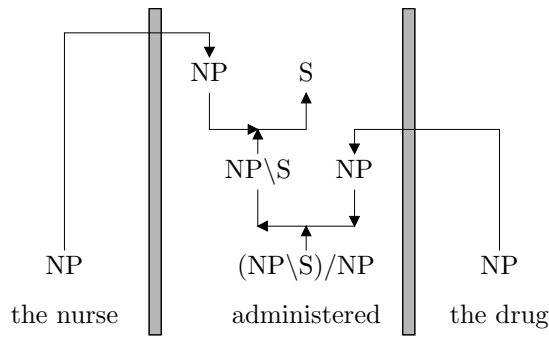


Figure 2. A transitive clause, with vertical bars marking divisions between the words.

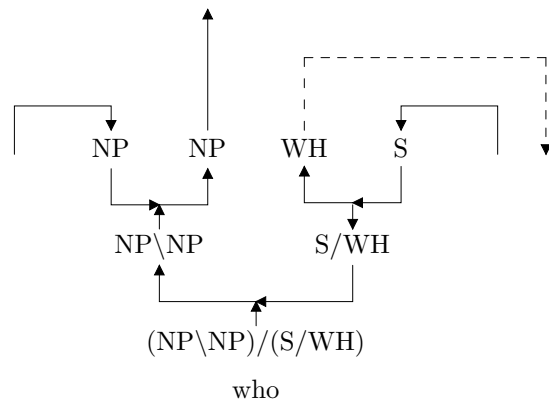


Figure 3. The partial proof net associated with the relative pronoun *who*.

point in processing are precisely the axiom links that must be cut in order to disconnect the proof net.

We end this section with some examples. Figure 2.1 shows the proof net for the simple transitive sentence *Kim admires Sandy*. The vertical bars indicate cuts made between the words in this sentence. Because exactly one axiom link crosses each cut, there is exactly one unsatisfied NP dependency at each of these locations.

Figure 2.1 shows the partial proof net associated with a relative pronoun. For simplicity NPs are treated as atomic units in this paper, so relative clauses are analysed as NP modifiers. (The interested reader can check that the relative complexity measures are not affected if relative clauses are analysed as N' modifiers). Thus a relative pronoun

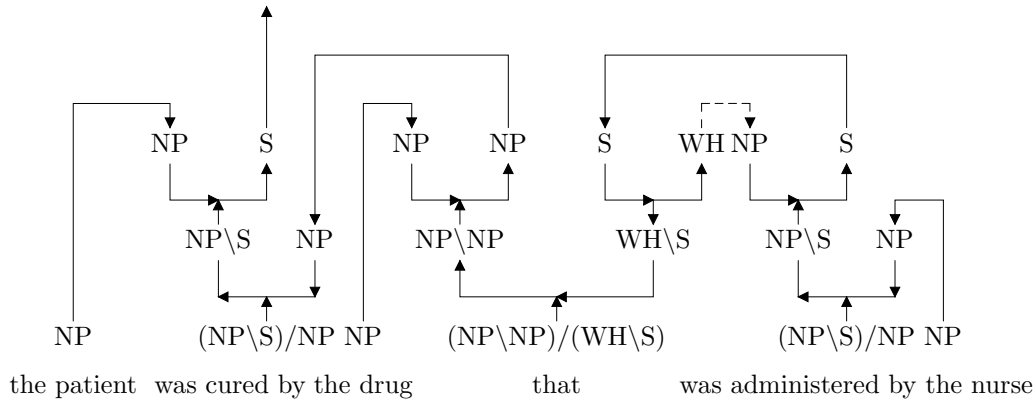


Figure 4. A proof net for a relative clause modifying the object.

is analysed as something which provides a WH resource to its right and consumes the S that it appears in. It also consumes an NP to its left and produces a new NP, i.e., it functions as an NP modifier.

3. Proof Nets, Dependencies and Processing

With the formal tools of proof nets now available, it is relatively simple to express many dependency-based theories of human sentence processing complexity as geometric constraints on proof nets.

First, consider the iterated object relative clauses shown in (4), which are usually described as right-branching constructions. While increased length does reduce acceptability, there does not seem to be any significant processing overload associated with such examples.

- (4) a. The patient was cured by the drug [that was administered by the nurse].
- b. The patient was cured by the drug [that was administered by the nurse [who was supervised by the doctor]].
- c. The patient was cured by the drug [that was administered by the nurse [who was supervised by the doctor [who was admired by the student]].

Figure 3 depicts the proof net for (4a). For simplicity in this proof net the passivised forms *was cured by* and *was administered by* are treated as transitive verbs, rather than being analysed into their component lexical items. At most 2 axiom links cross any cut in the proof nets for

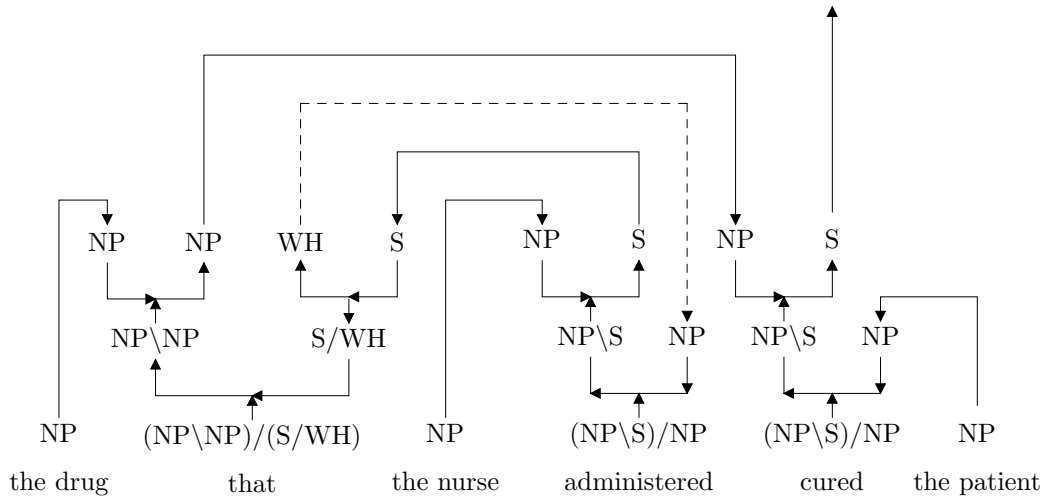


Figure 5. A proof net for a relative clause modifying the subject.

these examples (either two NP links or an NP and a WH link, depending on where the cut lies), which is consistent with processing complexity being independent of the number of such constructions involved.

Now consider the iterated subject relative clauses shown in (5), which are classic examples of center-embedded constructions. In contrast to the object relative clauses in (4), two or more levels of embedding renders the example virtually incomprehensible.

- (5) a. The drug [that the nurse administered] cured the patient.
- b. #The drug [that the nurse [that the doctor supervised] administered] cured the patient.
- c. #The drug [that the nurse [that the doctor [that the students admired] supervised] administered] cured the patient.

Figure 4 depicts the proof net for (5a). The maximal cut in the constructions (5) occurs just after the most embedded subject. At depth of embedding n the maximal cut crosses $3n + 1$ axiom links. The assumption that the human language processor can only keep track of a small number of such incomplete dependencies accounts for the increasing ill-formedness as the number of such constructions increases.

It is sometimes claimed that subject complement clauses such as (6) are easier to comprehend than corresponding subject relatives (Stabler1994). Figure 5 depicts the proof net for (6). The maximal cut for

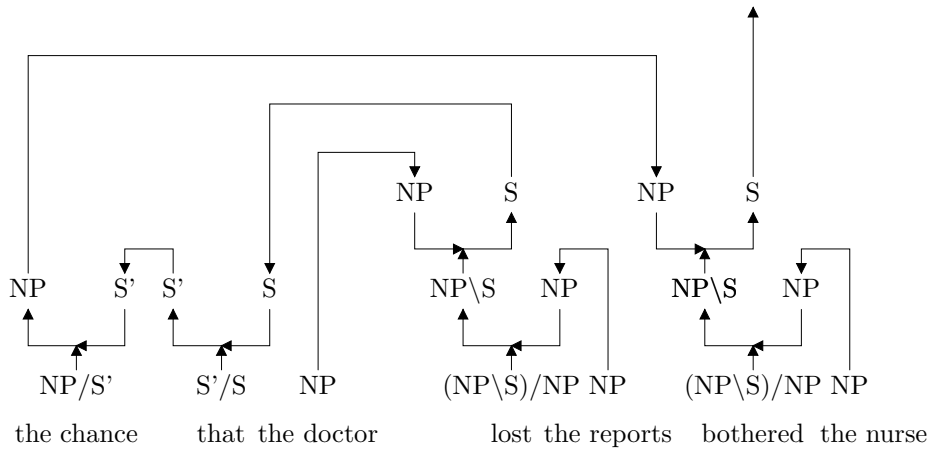


Figure 6. A proof net for a complement subject example

this example crosses one less axiom link than the corresponding relative clause example (5a), which is consistent with this putative difference.

(6) The chance [that the doctor lost the reports] bothered the nurse.

Thus far processing complexity seems proportional to the number of axiom links crossed by a maximal cut. The examples (2), reprinted here as (7), suggest that the matter is more delicate. Proof nets for these examples are schematically depicted in Figures 7 and 7 (only links crossed by the maximal cut are shown).

(7) a. The chance [_{SC} that the nurse [_{RC} who the doctor supervised] lost the reports] bothered the intern.

b. #The intern [_{RC} who the chance [_{SC} that the doctor lost the reports] bothered] supervised the nurse.

Gibson and Thomas1996 (Gibson and Thomas1996) present a hypothesis that accounts for this distinction, and which is expressed in terms of the internal states of an unspecified automaton that constructs a parse tree. They hypothesise that the human sentence processor overloads very quickly when it predicts a category whose features are subsumed by the features of one of the categories it is currently in the process of completing. In their approach a relative clause introduces a prediction for an S node with a WH feature, while a sentential complement introduces a prediction for an S node without a WH feature. They propose

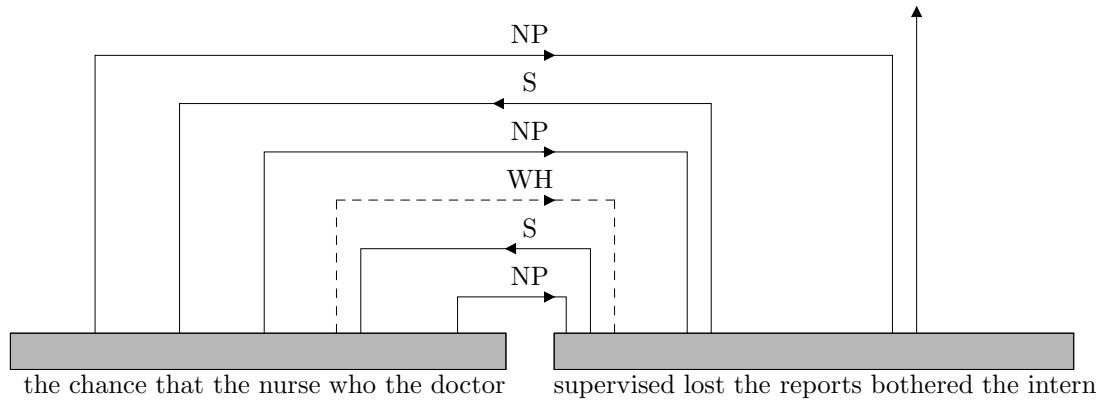


Figure 7. The schematic structure, including the maximal cut, of the proof net for a relative clause embedded within the matrix subject's complement.

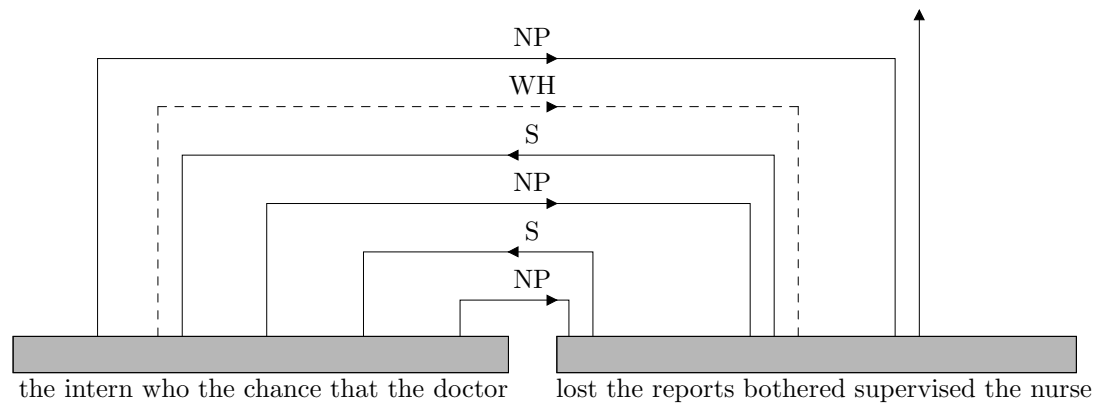


Figure 8. The schematic structure, including the maximal cut, of the proof net for a subject complement clause embedded within a subject relative clause.

that prediction for an S node without a WH feature is subsumed by a prediction for an S node with a WH feature, so only (7b) requires the introduction of a prediction that is subsumed by the features of a category currently being parsed.

This paper does not attempt to empirically evaluate this hypothesis, but merely points out that in the context of the current examples, this hypothesis can be restated as a constraint on the sequence of axiom links crossed by any cut in the proof net. To see this, note that Gibson and Thomas1996 (Gibson and Thomas1996)'s category predictions correspond to leftward-pointing axiom links in a proof net. Thus there

are two S predictions at the locations of maximal cut complexity in the proof nets in Figures 7 and 7.

We could directly encode Gibson and Thomas1996 (Gibson and Thomas1996)'s proposal in a proof net system by enriching the S type with WH features (possibly inherited via some percolation mechanism, as they suggest): in proof net terms their hypothesis would be that the human sentence processor overloads if a leftward pointing axiom link is embedded within another leftward pointing axiom link labelled with a subsuming type.

However, because WH-dependencies appear in the proof nets used here it is not necessary to enrich the S type with WH features, as the relevant S axiom links are easily identified geometrically in the proof net (they are the ones introduced by expansion of a S/WH type). In the proof nets depicted in Figures 7 and 7 they are the S axiom links adjacent to a WH axiom link.

4. Conclusion

The previous section has shown that theories of processing stated in terms of incomplete dependencies can often be restated as geometric constraints on proof nets. Such constraints would be more plausible if they can be shown to arise naturally from independently required mechanisms or processes, such as incremental semantic interpretation. I close the paper with some more technical remarks on this relationship.

As is well-known, one of the major attractions of resource logic accounts of natural language is that the syntax-semantics interface can take a particularly simple form, as the Curry-Howard correspondence pairs each syntactic operation with a corresponding semantic counterpart.

Strong incremental interpretation requires that every prefix of a sentence form a single semantic entity. Take linguistic objects to be pairs $X : \alpha$ of a type X and a λ -term α of type X , and let Γ, Δ be a string of lexical type/ λ -term pairs with an interpretation β , i.e.,

$$\Gamma, \Delta \vdash S : \beta$$

Then an *incremental interpretation* of the prefix Γ is an *interpolant* $X : \alpha$ that satisfies the pair of constraints:

$$\Gamma \vdash X : \alpha \quad \text{and} \quad X : \alpha, \Delta \vdash S : \beta.$$

For many resource logics, including LCG, it is possible to prove that such incremental interpretations always exist, i.e., if Γ, Δ has an interpretation then Γ has an incremental interpretation.

A natural question to ask is if there is any relationship between the interpolants that function as incremental interpretations and the proof net cuts used in the accounts of center embedding processing complexity.

Define the complexity of a type as one plus the number of binary connectives appearing in it (thus $(NP \setminus S) / NP$ has complexity 3). Then the complexity of an interpolant is never less than the number of axiom links crossed by the corresponding cut, since the number of atomic links constructed by expanding the interpolant using proof net connectives is given by the interpolant's complexity.

On the other hand, the minimum complexity of an interpolant may be greater than the number of atomic links crossed by a minimal cut. For example, consider the sequence of LCG types corresponding to a transitive clause in a SOV language.

$$NP : a, NP : b, \triangle NP \setminus (NP \setminus S) : r \quad (8)$$

The minimal cut at ' \triangle ' in a proof net for (8) crosses two NP axiom links. However, the smallest LCG interpolant at ' \triangle ' is

$$S / (NP \setminus (NP \setminus S)) : \lambda f.f(b)(a),$$

which has complexity 4.

Interestingly, this difference disappears if we enrich LCG's type with product ' \otimes ', with pair formation ' $\langle \cdot, \cdot \rangle$ ' as the corresponding semantic operation. With this extension the minimal interpolant is

$$NP \otimes NP : \langle a, b \rangle$$

which corresponds directly to the axioms crossed by the minimal cut.

References

- Chomsky, Noam. 1965. *Aspects of the Theory of Syntax*. The MIT Press, Cambridge, Massachusetts.
- Chomsky, Noam and George Miller. 1963. Finitary models of language users. In R. Luce, editor, *Handbook of Mathematical Psychology*, volume 2. John Wiley and Sons, New York.
- Gibson, Edward and James Thomas. 1996. The processing complexity of English center-embedded and self-embedded structures. In *The Proceedings of the NELS*. Graduate Student Association, University of Massachusetts, Amherst.
- Girard, Jean-Yves. 1995. Linear Logic: Its syntax and semantics. In Jean-Yves Girard, Yves Lafont, and Laurent Regnier, editors, *Advances in Linear Logic*. Cambridge University Press, Cambridge, England, pages 1–42.
- Girard, Jean-Yves, Yves Lafont, and Paul Taylor. 1989. *Proofs and Types*, volume 7 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge University Press, Cambridge, England.

- Lewis, R. 1996. A theory of grammatical but unacceptable embeddings. *Journal of Psycholinguistic Review*.
- Moortgat, Michael. 1997. Categorical type logics. In Johan van Benthem and Alice ter Meulen, editors, *Handbook of Logic and Language*. The MIT Press, Cambridge, Massachusetts, pages 93–178.
- Morrill, Glyn V. 1994. *Type-logical Grammar: Categorical Logic of Signs*. Kluwer Academic Publishers, Dordrecht.
- Roorda, Dirk. 1991. *Resource Logics: Proof-theoretical Investigations*. Ph.D. thesis, University of Amsterdam, September.
- Roorda, Dirk. 1992. Proof nets for Lambek calculus. *Journal of Logic and Computation*, 2:211–231.
- Stabler, Edward P. 1994. The finite connectivity of linguistic structure. In Chuck Clifton, Lyn Frazier, and Keith Rayner, editors, *Perspectives in Sentence Processing*. Lawrence Erlbaum, pages 303–336.
- van Benthem, Johan. 1995. *Language in Action: Categories, Lambdas and Dynamic Logic*. The MIT Press, Cambridge, Massachusetts.

Address for correspondence:

Department of Cognitive and Linguistic Sciences, Box 1978
Brown University
Providence, RI 02912, USA
Email: mj@lx.cog.brown.edu