

Binding Machines

António Branco*
University of Lisbon

Binding constraints form one of the most robust modules of grammatical knowledge. Despite their crosslinguistic generality and practical relevance for anaphor resolution, they have resisted full integration into grammar processing. The ultimate reason for this is to be found in the original exhaustive coindexation rationale for their specification and verification. As an alternative, we propose an approach which, while permitting a unification-based specification of binding constraints, allows for a verification methodology that helps to overcome previous drawbacks. This alternative approach is based on the rationale that anaphoric nominals can be viewed as binding machines.

1. Introduction

Since the so-called integrative approach to anaphor resolution was developed in the late 1980s (Carbonell and Brown 1988; Rich and LuperFoy 1988; Asher and Wada 1989), and its practical viability extensively tested (e.g., Lappin and Leass 1994; Mitkov 1997, 1998), it has been common wisdom that factors determining the antecedents of anaphors divide into filters and preferences. The former exclude impossible antecedents and help to circumscribe the set of antecedent candidates; the latter help to pick the most likely candidate, which will be proposed as the antecedent.

Binding constraints are a significant subset of such filters. As they delimit the relative positioning of anaphors and their possible antecedents in grammatical geometry, these constraints are crucial to restricting the search space for antecedents and enhancing the performance of anaphor resolvers.¹ From an empirical perspective, they stem from quite robust generalizations and exhibit a universal character, given their parameterized validity across natural languages. From a conceptual point of view, in turn, the relations among binding constraints involve nontrivial symmetry, which lends them a modular nature. Accordingly, they have been considered one of the most robust and intriguing grammar submodules, usually referred to as binding theory. However, in contrast to this, the formal and computational handling of binding constraints has presented considerable resistance.

Anaphor resolution typically builds on many sources of information—among them, information about the grammatical structure of the sentence—so that the different filters and preferences may be used. Consequently, it must in general be regarded as a postgrammatical process, in the sense that it is completed after sentences are parsed. Binding constraints, as a subset of the filters for anaphor resolution, are a special case

* Department of Informatics, Faculdade de Ciências de Lisboa, Campo Grande, 1700 Lisboa, Portugal.
E-mail: Antonio.Branco@di.fc.ul.pt.

1 See the Appendix for a specification of binding constraints. We adhere to the following terminological convention: anaphors divide into reflexives and nonreflexives; reflexives form a class that includes short-distance (ruled by Constraint A; e.g., *himself*) and long-distance reflexives (Constraint Z; e.g., Chinese *ziji*); nonreflexives include pronouns (Constraint B; e.g., *he*) and nonpronouns (Constraint C; e.g., *the student*).

in this respect. Given that they form a submodule of grammar, they are specified on a par with other grammatical submodules and constraints, and they are thus expected to be integrated already into the processing of grammar. Nevertheless, this integration cannot be considered to have been adequately achieved.

As we will discuss at length, the original methodology for verifying the compliance of grammatical representations with binding constraints requires extragrammatical processing steps delivering a forest of indexed trees to anaphor resolvers (Chomsky 1981). More recently, constraint-based grammatical frameworks either require special-purpose extensions of the description formalism, though ensuring only a partial handling of these constraints, as in Lexical-Functional Grammar (LFG; Dalrymple 1993), or do not offer a solution yet to integrate them into grammar, as in Head-Driven Phrase Structure Grammar (HPSG; Pollard and Sag 1994).²

Our primary goal here is thus to bridge the gap between the grammatical nature of binding constraints and their full integration into grammar processing. In particular, we aim at achieving this in such a way that a lean interface between grammar and reference processing emerges.

In Section 2, we first underline the distinction, seldom taken into account, between specification and verification of binding constraints. We then review advances proposed in the literature concerning the completion of the verification task. We observe that three major lines of progress can be identified: packing of anaphoric ambiguity, packing of nonlocal context, and lexicalization of binding constraints.

Building on these contributions, in Section 3 we argue that the remaining step forward is to harmonize these different advances. We suggest that a more accurate, semantics-driven comprehension of the nature of binding constraints is a relevant move toward this harmonization. On the basis of this revision, we introduce a methodology for verifying these constraints, which rests on the new concept of binding machine, to be defined.

In Section 4, in the light of this new methodology, we show how binding constraints can be given a unification-based specification and can be fully integrated into grammar.

In Section 5, we present an illustrative example and discuss in detail how binding constraints and reference-processing systems are coordinated, and how the previously identified drawbacks are overcome.

2. Advances in the Verification Task

In recent decades, great strides have been made toward an empirically adequate specification of binding constraints, this being an important research issue in theoretical linguistics. Many aspects of this issue—a parameterizable definition of local domain, the existence of a fourth constraint for long-distance reflexives, the possible subject-orientedness of some anaphors, and the degree of universality of binding constraints, to name just a few—have come under intense scrutiny.

In contrast, the verification task has been studied much less extensively. Even though important problems also remain to be solved in this more applied dimension

² The fragment of grammar developed and extensively discussed in Pollard and Sag (1994) is formally specified in its Appendix with the HPSG unification-based description language. Binding constraints escape such encoding. While noting that these constraints have yet to be accommodated in HPSG grammars, Breckenkamp's (1996) and Backofen et al.'s (1996) subsequent elaboration of this issue implies that some kind of essential limitation of the unification-based formalism might have been reached, a suggestion we seek to contradict here.

of the so-called binding theory, the issue of determining whether a given grammatical representation complies with binding constraints has not attracted similar attention.

In this section, we briefly review major advances reported in resolving this issue.

2.1 Exhaustive Coindexing for Filtering

The first formulation of a verification procedure, based on exhaustive coindexation, dates back to Chomsky (1980, Appendix; 1981, Section 3.2.3). The basics of this approach can be outlined as follows:

After the grammatical parsing of a sentence with n NPs has been completed, for every parse tree t :

- a. *Indexation*: Generate a new, annotated tree by assigning indices to the NPs in t .
- b. *Filtering*: Store this annotated tree if the indexation of NPs respects binding constraints; otherwise, delete it.
- c. *Iteration*: Repeat (a)–(b) until all type-different assignments of n possibly different indices have been exhausted.

As discussed in Correa (1988), this procedure is grossly inefficient: its complexity was shown in Fong (1990) to be of exponential order. Moreover, this approach is conceptually awkward, given that a submodule of the grammar, the set of binding constraints, is not operative during grammatical processing, but functions as an extragrammatical add-on.³

This proposal also disregards the need to interface grammar with systems for reference processing. The input for such systems will not be a grammatical representation to be refined vis-à-vis the preferences for anaphor resolution, but a forest of differently labeled trees that have to be internally searched and compared with each other by anaphor resolvers.

2.2 Packing Anaphoric Ambiguity

A first proposal for improving the exhaustive coindexation-driven methodology is due to Correa (1988), whose goal was to enhance the integration of binding constraints into grammar and obtain a tractable verification procedure.

Simplifying some details, the proposed algorithm can be outlined as follows:

Let t be a constituency tree where every NP has a type-distinct index. Start from the top node of t with two empty stacks, A and B , where indices will be collected, respectively local c-commanding⁴ indices and nonlocal c-commanding indices, while descending the tree. When an NP _{j} is found:

- a. *Copy*: Leave a copy of A (if NP _{j} is a short-distance reflexive) or B (if it is a pronoun) at the NP _{j} .

³ Correa (1988, page 123) observes that although the integration of binding constraints “into rules which may be used to derive structure that already satisfies the [constraints] is not a straightforward task,” that should be the path to follow, a point also strongly stressed in subsequent elaboration on this issue by Merlo (1993).

⁴ C-command is a configurational version of the command relation where x c-commands y iff the first branching node that dominates x dominates y (Barker and Pullum 1990).

- b. *Assign*: Take the first index i of the stack copied into the NP_j node, and annotate NP_j with $j = i$.
- c. *Collect*: Add index j to A in each sister node of NP_j .

When a local domain border is crossed:

- d. *Reset*: Reset B to $A \cup B$.

This algorithm has been given two different implementations, one by Correa (1988), the other by Ingria and Stallard (1989). Further elaboration by Giorgi, Pianesi, and Satta (1990) and Pianesi (1991) offers a variant in terms of formal language techniques, where the stack copied into pronouns contains the antecedent candidates excluded by Principle B.

The “do-it-while-parsing” approach of Correa’s implementation has the advantage of discarding a special-purpose postgrammatical module for binding. Nevertheless, this solution turns out to be dependent on a top-down parsing strategy. On the other hand, while Ingria and Stallard’s implementation is independent of the parsing strategy adopted, its independence comes at the cost of still requiring a special-purpose postgrammatical parsing module for binding.

Besides incorporating binding theory into grammar, Correa’s development inside the coindexation-driven methodology presents other significant improvements. If one disregards step (b)—a disguised recency preference mixed with binding constraints—and considers the result of verifying these constraints to be the assignment to an NP of the set of indices of its grammatically admissible antecedents, then it is possible to discard the proliferation of indexed trees as a way to express anaphoric ambiguity. Moreover, this packing of anaphoric ambiguity provides for a neat interface with anaphor resolvers, whose preferences will then pick the most likely antecedent candidate from the relevant stack of indices.

These advances permit a verification procedure of tractable complexity (Correa 1988, page 127; Giorgi, Pianesi, and Satta 1990, page 5). This results crucially from the move toward the lexicalization of the constraining effect of binding principles, a solution also adopted in subsequent proposals by other authors, as we will discuss below. The binding constraint of each anaphor is now enforced independently of how the surrounding anaphors happen to be resolved. This implies that there is no need to anticipate all the different resolutions for every relevant anaphor with a process of exhaustive coindexation. It also implies that cases of undesired transitive anaphoricity are handled by other filters during the anaphor resolution process.⁵

However, these positive results regarding the verification task seem to be obtained at the cost of some negative consequences regarding the specification task and empirical adequacy. The above algorithm is acknowledged not to be able to cope with

⁵ Consider the sentence *John said that he shaved him*. Ignoring how other anaphors are resolved, in the light of Binding Constraint B, *he* can take *John* as its antecedent, as empirically replicated in other minimally different examples such as *John_i said that he shaved Peter_i*; likewise, *him* can take *John* as its antecedent. A point worth noting is that, if *he* actually ends up resolved against *John*, the latter cannot be the antecedent of *him*, and vice versa. This specific resolution of *he* and *him*, out of the many possible resolutions, blocks two anaphoric links that would otherwise have been admissible. It induces a contingent violation of binding constraint B due to an accidental, transitive anaphoric relationship between *he* and *him*.

This issue is not discussed in Correa (1988), since this paper is strictly focused on syntax and binding. See footnote 13 below for a suggestion on how this issue may be handled in a grammatical framework integrating syntactic and semantic representations.

constraints involving nonlocal dependencies. It does not account for Principle C, and it only partially accommodates the anaphoric potential of anaphors complying with Principle B. As Stack B only contains indices of the nonlocal c-commanders—rather than all indices except those of the local c-commanders—the algorithm does not correctly account for the constraining effect of Principle B. Also this approach does not account for backward anaphora or crossover cases (Correa 1988, page 127; Ingria and Stallard 1989, page 268).⁶

2.3 Packing Nonlocality

Other improvements in the task of verifying binding constraints are due to Dalrymple (1993) and Johnson (1995). Instead of being concerned with packing ambiguity, they are concerned with packing nonlocality.

2.3.1 Trees in Nodes of Trees. Johnson's (1995) algorithm is embodied in Prolog code. Abstracting away from details associated with that format, it can be outlined as follows:

Let t be a constituency tree where every NP has a type-distinct index. For every NP_i in t , traverse the tree from NP_i upward until the top node is reached. When a locally c-commanding NP_j is found:

- a. Annotate NP_i with $i = j$ if NP_i is a short-distance reflexive.
- b. Annotate NP_i with $i \neq j$ if NP_i is a nonreflexive.

When a nonlocally c-commanding NP_j is found:

- c. Annotate NP_i with $i \neq j$ if NP_i is a nonpronoun.

Although this outline renders the algorithm in a bottom-up fashion, Johnson ingeniously develops an implementation of it that is independent of the parsing strategy by resorting to delaying mechanisms. Consequently, despite its postgrammatical flavor, this implementation does not require postgrammatical processing, thus incorporating the task of binding constraint verification into grammar processing.

These results are obtained with some auxiliary devices. Each node in the tree is "conceptualized as a pair consisting of a tree and a vertex in that tree" (Johnson 1995, page 62). Consequently, the whole tree where a given NP appears is locally accessible to be "walked up" since its replica is present at the pair (Category, Tree), which is the NP node itself.

This algorithm makes the verification of binding constraints more efficient because it does not resort to exhaustive indexation. However, it does so at the cost of highly complicating the grammatical representation, since the tree is replicated at each one of its nodes.

While avoiding exhaustive indexation, this approach does not fully eliminate the proliferation of trees. For a given ambiguous reflexive, with more than one admissible

⁶ See the Appendix for the notion of locality and local domain and other auxiliary notions in the definition of binding constraints.

Backward anaphora occurs in cases where the anaphor is resolved against an antecedent that occurs linearly after the anaphor, as in *If he_i is around, Peter_i will do it.*

An example of so-called crossover cases is the ungrammatical construction **Who_i did Peter think she_i saw?* or **Peter_i, he_i said you like*, where the fronted phrase is meant to be the antecedent of some pronoun c-commanding the position from which this phrase is displaced.

antecedent, each antecedent candidate corresponds to a different coindexation and, consequently, to a different tree. That is what generally happens with long-distance reflexives, whose antecedents can be found in any of the binding domains induced by the local or by the upward predicators, but it may also happen with short-distance ones, as in (2) below.

As to the interface with reference processing, problems arise with reflexives and nonreflexives, though of different nature. Reflexives, if ambiguous, give rise to proliferation of trees, thus requiring comparison between trees during subsequent anaphor resolution.

As to nonreflexives—pronouns and nonpronouns—their analysis does not give rise to proliferation of trees, but the representation of their ambiguity is not fully made explicit in the grammatical representation of the sentence being parsed. This is so because they end up associated with negative information, that is, information about what NPs cannot be their antecedents. The index of a pronoun is made unequal with the indices of its local c-commanders; it is not made equal with the indices of its grammatically admissible antecedents. The same holds for nonpronouns with respect to their c-commanders. Consequently, in this case, the task of determining the antecedent candidates that satisfy the relevant binding constraint of nonreflexives remains to be completed after grammatical processing is finished. This will involve some postgrammatical rescanning of the parse tree generated for extracting the indices that do not enter in the inequalities obtained during the parsing.

Finally, like Correa's approach, Johnson's does not account for backward anaphora, as only surface c-commanders are visible to the tree-climbing procedure.

2.3.2 Equations with Regular Expressions. The basic LFG account of binding, set forth by Dalrymple (1993), adopts a different approach to generalize over the possible nonlocality of intrasentential anaphoric dependencies. This approach makes crucial use of a special-purpose extension of the LFG description formalism, the so-called binding equations, which are lexically associated with anaphors. Building on Kaplan and Maxwell's (1988) proposal concerning functional uncertainty, binding equations are designed to encode the uncertainty concerning the long-distance path between the positions of the anaphor and its permissible antecedent in the grammatical structure.⁷ Given that uncertainty concerning long-distance dependencies involves a (possibly infinite) disjunction of possibilities, the basic idea is to encode such a disjunction in finite terms by the use of regular expressions over feature structures. An example of a binding equation encoding functional uncertainty is given in (1), preceded by an example with the corresponding long-distance subject-oriented reflexive, Chinese *ziji*.

- (1) Zhangsan_i yiwei [Lisi_j yiwei [... ziji_{i/j}/...]].
 Zhangsan_i thought [Lisi_j thought [... him_{i/j}/...]].
 $ziji: ((COMP^* OBJ \uparrow) SUBJ)_\sigma = \uparrow \sigma$

The right-hand side of the equation stands for the semantic representation (" σ ") of the anaphor (" \uparrow "), while the left-hand side stands for the semantic representation of the antecedent. The description of the antecedent indicates that the long-distance reflexive is an object and that this object is constrained to be part of a feature structure where

⁷ Koenig (1999) introduces a device in HPSG description language for stating inside-out constraints. This would help in developing an HPSG emulation of the LFG approach for the verification of binding constraints.

its antecedent may be one of the possibly many upward subjects. The Kleene operator “*” allows abbreviation of the set of paths consisting of zero or more occurrences of COMP—corresponding to possible successive clausal embeddings—followed by one occurrence of OBJ.

While regular expressions may be used in binding equations, such expressions are not necessary if the grammatical relation between the anaphor and its admissible antecedents does not involve a long-distance dependency. That is the case in (2), which displays the binding equation for the short-distance reflexive *himself*. Given that both the subject and the object are admissible antecedents for the reflexive, in the binding equation the use of the attribute GF, which stands for any grammatical function, underspecifies the grammatical functions of the admissible antecedents (Dalrymple 1993, Section 4.4.2).

- (2) John_i described Bill_j to himself_{i/j}.
himself: ((OBL_{Goal} ↑) GF)_σ = ↑ σ

Binding equations may also express negative constraints, as in (3), where the semantic representation of the pronoun is constrained to be different from that of its local coarguments.

- (3) *John_i described Bill_j to him_{i/j}.
him: ((OBL_{Goal} ↑) GF)_σ ≠ ↑ σ

As noted in Dalrymple (1993, Section 3.3), a few aspects of this approach for binding need to be fully worked out. For instance, the positive equations for reflexives do not require identity of indices of anaphorically related expressions, but instead impose identity of semantic representations. Without further elaboration, this will incorrectly enforce any type of anaphoric link (coreference, bound, bridging, e-type, etc.) to the sole mode of coreference. Another important issue is the account of nonlexical anaphoric NPs: it is not clear how this type of NP (e.g., anaphoric definite descriptions, ruled by Principle C) may be assigned the corresponding binding equation.

However these difficulties turn out to be resolved, the LFG approach for binding, though building on a different strategy for handling nonlocality, presents the same sort of problems as Johnson’s proposal.

The interfacing of grammar with reference-processing systems is problematic since the proliferation of representations is not avoided. Constructions with reflexives, if these are ambiguous, end up associated with several grammatical representations. In the case of long-distance reflexives, as exemplified in (1), these representations result from the possibly many solutions for the functional uncertainty encoded by the regular expression in the binding equation. In the case of short-distance reflexives, as exemplified in (2), they result from the different solutions for the unification of the different grammatical functions of the admissible antecedents with the attribute GF in the binding equation.

Likewise, the anaphoric capacity of pronouns and nonpronouns, typically ambiguous, is not explicitly captured in the final grammatical representation. These anaphors are lexically associated with negative equations, and for this type of equation there is only one possible solution, namely, the grammatical structure where the semantic representation of the anaphor is not identical to the semantic representations of any of the phrases complying with the description of the antecedent in the left-hand side

of the equation (Dalrymple 1993, Section 4.1.5). Therefore, for these anaphors the final grammatical representation provides no information about what their admissible antecedents are according to the relevant binding constraints.

3. A Semantics-Driven Approach

The contributions assessed above share a common point of departure with regard to the verification algorithm first proposed by Chomsky (1981), each addressing and solving some of its more significant drawbacks. The common move toward the lexicalization of binding constraints represents an important shift in the verification strategy: verifying binding constraints is not a matter of inspecting final grammatical representations, but instead a matter of some local operation triggered by information lexically associated with anaphors about their anaphoric class. This move has allowed binding constraint verification to be incorporated into grammar processing and permitted tractable verification procedures.

From the discussion in the previous section, it follows also that these contributions have been partially successful in overcoming other problems of the verification methodology based on exhaustive coindexation. Though partially successful, they have brought to the fore important dimensions of binding that have to be concomitantly accounted for. Accordingly, an alternative method for the verification of binding constraints has to find a way to harmonize all those different dimensions—lexicalization, anaphoric ambiguity packing, and nonlocal context packing—while providing adequate empirical coverage and neatly interfacing grammar with reference processing.

Against this background, a breakthrough depends, in our view, on reconsidering some primitives underlying the conception of binding constraints. In the previous section, we made a clear distinction between specification and verification of binding constraints, so that the latter task could be isolated and better assessed. We will argue now that further progress on the verification task depends on bridging this distinction and possibly changing the way the specification of binding constraints is understood.

3.1 Patterns in the Semantics of Anaphors

Binding constraints have generally been viewed as well-formedness conditions on syntactic representations, thus belonging to the realm of syntax. In line with Gawron and Peters (1990), however, we think these constraints should rather be understood as conditions on semantic representations, since they primarily delimit (nonlocal) aspects of semantic composition, rather than aspects of syntactic composition.⁸

Like other types of constraints on semantic composition, binding constraints impose conditions on the interpretation of certain expressions—anaphors, in the present case—based on syntactic geometry. However, this cannot be viewed as implying that they express grammaticality requirements. By replacing a pronoun with a reflexive in a given sentence, for instance, we do not turn a grammatical construction into an ungrammatical one, even if we assign to the reflexive the antecedent appropriately se-

⁸ As implied by the title of this section, and as will become clear in the following discussion, this does not mean that we are claiming that binding theory can be built without any reference to syntactic constructs.

In the argument in the following paragraphs, we are assuming a notion of semantic composition not in its strict sense, as used for example in Montague Grammar, but in the broader sense that the intermediate semantic representations of the expressions are composed from other representations, as used in Discourse Representation Theory (DRT). Note that reformulations of frameworks like DRT can be worked out that result in a semantic system adhering to strict compositionality; see Janssen (1997, Section 4.4) for references and a thorough discussion of this issue.

lected for the pronoun. In that case, we are simply asking the hearer to try to assign to that sentence a meaning it cannot express—just as if we were to ask whether someone could interpret *The red book is on the white table* as describing a situation where a white book is on a red table.

In this example, given how they happen to be syntactically related, the semantic values of *red* and *table* cannot be composed in such a way that the sentence could be used to describe a situation concerning a red table, rather than a white table. Likewise, in the sentence *John thinks Peter shaved him*, given how they happen to be syntactically related, the semantic values of *Peter* and *him* cannot be composed in such a way that this sentence could be used to describe a situation where John thinks that Peter shaved himself (i.e., Peter), rather than a situation where John thinks that Peter shaved other people (e.g., Paul, Bill, or John himself). The difference between these two cases is that in the former, the composition of the semantic contributions of *white* and *table* (for the interpretation of the NP *white table*) is constrained by local syntactic geometry, while in the latter, the composition of the semantic contributions of *John* and *him* (for the interpretation of the NP *him*) is constrained by nonlocal syntactic geometry.

This discussion leads us to consider that, semantically, an anaphor should be specified in the lexicon as a function whose argument is a suitable representation of the context—providing a semantic representation of the NPs available in the discourse vicinity—and its value is the set of the grammatically admissible antecedents for that anaphor. This rationale is in line with other approaches to the meaning of anaphors that, building in other sorts of arguments or research concerns, understand it also as a projection from some relevant representation of contexts to entities.⁹ But given the specific focus of the present study, what should be noted is that, all in all, there will be four such functions available to be lexically associated with anaphors, each corresponding to one of the four different classes of anaphors, in accordance with the four binding constraints A, B, C, and Z.¹⁰

3.2 Binding Machines

Given these considerations, we can show that this conceptual shift to a semantics-driven approach for the verification of binding constraints provides an adequate basis for harmonizing the advances put forward in the literature and discussed above.

To make this alternative rationale for binding perspicuous, we suggest envisioning an anaphoric NP as a *binding machine*, which operates by receiving an input, changing its internal state, and returning an output. More specifically, an anaphoric NP can be

⁹ See, among others, Gawron and Peters (1990), Lappin and Francez (1994), and the discussion in Jacobson (1999).

Adopting Löbner's (1987) duality criterion for quantification in natural language, and the formal tools he developed for the analysis of phase quantification, we showed in Branco (2000) that the four binding constraints can be seen as the effect of four binding quantifiers. These phase quantifiers can be viewed as being expressed by the nominals of the four binding classes, and they quantify over the reference markers organized in the obliqueness order.

A full-fledged account of the empirical support and justification for these results, and of their implications, is beyond the scope of this article. For an abridged presentation of the core argument, see Branco (1998).

¹⁰ As there are different grammatical frameworks, binding constraints have been specified under different versions. Some differences between versions are due just to this fact that binding constraints are supposed to be accommodated into different grammatical frameworks; some other differences, however, are real differences of specification in the sense that different variants may not have the same empirical coverage or be aimed at predicting the same (un)grammatical constructions. In the Appendix, we present a common and fairly well empirically tested version of binding theory given the current state of the art in this area, a version presently adopted in the HPSG framework. For an alternative, see for example Reinhart and Reuland (1993).

viewed as a binding machine that (1) takes a representation of its context; (2) updates its own semantic value in response both to its context and to its intrinsic anaphoric potential (i.e., in accordance with its binding constraint); and (3) contributes to the makeup of the context, which the other binding machines read as input (i.e., against which the other anaphoric NPs are interpreted).¹¹

The *output* of an anaphoric nominal n viewed as a binding machine is simply the incrementing of the context with a copy of its reference marker.¹²

The *internal state* of the machine after its operation is a representation of the contextualized anaphoric capacity of n under the form of the set of reference markers of the grammatically admissible antecedents of n . This internal state results when the binding constraint associated with n is applied to the input, and it is the interface point between grammar and reference processing. This set of reference markers collects the antecedent candidates, and its elements are submitted to other filters and preferences by the anaphor resolvers so that one of them ends up being chosen as the antecedent.

The *input*, in turn, is a representation of the aspects of the context relevant to help circumscribe the anaphoric potential of nominal anaphors. It is coded under the form of three lists of reference markers, \mathbf{A} , \mathbf{Z} , and \mathbf{U} . In list \mathbf{A} , the reference markers of the local o-commanders of n are ordered according to their relative grammatical obliqueness; \mathbf{Z} includes the o-commanders of n , possibly observing a multiclausal obliqueness hierarchy; and \mathbf{U} is the list of all reference markers in the discourse context, including those not linguistically introduced.

Given this setup, the contribution of binding constraints in circumscribing the anaphoric potential of nominals is explicitly acknowledged. The particular contextualized instantiation of that potential and the verification of binding constraints coincide and consist of a few simple steps. If n is a short-distance reflexive, its internal state is set up as \mathbf{A}' , where \mathbf{A}' contains the reference markers of the o-commanders of n in \mathbf{A} . If n is a long-distance reflexive, its semantic representation includes \mathbf{Z}' , such that \mathbf{Z}' contains the o-commanders of n in \mathbf{Z} . If n is a pronoun, $\mathbf{B} = \mathbf{U} \setminus (\mathbf{A}' \cup [\text{r-mark}_n])$ is encoded into its representation, where r-mark_n is the reference marker of n . Finally, if n is a nonpronoun, its updated semantics keeps a copy of $\mathbf{C} = \mathbf{U} \setminus (\mathbf{Z}' \cup [\text{r-mark}_n])$.

Besides adhering to an empirically grounded conception of binding constraints, this approach embodies, and harmonizes, the crucial contributions of previous proposals concerning the verification of these constraints. It assumes the lexicalization of binding constraints. Concomitantly, it builds on specific strategies for the packing of anaphoric ambiguity (viz., list of reference markers) and nonlocal context (viz., set of lists of reference markers). Moreover, it achieves this while avoiding the above-mentioned problems related to the proliferation of grammatical representations and to the interfacing of grammar with reference processing, as well as the problems of ensuring complete empirical coverage.

What remains to be discussed is whether, given this new format for the verification of binding constraints, they can still be specified and integrated into grammar processing with currently affordable formal and computational tools.

4. A Unification-Based Specification Exercise

This new approach to binding constraints can be integrated into grammar easily and in a principled manner. In what follows, we outline how these constraints can be specified and handled in a unification-based grammatical framework such as HPSG.

¹¹ This rationale is in line with the insights of Johnson and Klein (1990) concerning the processing of the semantics of nominals.

¹² See Kamp and Reyle (1993) for the notion of reference marker.

As a proposal for that integration, we designed an extension to the Underspecified Discourse Representation Theory (UDRT) semantics component for HPSG developed by Frank and Reyle (1995). This component is encoded as the value of the feature $\text{CONT}(\text{ENT})$, which is now extended with the feature $\text{ANAPH}(\text{ORA})$; see (4). This new feature keeps information about the anaphoric potential of the corresponding nominal n : its subfeature $\text{ANTEC}(\text{EDENTS})$ keeps a record of how that potential is updated when the anaphor enters a grammatical construction; and its subfeature $\text{R}(\text{EFERENCE})\text{-MARK}(\text{ER})$ indicates the reference marker of n , to be contributed to the context.

Similarly, and still assuming Pollard and Sag's (1994) feature geometry as a starting point, the NONLOC value is also extended with a new feature, $\text{BIND}(\text{ING})$, with subfeatures LIST-A , LIST-Z , and LIST-U . These lists provide a specification of the relevant context and correspond to the lists **A**, **Z**, and **U** above. Subfeature LIST-LU is a fourth, auxiliary list for encoding the contribution of local context to the global, nonlocal context.

The SYNSEM value of a pronoun, for instance, can now be designed as shown in (4).

$$(4) \left[\begin{array}{l} \text{LOC} \mid \text{CONT} \\ \text{NONLOC} \mid \text{BIND} \end{array} \left[\begin{array}{l} \text{LS} \left[\begin{array}{l} \text{L-MAX} \quad [1] \\ \text{L-MIN} \quad [1] \end{array} \right] \\ \text{SUBORD} \quad \{ \} \\ \text{CONDS} \quad \left\{ \begin{array}{l} \text{LABEL} \quad [1] \\ \text{ARG-R} \quad [2] \end{array} \right\} \\ \text{ANAPH} \quad \left[\begin{array}{l} \text{R-MARK} \quad [2] \\ \text{ANTEC} \quad [5] \textit{principleB}([4],[3],[2]) \end{array} \right] \\ \text{LIST-A} \quad [3] \\ \text{LIST-Z} \quad \textit{list} \\ \text{LIST-U} \quad [4] \\ \text{LIST-LU} \quad [2] \end{array} \right] \right]$$

Given this feature structure, the binding constraint associated with pronouns is specified as the relational constraint *principleB*. This relational constraint returns list **B** as the value of ANTEC . It is defined to take (in the first argument) all markers in the discourse context, given in LIST-U value, and remove from them both the local o-commanders of the pronoun (included in the second argument) and the marker corresponding to the pronoun (in the third argument).

The SYNSEMS of other anaphors, ruled by Principles A, C, and Z, are similar to the one above.¹³ The only difference lies in the relational constraint in the ANTEC value, which encodes the appropriate binding constraint and returns the updated anaphoric potential under the form of list **A'**, **C**, or **Z'**, respectively, as discussed in the previous section.

Turning to the specification of the context (i.e., the values of LIST-A , LIST-Z , LIST-U , and LIST-LU), this is handled by means of a new HPSG principle, which can be termed the Binding Domains Principle. This principle consists of three clauses constraining

¹³ Binding constraints for nonlexical anaphoric nominals are lexically stated in the corresponding determiners.

A constraint for pronominal anaphoric transitivity may also be introduced at the lexical representation of pronouns, by including in the CONDS value in (4) Discourse Representation Structure conditions expressing that

$$\forall r_a, r_b (([2] = \textit{anaph}r_a \wedge r_b = \textit{anaph}r_a) \Rightarrow ([r_b] \cup [5] = [5])).$$

signs and their values with respect to these lists of reference markers. Due to space limitations, we illustrate this principle simply by stating Clause I, which constrains LIST-U and LIST-LU.¹⁴

- (5) Binding Domains Principle, Clause I
- a. In every sign, the LIST-LU value is identical to the concatenation of the LIST-LU values of its daughters.
 - b. In a sign of sort *discourse*, the LIST-LU and LIST-U values are token identical.
 - c. In a non-NP sign, the LIST-U value is token identical to each LIST-U value of its daughters.
 - d. In an NP sign *k*,
 - i. In Spec-daughter, the LIST-U value is the result of removing the elements of the LIST-A value of Head-daughter from the LIST-U value of *k*;
 - ii. In Head-daughter, the LIST-U value is the result of removing the value of R-MARK of Spec-daughter from the LIST-U value of *k*.

LIST-LU collects, up to the outermost sign of sort *discourse*, all the markers contributed by the different NPs for the context. At this sign, they are passed to LIST-U, by means of which they are propagated to every NP. The HPSG ontology was extended with the sort *discourse*, which corresponds to sequences of sentential signs and at whose signs reference markers from the nonlinguistic context may be introduced in the semantic representation.¹⁵ Subclause (d) is meant to avoid what is known in the literature as the i-within-i effect.

5. Example and Discussion

The above unification-based specification of binding constraints, while ensuring their integration into grammar, allows the binding module to be suitably hooked up with systems of reference processing. Feature ANTEC is the interface point between them.

¹⁴ Clauses II and III constrain LIST-A and LIST-Z, respectively. Roughly, Clause II ensures that the LIST-A value is passed from the lexical head to its successive projections, and also from the head-daughters to their arguments. Note that exemption occurs when $principleA(\square, \square)$ is the empty list, in which case the reflexive should find its antecedent outside any binding constraint (Pollard and Sag 1994, Chapter 6).

Clause III ensures that, at the top node of the grammatical representation, LIST-Z is set up as the LIST-A value of that node, and that LIST-Z is successively incremented at the suitable downstairs nodes by appending its value with the LIST-A value of those nodes.

At the lexical entry of a predicator, LIST-A is defined as the concatenation of the R-MARK values of its subcategorized arguments specified in the ARG-S value.

For a detailed specification of the Binding Domains Principle, see Branco (2000).

¹⁵ Reference markers can be introduced linguistically, by the utterance of the corresponding expressions, or nonlinguistically, by means of their cognitive availability in the context of the discourse. Theories of natural language semantics can be used to represent these two types of reference markers. Nevertheless, only a global theory encompassing natural language and cognition seems to be able to pursue the ambitious goal of providing an integrated account of how both types of markers, and not only those linguistically evoked, are introduced into semantic representation.

We are following a distinction between the notions of anaphor resolution and reference processing commonly assumed in the literature. Anaphor resolution is seen as being concerned with the task of identifying the antecedents of anaphors. It is therefore part of a reference-processing system, whose overall goal, in turn, is to determine the interpretation of the anaphors. This involves determining the appropriate semantic type of the anaphoric link between an anaphor and its antecedent (coreference, bridging, e-type, bound anaphora, etc.) and providing a suitable semantic representation for this link.

Being the interface point between grammatical representation and reference processing, the list value of the feature ANTEC has just to be reduced by anaphor resolvers, given the relevant preferences and filters other than binding constraints, until the most likely antecedent is isolated. It is thus a process concerning selection in a list, rather than search in a set of indexed trees.

As to reference processing in general, the specification suggested in the previous section provides a suitable framework for the correct representation of the semantically different types of anaphoric links, the range of options not being restricted to coreference only. After the anaphor has been resolved, the reference marker of the anaphor and the reference marker selected as the antecedent can be related in accordance with the mode of anaphora determined by the reference-processing system.

This semantic relation between anaphorically related reference markers can be represented simply as another DRS condition in the CONDS value. This makes possible a mainstream DRT representation for the resolved anaphoric link, thus building on the substantial number of already worked out solutions available in the literature for DRT-based semantic representation of anaphora.¹⁶

This specification of binding theory for HPSG was tested with a computational implementation using ProFIT (Erbach 1995). In this implementation, the relational constraints corresponding to binding principles were straightforwardly encoded by means of Prolog predicates associated to the lexical clauses of anaphoric expressions, and defined in terms of simple auxiliary predicates ensuring the component operations of list appending, list difference, and so on. It is worth noting that some of these predicates have arguments—for example, the LIST-U value, whose value is computed when the whole relevant grammatical representation is built up. This is a consequence of packing nonlocal information in such lists. As in Johnson's approach, it requires that some delaying device be used, which in this computational grammar was done by resorting to the Prolog built-in predicate `freeze/2`.

For the sake of the example, consider the following multiclausal sentence from Portuguese displaying backward anaphora between a topicalized reflexive and a pronoun:

- (6) De si próprio, cada estudante disse que ele gosta.
of him self every student said that he likes
'Himself, every student said that he likes.'

An abridged version of the grammatical representation produced by the implemented grammar for this sentence is presented in Figure 1, where the feature structures below the tree correspond to partial grammatical representations of the constituents

¹⁶ See Kamp and Reyle (1993) for a comprehensive rendering of DRT, and Branco (2000, Chapter 5) for an overview concerning the semantic representation of different modes of anaphora.

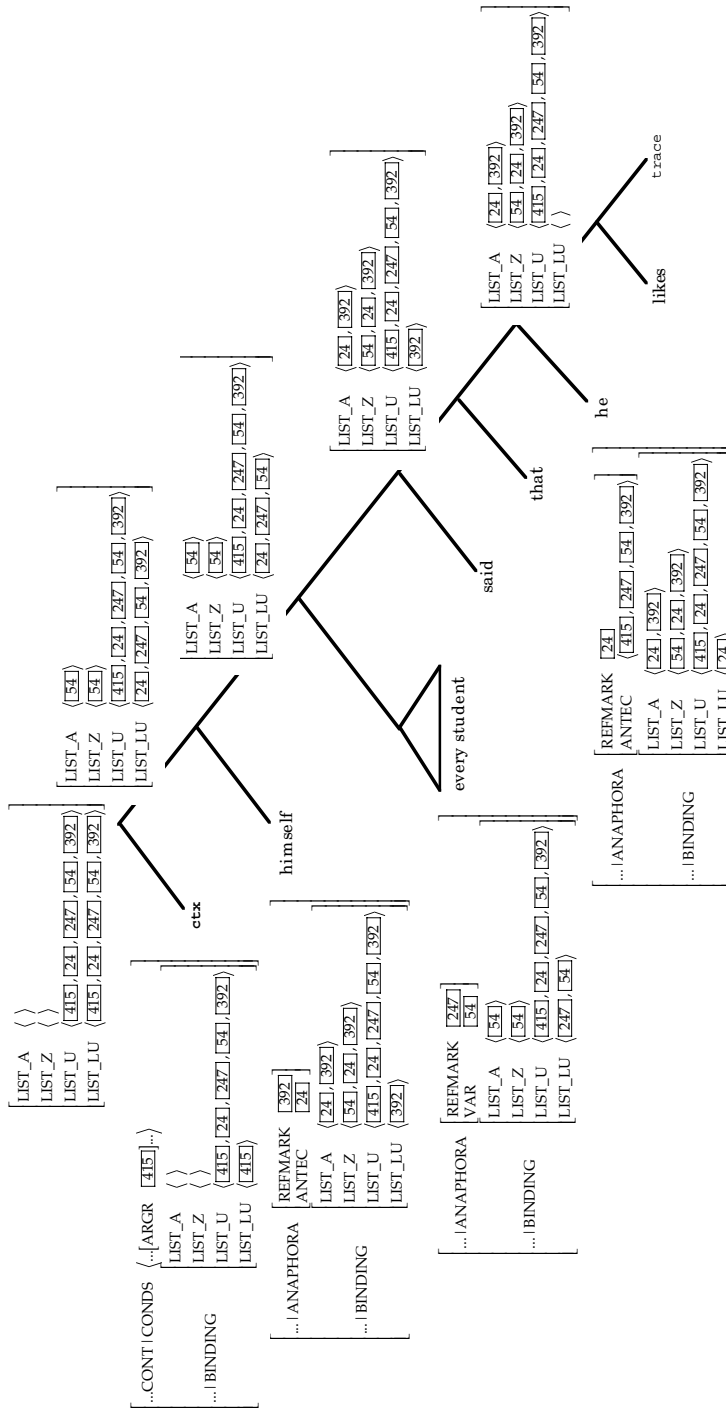


Figure 1
Abridged grammatical representation for the example sentence (9).

in the leaves of the tree, while the ones above the tree correspond to partial representations of some nonterminal nodes.

First, consider LIST-Z. In the outer nodes of the matrix clause, due to the effect of the Binding Domains Principle, Clause III, the LIST-Z value is obtained from the value of LIST-A, with which it is token identical, comprising the list with a single element $\langle 54 \rangle$. In the nodes of the embedded clause, the LIST-Z value is the concatenation of that upper LIST-Z value and the LIST-A value $\langle 24, 392 \rangle$ in the embedded clause, from which the list $\langle 54, 24, 392 \rangle$ is the result. LIST-A values are obtained from the representation of the subcategorization frames of the verbal predicates.

Next, consider LIST-LU. Reading upward, note that at each higher level in the constituency representation, the list gets longer; by the effect of the Binding Domains Principle, Clause I, the LIST-LU value at a given node gathers the reference markers of the nodes dominated by it. At the discourse top node, LIST-LU includes all the reference markers of the NPs in the example, the list $\langle 415, 24, 247, 54, 392 \rangle$. The Binding Domains Principle, Clause I, also ensures that this list of all reference markers is passed to the LIST-U value of the top node and that it is then percolated down to all relevant nodes of the grammatical representation.

Taking a closer look at the NPs, it is easy to check that every phrase contributes to the global anaphoric potential of its linguistic context by passing the tag of its reference marker into its own LIST-LU. In the case of the quantificational NP *every student*, two tags are passed, corresponding to the REFMARK value, providing for e-type anaphora, and the VAR value, providing for bound anaphora interpretations. And in the case of the *ctx* node, to illustrate how the nonlinguistic context may be taken into account in the linguistic representation, the reference marker $\langle 415 \rangle$ is obtained from the set of semantic conditions that conventionally may capture the nonlinguistic context.

On the other hand, the context also contributes to establishing the anaphoric potential of each NP. This is ensured by the different clauses of the Binding Domains Principle, which enforce the presence of suitable values of LIST-A, LIST-Z, and LIST-U at the different nodes.

Finally, token identity is ensured between the ANTEC value and the outcome of the different relational constraints that are lexically associated with each NP and express binding constraints. The value of ANTEC is a list that, at this stage of anaphor resolution, records the grammatically admissible antecedents of the corresponding anaphor only in the light of binding constraints.

6. Conclusions

Departing from the coindexation-driven approach for encoding anaphoric dependencies in grammatical representations, we have proposed an alternative methodology where binding constraints are viewed as contributing to circumscribing their contextually determined semantic value. This semantics-driven approach allows a principled integration of binding constraints into grammar that supports both a specification format and a verification methodology free from previous difficulties. Importantly, it also permits a neat interface between the grammatical module of binding and systems of reference processing.

Appendix

In this article, we consider the version of binding constraints formulated within Head-Driven Phrase Structure Grammar (Pollard and Sag 1994, Chapter 6). Recent developments indicate that there are four binding constraints (Xue, Pollard, and Sag 1994;

Branco and Marrafa 1999). Here, the definition of each binding constraint is followed by an illustrative example.

- (7) *Principle A*
 A locally o-commanded short-distance reflexive must be locally o-bound.
 Lee_i thinks [Max_j saw himself_{*i/j}].
- (8) *Principle Z*
 An o-commanded long-distance reflexive must be o-bound.
 [O amigo do Rui_i]_j acha que o Pedro_k gosta dele próprio_{*i/j/k}.
 [the friend of_the Rui] thinks that the Pedro likes of_he PRÓPRIO
 '[Rui's friend]_j thinks that Pedro_k likes him_j/himself_k.' (Portuguese)
- (9) *Principle B*
 A pronoun must be locally o-free.
 Lee_i thinks [Max_j saw him_{i/*j}].
- (10) *Principle C*
 A nonpronoun must be o-free.
 [Kim_i's friend]_j thinks [Lee saw Kim_{i/*j}].

These constraints are defined on the basis of some auxiliary notions.

The notion of *local domain* involves the partition of sentences and associated grammatical geometry into two zones of greater or lesser proximity with respect to the anaphor. The exact definition of the boundary separating the local from the nonlocal domain may vary from language to language. Typically, the local domain tends to correspond to the structure in the grammatical representation that is affected by the selectional capacity and requirements of a predicator.

O-command is a partial order under which, in a clause, the subject o-commands the direct object, the direct object o-commands the indirect object, and so on, following the usual obliqueness hierarchy of grammatical functions, while in a multiclausal sentence, the upward arguments o-command the successively embedded arguments.

The notion of *o-binding* is such that x o-binds y iff x o-commands y and x and y are coindexed, where coindexation is meant to represent anaphoric links.

Acknowledgments

I am grateful to Hans Uszkoreit for advice and helpful discussion, and to Mark Johnson for clarifying criticisms. I am solely responsible for remaining errors.

The results presented here were obtained while I was on leave at the Language Technology Group of the DFKI-German Research Center on Artificial Intelligence, Saarbrücken, Germany, whose hospitality and enthusiastic atmosphere I was very

fortunate to enjoy and I hereby gratefully acknowledge.

References

- Asher, Nicholas and Hajime Wada. 1989. A computational account of syntactic, semantic and discourse principles for anaphora resolution. *Journal of Semantics*, 6:309–344.
- Backofen, Rolf, Tilman Becker, Jo Calder, Joanne Capstick, Luca Dini, Jochen Dörre,

- Gregor Erbach, Dominique Estival, Suresh Manandhar, Anne-Marie Mineur, Gertjan van Noord, Stephan Oepen, and Hans Uszkoreit. 1996. Final report of the EAGLES Formalisms Working Group. Technical Report, EAGLES-Expert Advisory Group on Language Engineering Standards, Luxemburg.
- Barker, Chris and Geoffrey K. Pullum. 1990. A theory of command relations. *Linguistics and Philosophy*, 13:1–34.
- Branco, António. 1998. The logical structure of binding. In *Proceedings of the 36th Annual Meeting of the Association for Computational Linguistics and 17th International Conference on Computational Linguistics (ACL/COLING'98)*, pages 181–187.
- Branco, António. 2000. *Reference Processing and Its Universal Constraints*. Edições Colibri, Lisbon.
- Branco, António and Palmira Marrafa. 1999. Long-distance reflexives and the binding square of opposition. In Gert Webelhuth, Jean-Pierre Koenig, and Andreas Kathol, editors, *Lexical and Constructional Aspects of Linguistic Explanation*. CSLI Publications, Stanford, CA, pages 163–177.
- Bredenkamp, Andrew. 1996. *Towards a Binding Theory for Head-Driven Phrase Structure Grammar*. Ph.D. thesis, University of Essex.
- Carbonell, Jaime and Ralf Brown. 1988. Anaphora resolution: A multi-strategy approach. In *Proceedings of the 12th International Conference on Computational Linguistics (COLING'88)*, pages 96–101.
- Chomsky, Noam. 1980. On binding. *Linguistic Inquiry*, 11:1–46.
- Chomsky, Noam. 1981. *Lectures on Government and Binding*. Foris, Dordrecht.
- Correa, Nelson. 1988. A binding rule for government-binding parsing. In *Proceedings of the 12th International Conference on Computational Linguistics (COLING'88)*, pages 123–129.
- Dalrymple, Mary. 1993. *The Syntax of Anaphoric Binding*. CSLI Publications, Stanford, CA.
- Erbach, Gregor. 1995. Prolog with features, inheritance and templates. In *Proceedings of the 7th Conference of the European Chapter of the Association for Computational Linguistics (EACL'95)*, pages 180–187.
- Fong, Sandiway. 1990. Free indexing: Combinatorial analysis and a compositional algorithm. In *Proceedings of the 28th Annual Meeting of the Association for Computational Linguistics (ACL'90)*, pages 105–110.
- Frank, Anette and Uwe Reyle. 1995. Principle based semantics for HPSG. In *Proceedings of the 7th Conference of the European Chapter of the Association for Computational Linguistics (EACL'95)*, pages 9–16.
- Gawron, Jean Mark and Stanley Peters. 1990. *Anaphora and Quantification in Situation Semantics*. CSLI Publications, Stanford, CA.
- Giorgi, Alessandra, Fabio Pianesi, and Giorgio Satta. 1990. A computational approach to binding theory. In *Proceedings of the 13th International Conference on Computational Linguistics (COLING'90)*, pages 1–6.
- Ingria, Robert and David Stallard. 1989. A computational mechanism for pronominal reference. In *Proceedings of the 27th Annual Meeting of the Association for Computational Linguistics (ACL'89)*, pages 262–271.
- Jacobson, Pauline. 1999. Binding without pronouns. In *Formal Grammar Conference 1999: Symposium on Grammatical Resources and Grammatical Inference*. Abstract.
- Janssen, Theo. 1997. Compositionality. In Johan van Benthem and Alice ter Meulen, editors, *Handbook of Logic and Language*. Elsevier, Amsterdam, pages 417–474.
- Johnson, Mark. 1995. Constraint-based natural language parsing. Course notes, 7th European Summer School in Logic, Language and Information, Barcelona.
- Johnson, Mark and Ewan Klein. 1990. Discourse, anaphora and parsing. In *Proceedings of the 28th Annual Meeting of the Association for Computational Linguistics (ACL'90)*, pages 669–675.
- Kamp, Hans and Uwe Reyle. 1993. *From Discourse to Logic: Introduction to Modeltheoretic Semantics of Natural Language, Formal Logic and Discourse Representation Theory*. Kluwer, Dordrecht.
- Kaplan, Ronald and John Maxwell. 1988. An algorithm for functional uncertainty. In *Proceedings of the 12th International Conference on Computational Linguistics (COLING'88)*, pages 297–302.
- Koenig, Jean-Pierre. 1999. Inside-out constraints and description languages for HPSG. In Gert Webelhuth, Jean-Pierre Koenig, and Andreas Kathol, editors, *Lexical and Constructional Aspects of Linguistic Explanation*. CSLI Publications, Stanford, CA, pages 265–280.
- Lappin, Shalom and Nissim Francez. 1994. E-type pronouns, I-sums, and donkey anaphora. *Linguistics and Philosophy*, 17:391–428.
- Lappin, Shalom and Herbert Leass. 1994. An algorithm for pronominal anaphora resolution. *Computational Linguistics*, 20:535–561.

- Löbner, Sebastian. 1987. Quantification as a major module of natural language semantics. In Jeroen Groenendijk, Dick de Jong, and Martin Stokhof, editors, *Studies in DRT and the Theory of Generalized Quantifiers*. Foris, Dordrecht, pages 53–85.
- Merlo, Paola. 1993. For an incremental computation of intra-sentential coreference. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI'93)*, pages 1216–1221.
- Mitkov, Ruslan. 1997. Factors in anaphora resolution: They are not the only things that matter. In *Proceedings of the ACL/EACL'97 Workshop on Operational Factors in Practical, Robust Anaphora Resolution*, pages 36–51.
- Mitkov, Ruslan. 1998. Robust pronoun resolution with limited knowledge. In *Proceedings of the 36th Annual Meeting of the Association for Computational Linguistics and 17th International Conference on Computational Linguistics (ACL/COLING'98)*, pages 869–875.
- Pianesi, Fabio. 1991. Indexing and referential dependencies within binding theory: A computational framework. In *Proceedings of the 5th Conference of the European Chapter of the Association for Computational Linguistics (EACL'91)*, pages 39–44.
- Pollard, Carl and Ivan Sag. 1994. *Head-Driven Phrase Structure Grammar*. University of Chicago Press, Chicago.
- Reinhart, Tanya and Eric Reuland. 1993. Reflexivity. *Linguistic Inquiry*, 24:657–720.
- Rich, Elaine and Susann LuperFoy. 1988. An architecture for anaphora resolution. In *Proceedings of the 2nd Conference on Applied Natural Language Processing (ANLP'88)*, pages 18–24.
- Xue, Ping, Carl Pollard, and Ivan Sag. 1994. A new perspective on Chinese *ziji*. In *Proceedings of the 12th West Coast Conference on Formal Linguistics (WCCFL'94)*. CSLI Publications, Stanford, CA.