Title: Anaphor resolution: Is the search optimization rationale flawed?

Auhtor's name: António Horta Branco

Affiliation: Universidade de Lisboa, Departamento de Informática

Postal address: Faculdade de Ciências, Departamento de Informática, Campo Grande, 1749-016 Lisboa, Portugal

Telephone: +351 21 757 06 06

Fax: +351 21 757 00 84

Email: Antonio.Branco@di.fc.ul.pt

Homepage: www.di.fc.ul.pt/~ahb

Title: Anaphor Resolution: Is the search optimization rationale flawed?

Abstract: We discuss the primitives underlying mainstream cognition-driven approaches to anaphora and isolate the search optimization rationale behind the constraints on anaphor resolution. This rationale implies some predictions (i) about the existence of natural classes of anaphors such that every element in each class has the same set of admissible antecedents, and (ii) about the fact that the corresponding sets of admissible antecedents bear specific relations among them. The four binding classes are the naturally occurring classes of anaphors satisfying (i), and we verify that the predictions concerning (ii) are not empirically satisfied by the corresponding sets of admissible antecedents. We discuss in which sense this negative result casts doubts on the empirical support of the search optimization rationale for cognitive accounts of anaphor resolution.

Resumo: Discutimos as assunções elementares que presidem às análises da anáfora de inspiração cognitiva mais proeminentes e isolamos o rationale baseado na optimização da busca subjacente às restrições relativas à resolução de anáforas. Este rationale conduz a algumas previsões (i) acerca da existência de classes naturais de anáforas tais que cada elemento dessas classes tem o mesmo conjunto de antecedentes admissíveis e (ii) acerca do facto de os conjuntos de antecedentes admissíveis correspondentes exibirem relações específicas entre si. As quatro classes de ligação (*binding classes*) constituem as classes de anáforas que ocorrem naturalmente e satisfazem (i), e verificamos que as previsões relativas a (ii) não são empiricamente satisfeitas pelos correspondentes conjuntos de antecedentes admissíveis. Discutimos em que medida este resultado negativo levanta dúvidas acerca da

plausibilidade empírica do rationale baseado na optimização da busca para as análise da resolução de anáforas de base cognitiva.

Keywords: : Natural language, Cognitive models of anaphora, Anaphor resolution, Binding principles.

1. Introduction

Mainstream cognitive models of nominal anaphor resolution envisage this linguistic process as a particular case of search optimization, where a “divide to conquer” processing strategy is adopted: The search space of working memory with antecedent candidates is "sectioned", each "section" containing the admissible antecedents for anaphors of different classes.

In Section 2 we discuss the details of the shared rationale behind these models. In particular, we see that the natural metrics that has been generally assumed for sectioning the search space for anaphor resolution is attentional prominence: Given the items that happen to be in short-term memory at a certain moment, anaphors of different classes are resolved against antecedent candidates with different degrees of attentional prominence.

In Section 3, we isolate central predictions implied by this rationale, concerning the relationship among the sets of admissible antecedents. We discuss also how attempts that seek to test such predictions have faced problems in finding objective criteria to isolate the relevant sets of antecedent candidates, and argue that a suitable way forward is to consider the binding classes of anaphors and the corresponding sets of admissible antecedents.

Binding classes are intensionally defined by means of constraints that are valid across different languages, known as binding constraints or principles. The definition of binding constraints capture empirical generalizations and are primarily aimed at delimiting the relative positioning of anaphors and their admissible antecedents in the grammatical geometry of linguistic constructs. Each anaphor belongs to the class of either short-distance reflexives (complying with Principle A), long-distance reflexives (Principle Z), pronouns (Principle B) or non-pronouns (Principle C) because, with respect to any syntactic position, it has the same set of admissible antecedent candidates of the anaphors of its class.

In Section 4, we check the predictions underlined in the previous Section against the observed sets of antecedent candidates for the different binding classes, and discuss why these predictions are not confirmed. We will elaborate on how this militates against the insight embodied in the mainstream rationale referred to above.

Finally, in Section 5, we discuss the implications of this negative result both for a cognitive account of binding classes and for the empirical justification of current mainstream cognitive models of anaphor resolution.

2. Models for Anaphor Resolution

Anaphor resolution has been a major topic of inquiry for the research into the nature of linguistic processes. In the present Section, we seek to isolated what we think to be the common rationale behind the accounts proposed in the literature on cognitive models of anaphor resolution. The following quotation from (Gundel *et al.*, 1993) on the essential polysemy and polymorphism of natural language provides a suitable starting point: “One of the more interesting facts about human language is that we can use different forms to refer to the same thing, and the same form can be used to refer to many different things” (p. 276).

2.1. Polysemy and Interpretation

Polysemy latu sensu appears as a convenient solution in a representational system that has finite resources to represent a virtually infinite number of entities. As there might be an undetermined number, say, of students with yellow t-shirts, it would be unbearable to a finite mind to have a specific representational device for each such student rather than the polysemous form *the student with a yellow t-shirt*. Polysemy, the feature that “the same form can be used to refer to many different things”, is a helpful device to bring the linguistic system, or at least its lexicon, to a "handy" dimension.

Although polysemy appears as a convenient feature of the linguistic system, human language is not maximally reduced to a single hyper-polysemous blurb-lexeme that would be used to refer to any possible entity. The reason for this is to be found in the fact that, the gain elicited by polysemy is obtained at the cost of a non negligible side-effect: If the system counts on polysemy, it has also to count on the help of an interpretive task to assign to a polysemous expression, in a specific occurrence, the entity actually referred to by that expression. The benefit brought by a higher level of polysemy raises thus also important costs due to this non negligible interpretation overhead. As a result of these opposite trends, the system finds some balance point under the motto "enough polysemy for affordable interpretation processing".

2.2. Interpretation and polymorphism

Polymorphism latu sensu on the other hand, appears as a convenient solution in a representational system that uses some interpretative schema to handle polysemy. As speakers keep referring, say, to the same student with yellow t-shirt, they can avoid going through the whole interpretive process of deciding which one of the possibly indefinite number of referents of *the student with a yellow t-shirt* should be picked out: This can be done by using a different, anaphoric form (*he*, *this guy*, *the student*,…) which is interpretively parasitic of the initial reference and simply signals that the same entity is being referred to. Polymorphism, the feature that "we can use different forms to refer to the same thing", is a helpful device to lower the processing effort involved in the interpretation process needed for taming polysemy.

Anaphora should thus be understood not just as a manifestation of polymorphism, but as a phenomenon emerging at the juncture of polysemy and polymorphism. Anaphoric polymorphism appears as a convenient solution in a system that has to handle polysemy in real time.

2.3. Polymorphism and anaphor resolution

Looking at anaphora under this broad perspective permits a straightforward justification of the correlation, frequently mentioned in the literature, between anaphors and expressions that have weaker semantic content. But more important for the point to be made in the present article, this broad perspective allows us also to bring to light a shared rationale of current proposals for a cognitive grounding of anaphor resolution.

Anaphora is used to avoid repeatedly going through the whole interpretive, polysemy reducing process. As speakers keep referring, say, to the same person already referred to by *the student with yellow t-shirt* they use some anaphoric form, e.g. *he*, and avoid going again through the whole interpretive process of deciding which one of the possibly indefinite number of referents of *the student with a yellow t-shirt* should be picked out: This is done via the so called resolution of the anaphor *he*, by means of which the identification of the antecedent of the anaphor — *the student with a yellow t-shirt* in this example — takes place and the corresponding referent can be used to assign an interpretation to *he*.

Cognition-driven approaches to anaphor resolution have typically seen this phenomenon as a case where a cognitive process is reformatted in terms of a simpler process: The cognitive search in a large, long-term or semantic memory involved in context-driven interpretation is taken over by a search for the recently activated antecedent of the anaphor in the shorter, working or short-term memory — only in case a suitable antecedent fails to be available in the working memory, the anaphor resolution process proceeds to a search in the semantic memory. Anaphora is thus seen as a case of search optimization by means of the reduction of the search space for interpretation.

2.4. Anaphor resolution and differentiation of anaphoric capacity

This search optimization rationale for anaphor resolution has been further explored in cognitive models of anaphora.

Differentiation of anaphoric capacity is used to avoid going through the scanning of the whole working memory in the anaphor resolution process. As speakers refer again, say, to the same person already referred to by *the student with yellow t-shirt*, the specific anaphoric form they use, e.g. *the student*, *he* or *himself*, depends on the relative position of the representation of the referent of *the student with yellow t-shirt* in the working memory: Different types of anaphors have thus been assumed to pick referential items from different "sections" of the relevant search space.

First, as we saw in the previous subsection, in order to reduce interpretive overhead due to polysemy by means of anaphoric polymorphism, the search optimization is implemented in terms of the reduction of the size of the search space, with the transfer of the search process from the long term memory to the short term memory whenever possible. Second, in order to reduce the scanning steps in the working memory for anaphoric resolution, the search optimization is implemented in terms of a "divide to conquer" strategy: The search space for finding antecedents for anaphors is "sectioned", each section being reserved to be searched for the resolution of anaphors of a specific type.

For this schema to work, there has to be some feature that discriminates different items in the working memory from one another and induces a partial order over them. This order is typically established according to the attentional prominence that each such item bears. Attentional prominence is assumed to reflect a natural metrics for "distance" in the relevant cognitive search space, with less attentionally prominent items being the ones that take longer to be retrieved.

When focussing in the anaphor resolution process under the broad angle just presented, one is not concerned with the specific processes by which items are brought into working memory, how they happen to be ordered, reinforced, removed, etc. On the other hand, we think also that no rejection of any of the different cognitive models of anaphora or reference argued for in the literature (vd. Festas, 2003) is implied. Given the items possibly in working memory at a certain moment, one is rather focussed on the fact that anaphors of a given type can thus be resolved against items with a certain attentional prominence, while anaphors of another type are resolved against items with some other degree of attentional prominence.

Skimming through the literature, one finds different proposals concerning the number of sections into which the search space for anaphor resolution is expected to divide. Just a few examples: Authors like Guindon (1985) or Givón (1992) discuss a division, respectively, into two and three "sections". Gundel *et al.* (1993), in turn, proposes a schema that may extend the division up to six "sections", depending on the specific language at stake.

3. Implications

According to the essential tenets of this rationale, different sorts of anaphors — whose antecedent entities are to be found in different “sections” of the search space — are expected to have different sets of admissible antecedent entities.

3.1. Predictions

The strong prediction is that anaphors of different types have different, *disjoint* sets of antecedents. This claim can be found, for instance, in (Garrod and Sanford, 1982).

Another, weaker but also plausible prediction in this connection is that, if the different sets of admissible antecedents turn out not to be disjoint, they would at least be expected to be *successively included* within each other. If we admit that an anaphor is of a given type such

that it is sensitive to items with a certain degree of attentional prominence, it is not a contradiction to accept that this anaphor may also be sensitive to items with a higher degree of prominence. This is the intuition behind the approach, for instance, of Gundel *et al.* (1993, 1998).

The search optimization rationale for anaphor resolution — with the assumed correlation between anaphoric forms and attentional prominence of antecedent candidates — can thus be seen as inducing a delimitation of anaphors into different natural classes. These classes are circumscribed in terms of the antecedents that the corresponding anaphors admit: A given class of anaphors is defined because every anaphor in that class can be resolved against the same set of antecedents.

The point worth stressing here is that this establishes a very interesting and self-contained line of empirical inquiry: If we succeed in isolating different sets of admissible antecedents, then we will succeed in isolating natural, cognitively motivated classes of anaphors. This line of inquiry is one of major relevance also because, if we find such natural classes of anaphors, then we are providing a piece of empirical support of paramount importance for the whole conjecture embodied in the search optimization rationale.

3.2. Fuzzy delimiters

A first step towards pursuing this research path is to find a methodological device that allows to categorize items according to their attentional prominence. This involves finding a suitable scale of the attentional prominence of admissible antecedent entities. Besides, we need also objective criteria to decide which item in the scale a given anaphor should be put in correspondence with. The pursuing of these goals has been reported at various places in the literature, cf. among others, (Prince, 1981) and (Gundel *et al.* 1993).

The scale used to evaluate the attentional status of the cognitive item against which a given anaphor is resolved is typically defined by means of a set of keywords, like "familiar", "activated", "evoked", "uniquely identifiable", "brand new", etc. These keywords come with informal definitions under the form of example sentences and a discussion of some cases to which they may apply. The keywords come also with a hierarchy, where the relative positioning of each keyword in the scale is defined vis a vis the other keywords.

This sort of approach to define a scale of attentional prominence seems to be flawed, in our view, in some crucial aspects.

There is not an empirical justification for the number of required keywords, i.e. of the distinct degrees of relevant attentional prominence.

Keywords are defined in such a way that the boundaries between the degrees of prominence they are supposed to delimit are not clear.

Above all, there is no empirically well defined criteria to unequivocally decide which point of the scale is the antecedent of an anaphor in a specific occurrence in correspondence with.

4. Experiment

These recurrent shortcomings represent a considerable drawback for the goal of finding empirical support to the search optimization rationale of anaphor resolution. The line of argument we would like to explore in the present article is that overcoming this drawback may involve changing the angle from which the correlation between natural classes of anaphors and search optimization could or should be addressed.

Instead of in the first place looking for objective criteria to identify attentional status of items and then trying to use them to possibly delimit classes of anaphors, we should take into account naturally occurring classes of anaphors — empirically motivated precisely on the basis of differences concerning the classes of their admissible antecedents — and try to clarify

the possible cognitive underpinnings of such classes. In particular, one should discuss whether and how such classes may fit into a search optimization rationale for anaphor resolution.

4.1. Binding classes

The most notorious classes of anaphors obtained via grouping of the corresponding sets of admissible antecedents are the so called binding classes. Each of these classes contains all and only the anaphors that may pick an antecedent from the same set of admissible antecedents. A classical contrast permitting to illustrate the kind of difference at stake is the one between *Peter said John described Tom to himself* and *Peter said John described Tom to him*: While *himself* have *John* and *Tom* as admissible antecedents but not *Peter*, *him* has Peter as admissible antecedent (and possibly other antecedents introduced in the discourse or the context), but not *John* or *Tom*. Accordingly, *himself* and *him* are said to belong to different binding classes, the former to the class of the so called short-distance reflexives, the latter to the class of the so called pronouns.

The members of a given binding class can be intensionally characterized as those anaphors that are ruled by a specific binding constraint, with this constraint expressing an objective criterion to categorize anaphors according to one of the different available binding classes. Such binding constraints capture empirical generalizations and are aimed at delimiting the relative positioning of anaphors and their admissible antecedents in grammatical geometry.

Since their first formulation in (Chomsky, 1980, 1981), the definition of binding principles has been the focus of intense research, from which a binding theory of increased empirical adequacy has emerged. From an empirical perspective, binding constraints, or binding principles, 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 non-trivial symmetry, which lends them

a modular nature. Accordingly, they have been considered one of the most robust modules of grammatical knowledge, usually known under the term of "binding theory".[1]

Recent developments of (Pollard and Sag, 1994), in particular (Xue *et al.*, 1994, Branco and Marrafa, 1999, Branco 2000), indicate that there are four binding constraints. Below, the definition of each principle is illustrated by an example with relevant contrasts:[2]

Principle A

If a short-distance reflexive is locally o-commanded, it must be locally o-bound.

{...*ant*$_c$..} [*O amigo do Lee*$_i$]$_j$ *acha que* [*o vizinho do Max*$_k$]$_l$ *gosta de si próprio*$_{*c/*i/*j/*k/l}$. (Portuguese)

[Lee$_i$'s friend]$_j$ thinks [Max$_k$'s neighbour]$_l$ likes himself$_{*c/*i/*j/*k/l}$.

Principle Z

If a long-distance reflexive is o-commanded, it must be o-bound.

{...*ant*$_c$..} [*O amigo do Lee*$_i$]$_j$ *acha que* [*o vizinho do Max*$_k$]$_l$ *gosta dele próprio*$_{*c/*i/j/*k/l}$.

[Lee$_i$'s friend]$_j$ thinks [Max$_k$'s neighbour]$_l$ likes him$_{*c/*i/j/*k}$/himself$_l$.

---

[1] Vd. Dopkins and Nordlie, 1995 and van der Lely and Stollwerck, 1997 and references cited therein for an overview of psycholinguistic research on binding constraints.

[2] Coindexation marks anaphoric links between anaphors and their tentative antecedent(s); indexes prefixed by '*' mark non admissible anaphoric links; and '{...*ant*$_c$...}' represent tentative antecedents available outside the sentence, in the discourse or in the context.

We are using examples of Portuguese, a language with anaphors of each of the four binding classes. Some languages may not have anaphors of every binding type.

Principle B

A pronoun must be locally o-free.

{...*ant*$_c$..} [*O amigo do Lee*$_i$]$_j$ *acha que* [*o vizinho do Max*$_k$]$_l$ *gosta dele*$_{c/i/j/k/*l}$.

[Lee$_i$'s friend]$_j$ thinks [Max$_k$'s neighbour]$_l$ likes him$_{c/i/j/k/*l}$.

Principle C

{...*ant*$_c$..} A non-pronoun must be o-free.

[*O amigo do Lee*$_i$]$_j$ *acha que* [*o vizinho do Max*$_k$]$_l$ *gosta do rapaz*$_{c/i/*j/k/*l}$.

[Lee$_i$'s friend]$_j$ thinks [Max$_k$'s neighbour]$_l$ likes the boy$_{c/i/*j/k/*l}$.

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 less proximity with respect to the anaphor. Typically, the local domain coincides with the predication domain of the predicator subcategorizing the anaphor. In some cases, there may be additional requirements that the local domain is circumscribed by the first upward predicator that happens to be finite, bears tense or indicative features, etc.[3] For instance, in the example *Lee's friend thinks* [*Max's neighbour likes him*] the local domain is of *him* indicated between square brackets.

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, being that in a multi-clausal sentence, the upward arguments o-command the successively embedded arguments. For instance, in the example *The girl who said that Peter knows Max thinks Max's neighbour likes him*, we get

[3] For details, see (Dalrymple, 1993).

the following o-command relations: *The girl who Peter said that Max knows Max* < *Max's neighbour* < *him*, and *Peter* < *Max* < *who*.

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.[4] For instance, in the example *Lee's friend thinks Max's neighbour likes himself*, *Lee's friend* (non locally) o-binds *himself*, *Max's neighbour* locally o-binds it, and *Lee* and *Max* does not o-bind it.

Note that, given their conditional definition, Principles Z and A are complied with if the reflexives are in so called *non exempt positions*, that is if they are, respectively, o-commanded and locally o-commanded.

It is now well established in the literature that there is a distinction between constraints for anaphor resolution (excluding tentative antecedents from the set of admissible antecedent candidates) and preferences (making the resolution process to converge on the actual antecedent). Binding constraints are thus to be counted in the set of such constraints, though they are not the only ones.[5]

---

[4] There are anaphors that are subject-oriented, in the sense that they only take antecedents that have the grammatical function Subject. Some authors (e.g. Dalrymple, 1993) assume that this should be seen as an intrinsic parameter of binding constraints and aim at integrating it in their definition. In this point we follow previous results of ours reported in (Branco, 1996), where the subject-orientedness of anaphors is argued to be, not an intrinsic feature of binding constraints, but one of the surfacing effects that result from the non linear obliqueness hierarchy associated with some predicators (or to all of them in some languages).

[5] For details on the distinction between constraints and preferences in anaphor resolution, and their listings, see Carbonell and Brown, 1988; Rich and Luperfoy, 1988; Asher and Wada, 1989; Lappin and Leass, 1994; Mitkov, 1997; Branco, 2000.

4.2. Sets of admissible antecedents

As discussed above, the search optimization rationale for anaphor resolution implies some predictions concerning the relations between the different natural classes of admissible antecedents for anaphors. These classes are expected to be either disjoint — strong prediction —, or successively included within each other — weak prediction. Given the binding classes just presented, we can now check if they conform to these predictions. For each of the four binding classes, we delimit the corresponding sets A, B, C and Z of admissible antecedents and then check how they relate to each other.

In order to proceed with this test, first, we have to fix a non exempt position **x** in a generic multi-clausal grammatical structure, like the one used above for the examples illustrating the different binding principles, that can be schematically represented as

{...disc/cont*...}...nloc*...[...noc*...]...[...loc*...[...noc*...]...**x**...]$_{\text{LocalDom}}$...

where `nloc*`, `noc*` and `loc*` stand, respectively, for positions of non-local o-commanders, non-o-commanders and local o-commanders. Second, we have to successively instantiate **x** with an anaphor from each different binding class. We will then be able to observe what are the relations among the sets of admissible antecedents of each binding class.

If we assume that **x** is any anaphor complying with principle A, we see that the admissible antecedents of **x** form the set of its local o-commanders, which we can call the set A.

In case **x** is an anaphor complying with principle Z, the set Z of its admissible antecedents is made of its o-commanders.

When **x** is an anaphor ruled by Principle B, the set B of its admissible antecedents contains all the antecedents that are non-local o-commanders of **x**.

Finally, the set C of the admissible antecedents of **x** when this is an anaphor complying with principle C has all the items that are non-o-commanders of **x**.

Given the definitions of the o-command relation, and from a maximally generic point of view, the formal relations between these four sets of admissible antecedents are the following:

$$A \subset Z \;\&\; A \cap B = \varnothing \;\&\; A \cap C = \varnothing$$

$$Z \cap B \neq \varnothing \;\&\; Z \cap C = \varnothing$$

$$B \subset C$$

By representing each of these sets in graphical terms,

Insert Table 1 about here

the relations among them can be rendered in a more perspicuous way by means of the following diagram:

Insert Table 2 about here

It is straightforward to see that the admissible antecedents of short-distance reflexives are admissible antecedents of long-distance reflexives ($A \subset Z$); some admissible antecedents of long-distance reflexives are admissible antecedents of pronouns ($Z \cap B \neq \varnothing$); and the admissible antecedents of non-pronouns are admissible antecedents of pronouns ($C \subset B$).

From another perspective, this amounts to say that for a given possible antecedent of an anaphor in position **x**, it is the case that there are always at least two different types of

anaphors that can fill **x** and take that antecedent.[6] Or alternatively, for a given anaphor interpreted against a given antecedent, that anaphor can always be replaced at least by another one of a different binding type that can take the same antecedent.

In any case, what is crucial to note for our experiment is that the sets of admissible antecedents per anaphor type are not mutually disjoint. They are neither successively included within each other.

This does not match either the strong or the weak prediction implied by the search optimization rationale for anaphor resolution.

5. Discussion

In this paper we sought to isolate the search optimization rationale behind the constraints expected to bear on anaphor resolution.

This rationale implies some predictions about the existence of natural classes of anaphors such that, with respect to any position of occurrence, every element in each such class have the same set of admissible antecedents. In particular, it implies that these sets of admissible antecedents are predicted to exhibit certain relations among them: They are expected either to be disjoint, or at least to be successively included within each other.

Given the current state of the art of the research on anaphora, the four binding classes are the naturally occurring classes of anaphors satisfying the criterion pointed out above: Each binding class contains all and only the anaphors that, for a given grammatical position, have the same set of admissible antecedents.

[6] If one considers also exempt syntactic positions, then even reflexives have possible antecedents that may also be antecedents of pronouns and non-pronouns.

The result we argued for in the present paper is that the four sets of admissible antecedents of the four binding classes do not conform to the predictions underlined above: They are neither disjoint nor successively included in each other: While, two of them, A and C, are strictly included in the other two, Z and B — with A ⊂ Z and C ⊂ B —, the latter are not disjoint neither included in one another.

Given that the relations predicted by the search optimization rationale are not observed for the sets of admissible antecedents corresponding to binding classes, these natural classes are not offered any principled explanation by this rationale. Moreover, given that these are objectively determined natural classes of anaphors across natural languages, this result casts serious doubts that the search optimization rationale may provide a clear-cut justification for anaphor resolution and its constraints.

This should not be seen as forcing the inference that cognitively rooted factors (such as attentional prominence associated with recency of mention, just to refer an example) do not play an important role in anaphor resolution, at least as preference factors. Nor should it be seen as implying that binding constraints have been proved not to have any cognitive justification.

This result, if correct, shows that current cognitive models of anaphor resolution, crucially based on the search optimization rationale, make predictions that are infirmed by the very significant empirical generalizations embodied in the definition of binding classes.

6. References

Asher, N. and H. Wada (1988). A computational account of syntactic, semantic and discourse principles for anaphora resolution, *Journal of Semantics*, *6*, 309-344.

Branco, A. (1996). Branching split obliqueness at the syntax-semantics interface. In *Proceedings, 16th International Conference on Computational Linguistics (COLING96)*, 149-156.

Branco, A. and P. Marrafa (1997). Subject-oriented and non subject-oriented long-distance anaphora: an integrated approach. In *Proceedings, 11th Pacific-Asia Conference on Language, Information and Computation (PACLIC-96)*. Seoul: Kyung Hee University, 21-31.

Branco, A. (2000), *Reference Processing and its Universal Constraints*, Edições Colibri, Lisbon.

Carbonell, J. and R. Brown (1988). Anaphora resolution: a multi-strategy approach, In *Proceedings, The 12th International Conference on Computational Linguistics (COLING88)*, 96-101.

Chomsky, N. (1980). On binding. *Linguistic Inquiry*, *11*, 1-46.

Chomsky, N. (1981). *Lectures on Government and Binding*, Foris, Dordrecht.

Dalrymple, M. (1993). *The Syntax of Anaphoric Binding*. CSLI Publications, Stanford.

Dopkins, S. and J. Nordlie (1995). Processes of anaphor resolution. In R. Lorch and E. O'Brien (eds.), *Sources of Coherence in Reading*, Lawrence Erlbaum, Hillsdale, 145-157.

Festas, I. (2003). Resolução anafórica nas diversas teorias acerca da compreensão do discurso. O papel dos processos *bottom-up* e *top-down*. *Psychologica*, this volume.

Garrod, S. and A. Sanford (1982). The mental representation of discourse in a focussed memory system: Implications for the interpretation of anaphoric noun phrases. *Journal of Semantics*, *1*, 21-41.

Givón, T. (1992). The grammar of referential coherence as mental processing instructions. *Linguistics*, *30*, 5-55.

Gundel, J., N. Hedberg and R. Zacharski (1993). Cognitive status and the form of referring expressions in discourse. *Language*, *69*, 274-307.

Gundel, J. (1998). Centering theory and the givenness hierarchy: Towards a synthesis. In M. Walker, A. Joshi and E. Prince (eds.), 183-198.

Guindon, R. (1985). Anaphora resolution: short-term memory and focusing. In *Proceedings, 23rd Annual Meeting of the Association for Computational Linguistics (ACL85)*, 218-227.

Lappin, S. and H. Leass (1994). An algorithm for pronominal anaphora resolution, *Computational Linguistics*, *20*, 535-561.

Mitkov, R. (1997). Factors in anaphora resolution: they are not the only things that matter. a case-study based on two different approaches. In *Proceedings of the ACL/EACL97 Workshop on Operational Factors in Practical, Robust Anaphora Resolution.* Association for Computational Linguistics.

Pollard, C. and I. Sag (1994). *Head-Driven Phrase Structure Grammar*. The University of Chicago Press, Chicago.

Prince, E. (1981). On the reference of indefinite-*this* NPs. In A. Joshi, B. Webber and I. Sag (eds.), *Elements of Discourse Understanding*. Cambridge University Press, Cambridge.

Rich, E., and S. LuperFoy (1988). An architecture for anaphora resolution, In *Proceedings, 2nd Conference on Applied Natural Language Processing*, 18-24.

van der Lely, H. and L. Stollwerck (1997). Binding theory and grammatical specific language impairment in children. *Cognition*, *62*, 245-290.

Walker, M., A. Joshi and E. Prince. (eds.) (1998a). *Centering In Discourse.* Oxford University Press, Oxford.

Walker, M., A. Joshi, and E. Prince (1998b). Centering in naturally occurring discourse: An overview. In M. Walker, A. Joshi and E. Prince (eds.) (1998a), 1-30.

Xue, P., C. Pollard, and I. Sag (1994). A new perspective on chinese *ziji*. In *Proceedings of the West Coast Conference on Formal Linguistics (WCCFL'94)*, CSLI Publications, Stanford.

Table 1:

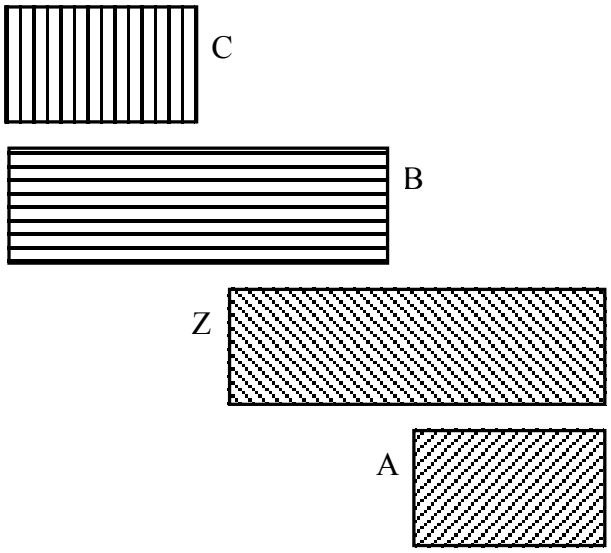

Table 2: