Binding Classes and Search Optimisation in Anaphor Resolution

António Branco

University of Lisbon
Faculdade de Ciências de Lisboa, Dep. de Informática,
Campo Grande, 1700 Lisboa, Portugal
www.di.fc.ul.pt/~ahb

Abstract

In this paper we study binding constraints under the perspective of their cognitively grounded justification. After delimiting the mainstream rationale and primitives of cognition-driven approaches to anaphora and anaphor resolution, we analyze the reasons why this rationale does not lead to a principled account of the binding constraints. More crucially, we discuss in which sense this negative result casts doubts on the empirical support of this rationale itself. Finally, we discuss how recent results on the underlying quantificational nature of binding constraints point towards an alternative rationale for a cognitively grounded explanation of binding facts.

1. Introduction

Cognition-driven approaches to anaphor resolution tend to envisage this phenomenon as a case where a “divide to conquer” processing strategy is adopted: The search space of antecedent candidates is "sectioned", each "section" containing the admissible antecedents for the anaphors of a given class.

In Section 2 we discuss the details of the rationale behind such approaches. In particular, we will see that the natural metrics that has been generally assumed for sectioning this search space 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.

Against this background, the issue we discuss in Section 3 is whether this rationale for anaphora processing provides for a suitable understanding of the possible cognitive underpinnings of binding constraints.

Binding constraints delimit the relative positioning of anaphors and their admissible antecedents in the grammatical geometry. Short-distance reflexives (Principle A), long-distance reflexives (Z), pronouns (B) and non-pronouns (C) belong to different classes because they have different sets of antecedent candidates when occurring in the same syntactic position.

The empirical issue we would like to bring to the fore is that, for a given syntactic position filled in by anaphors of different binding classes, the corresponding sets of antecedent candidates are not disjoint; they are not either successively included within each other.

We will elaborate on how this militates against the insight implied by the rationale describe above: The "sectioning" of the search space for anaphor resolution induced by binding constraints can hardly be understood as being shaped by an optimization strategy in the search for antecedents like the one referred to above.

Finally, in section 4, on the basis of recent results on the underlying quantificational nature of binding constraints, we discuss how the evidence from binding classes may help to shed light on a different paradigm of the possible cognitive underpinnings of anaphor resolution.

2. Search in a prominence scale

Research on reference processing and anaphor resolution gave rise to a significant volume of work and results in the literature.[1] Like in the case of other linguistic phenomena, the search for eventual cognitive underpinnings of anaphora appears to be a natural endeavor for the scientific effort to uncover possibly more fundamental and empirically motivated explanations for the language facts at stake.

2.1. Polysemy, polymorphism and anaphora

In this connection, it is worth considering the first sentence of (Gundel *et al.*, 1993) on the essential polysemy and polymorphism of natural language: “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).

This equation helps to emphasize that anaphora is typically seen as an eloquent case of one of these two facets, i.e. “different forms refer to the same thing”. Anaphoric links are established between different expressions — different forms — that either refer to the same entity or are in an essential way interpretively intertwined in sustaining reference to a given entity and possibly cospecifying its properties.

Although this facet is certainly the most noticeable one, a more comprehensive picture of anaphora in cognitive terms can be drawn when the other facet of the pair polysemy vs. polymorphism is also taken into account.

Because “the same form can be used to refer to many different things”, given the entities that a polysemic expression can be used to refer to, there certainly exists a non-negligible interpretive task to assign that expression, in a specific occurrence, the entity actually referred to at that occurrence. In order to minimize this interpretive effort, when an entity already referred to is referred to

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

again (as such, in part, as part of, etc.), anaphors are used to avoid going through that interpretive process again.

Anaphora should thus be understood not just as a singular manifestation of the circumstance that “different forms may refer the same thing”, but as a phenomenon emerging at the juncture of polysemy and polymorphism.

Polysemy 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 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 polysemic form *the student with a yellow t-shirt*. Anaphoric polymorphism, in turn, appears as a convenient solution in a system that has to handle polysemy in real time: 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 form which is interpretively parasitic and simply signals that the same entity is being referred to.

2.2. Search optimization: divide to conquer

Setting the analysis of anaphora under this wider perspective provides for a straightforward justification of the correlation, frequently mentioned in the literature, between anaphors and expressions which have weaker semantic content. Besides, it also allows to bring to light a shared rationale of current proposals for a cognitive grounding of anaphora and anaphor resolution.

Cognition-driven approaches to anaphor resolution have typically seen, in one way or other, this issue as a case where a cognitive process is reformated in terms of a simpler process.

Some entities may be recurrently referred to. In these cases of recurrent reference, the task of polysemy reduction is taken over by a task of anaphoric resolution. The cognitive search in a large long-term or semantic memory is avoided by means of a search in the shorter, working or short-term memory. Anaphora is thus seen as a case of search optimization, where there is a reduction of the cognitive search space for semantic interpretation.

Building further on this rationale for anaphora as search optimization, different types of anaphors have been assumed to pick referents from different “sections” of the relevant search space. While interpretive overhead due to polysemy is reduced by means of anaphoric polymorphism, anaphoric resolution, in turn, is facilitated by 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 a feature to discriminate different items in relevant working memory from one another. This is ensured by assuming that they enter a partial order and that they are distinguished from one another given their relative positioning in it. This order is typically established according to the attentional prominence that each such item is assumed to bear. Note that attentional prominence reflects a natural metrics for "distance” in the relevant cognitive search space, with less attentionally prominent items being the ones which take longer to be retrieved, i.e. whose search time for an antecedent is shorter.

Given the items possibly in relevant working memory, anaphors of a given type can thus pick up items with a certain attentional prominence, while anaphors of another type pick items with some other degree of attentional prominence in the relevant cognitive search space.

Skimming through the literature, it is possible to find 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.

2.3. Natural classes of anaphors from natural classes of antecedents

In line with the essential tenets of the rationale above, one would then expect that different sorts of anaphors — whose antecedent entities are to be found in different “sections” of the search space — have different sets of admissible antecedent entities.

The strong prediction would thus be that anaphors of different types have different, *disjoint* sets of antecedents, a position that can be found e.g. in (Garrod and Sanford, 1982).

Another, weaker but plausible prediction in this connection would be that, if the different sets of 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 searches or is sensitive to items with a certain degree of attentional prominence, it is not a contradiction to accept that that anaphor may also be sensitive to items with a higher degree of prominence. This is the intuition behind the approach of Gundel *et al.* (1993, 1998).

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

The point worth stressing then is that this sets up a very interesting and self-contained line of empirical inquiry: If we succeed in isolating different classes 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 optimisation rationale.

3. Antecedent accessibility through grammatical structure

In the previous section we sought to clarify the primitives of the rationale underlying cognition-based approaches to the issue of anaphor resolution. The implications of this rationale were explored up to the point were a relevant line of research emerged which may bring important empirical support to it. A major goal in this line of research is to inquire about possible natural classes of anaphors defined on the basis of natural classes of antecedent entities, where the latter be delimited in terms of attentional prominence.

3.1. Fuzzy delimiters

A first step in this direction is thus to find a methodological device that allows to categorize a given item according to its attentional prominence. This involves finding a suitable scale of the attentional prominence of admissible antecedent entities and objective criteria to decide where in the scale an anaphor should be put in correspondence with. The pursuing of this goal 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 corresponding to a given anaphor is typically defined by means of a set of keywords, like "familiar", "activated", "evoked", "uniquely identifiable", "brand new", etc. These keywords come with definitions under the form of examples 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, that is of 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 objective criteria to unequivocally decide which point of the scale is an anaphor in correspondence with.

These shortcomings in defining a methodological device that allows to categorize a given item according to its attentional prominence represent a drawback for the goal of finding empirical support to the search optimization rationale of anaphor resolution. This does not mean, however, that this might be seen as empirical justification to reject the conjectures embodied in such rationale.

Our point in bringing to light the above deadlock is not to gather evidence aiming at dismissing a cognition-driven rationale as such. Our line of argument is rather that overcoming this deadlock may involve changing the angle from which we are addressing the correlation between natural classes of anaphors and search optimization. Instead of in first place looking at objective criteria to identify attentional status and then trying to use them to possibly delimit classes of anaphors, we should take into account actual natural classes of anaphors — empirically motivated on the basis of differences in classes of admissible antecedents — and try to clarify their eventual cognitive underpinnings. In particular, we should discuss whether and how such classes may fit into a search optimization rationale for anaphor resolution.

3.2. Binding classes

The most notorious classes of anaphors obtained via grouping of corresponding sets of antecedents are binding classes. Each of these binding classes contains anaphors that may pick an antecedent from the same set of admissible antecedents. The members of a given class are intensionally characterized as those anaphors that obey a specific binding constraint, with this constraint expressing an objective criterion to categorize anaphors.

Binding constraints delimit the relative positioning of anaphors and their admissible antecedents in grammatical geometry. In the context of our discussion, a noteworthy point is that, in first place, binding constraints correlate the interpretation of anaphors with grammatical structure, not with the eventual cognitive status of the cognitive representation of the admissible antecedents.[2]

From an empirical perspective, binding constraints 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 as binding theory.

Since their first formulation in (Chomsky, 1980, 1981), the specification of binding constraints has been the focus of intense research in last decades, from which a binding theory of increased empirical adequacy has emerged. Recent developments of (Pollard and Sag, 1994) indicate that there are four of such constraints (vd. Xue *et al.*, 1994, Branco and Marrafa, 1999):

Principle A
A locally o-commanded short-distance reflexive must be locally o-bound.
Lee$_i$ thinks [*Max$_j$ saw himself$_{*i/j}$*].

Principle Z

[2] In the framework of centering theory, besides being correlated with the attentional status of the cognitive representation of their admissible antecedents, anaphors are correlated also with discourse and grammatical structure (vd. Walker *et al.*, 1998b). As this theory unfolds its predictions basically for clauses with single occurrences of pronouns in the clausal predication domain, and for pronouns in anaphoric links across adjacent sentences, no grouping of natural classes of anaphors of the kind we are concerned here is implied by those predictions.
For a discussion on how binding theory can be extended to phenomena at the discourse level see (Branco, 2000, p.289-294).

An o-commanded long-distance reflexive must be o-bound.
[*O amigo do Rui*$_i$]$_j$ *acha que o Pedro*$_k$ *gosta dele próprio*$_{*i/j/k}$. (Portuguese)
[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$_k$.'

Principle B
A pronoun must be locally o-free.
Lee$_i$ *thinks* [*Max*$_j$ *saw him*$_{i/*j}$].

Principle C
A non-pronoun 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 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.

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 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.[3]

3.3. Sets of admissible antecedents

As discussed above in Section 2.3, the search optimization rationale for anaphor resolution implies some predictions concerning the relations between the different natural classes of admissible antecedents. These are expected to be either disjoint — strong prediction —, or successively included within each other — weak prediction.

Given the natural binding classes just presented, we can now check if they conform to these predictions. For each different binding class we delimit the corresponding set of admissible antecedents and then check out how they relate to each other.[4] In order to proceed with this testing we just have to fix a position in a generic multiclausal grammatical structure and successively instantiate that position with an anaphor **x** from each different binding class. This way we will be able to collect the four sets of admissible antecedents and observe what are the relations among them.

Accordingly, if we assume that **x** is an A-anaphor complying with principle A, we see that its admissible antecedents form the set of the local o-commanders of **x**, which we can call the set A. In case **x** is a Z-anaphor, the set Z of its admissible antecedents is made of its o-commanders. When **x** is a B-anaphor, the set B of admissible antecedents contains all the antecedents that are not local o-commanders of **x**. Finally, the set C of the admissible antecedents of **x** when this is a C-anaphor has all the antecedents that are not o-commanders of **x**.

Summing up, the relations between these four sets 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 these sets in graphical terms,

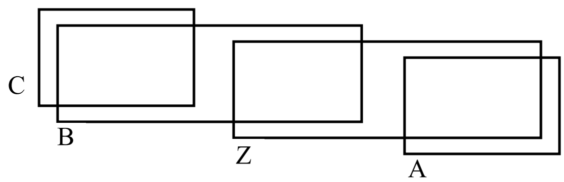

the relations among them can be rendered in a more perspicuous way by means of the following diagram:

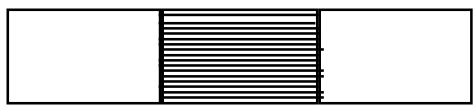

It is straightforward to see that the possible antecedents of short-distance reflexives are possible antecedents for long-distance reflexives; some possible

---

3 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).

4 As stressed above, under binding constraints, anaphors are correlated with the position in grammatical geometry of their admissible antecedents, not with some feature concerning their attentional status. This difference, however, does not imply that the search optimisation rationale could not be a suitable cognitive justification for binding classes, nor that these may not provide suitable empirical support to this rationale. It just imply that the search space where the search for antecedents is operative might be of a nature different from the one discussed in Section 2.

antecedents of long-distance reflexives are possible antecedents of pronouns; and the possible antecedents of non-pronouns are possible antecedents of pronouns.

From another, more general perspective, for a given possible antecedent, it is the case that there are always at least two different types of anaphors that can take it.[5] 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 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.

An important conclusion thus resulting from this discussion is that the search optimisation rationale does not seem to offer a justification of binding classes, or equivalently of binding constraints, of a more deeply rooted nature or in terms of a higher level of abstraction.[6]

4. Looking for other sort of cognitive underpinnings

Given the well established empirical grounding of binding constraints and their central role in the set of known filters for anaphor resolution,[7] the problems now uncovered cast non negligible doubts on the explicative and heuristic value of models for anaphora primarily based on search optimization in some cognitive space. This should not be taken, however, as implying that cognitively rooted factors (such as attentional prominence associated with recency of mention) do not play an important role in anaphor resolution, at least as preference mechanisms. Nor should it be seen as implying that binding constraints have been proved not to have any cognitive justification.

Instead, these negative results should be taken, in our view, as showing that cognitive underpinnings of binding classes seem not to be found in prima facie assumptions about cognitive life, but are perhaps more deeply entangled in other possibly non-trivial aspects of cognition.

4.1. The quantificational structure of binding constraints

This suggestion is supported by the results from some previous work of ours, which may provide interesting hints as to the possible, though more remote cognitive grounding of binding constraints.

Building on the existence of a fourth binding principle, for long-distance reflexives, we showed in (Branco and Marrafa, 1996, 1999) that the binding principles can be arranged into a square of logical oppositions.

This stimulated research on the eventual quantificational character of binding constraints. Adopting Löbner's duality criterion for quantification in natural language, and the formal tools he developed for the analysis of phase quantification, we showed in (Branco, 1998a,b, 2000) that the four binding constraints can be seen as the effect of four binding quantifiers, Q_A, Q_Z, Q_B and Q_C. These phase quantifiers are expressed by the nominals of the four binding classes, and quantify over the reference markers organized in the obliqueness order axis.

This proposal lends support to a radically new understanding of the formal nature of binding constraints, and to interesting explanations of some related issues such as the exemption occurrences and logophoric behavior of reflexives.

Giving a full-fledged account of the empirical support and justification as well of the implications of these results is clearly out the scope and the space constraints of the present paper. This does not hinder, nevertheless, a clear statement of the reasons behind our suggestion that these results might hint at a cognitive grounding of binding constraints: Given the quantificational structure underlying them, binding constraints are a manifestation of quantification, a universal semantic module of natural language, which is arguably rooted in some cognitive invariant.

4.2. Quantification, semantic universals and cognition

In order to lend relevant exploratory plausibility to this suggestion that the quantificational nature of binding may point towards a new type of cognitive grounding to binding classes, we should briefly review a couple of successive abstraction steps.

The first relevant step is the one that takes into account the abstraction involved in grouping together quite different linguistic phenomena into the single semantic class of quantification. Following Löbner's words, quantification involves "a seemingly very comprehensive

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

[6] van Hoek (1997) designs a cognitive grammar approach for binding constraints. In this approach, the basic assumption is that the shape of the conceptual structure associated with a given grammatical construction is isomorphic to the shape of its obliqueness hierarchy. Inasmuch as binding constraints are assumed to be somehow primitives of the binding system applying over this cognitive structure, this approach is not rooted in any search optimisation rationale. While this move makes it possible to provide the empirically correct predictions, in our view, it falls short of enhancing a principled cognition-based understanding of binding constraints, independently motivated from the binding facts proper.

[7] Since the so called integrative approach to anaphor resolution was set up in late eighties ((Carbonell and Brown, 1988), (Rich and LuperFoy, 1988), (Asher and Wada, 1988)), and its practical applicability extensively checked up, (cf. (Lappin and Leass, 1994), (Mitkov, 1997), (Mitkov, 1998) among others) it became common wisdom that factors for determining the antecedents of anaphors divide into filters or constraints, and preferences or heuristics. The first exclude impossible antecedents and help to determine the set of antecedent candidates; the latter help to pick the most likely candidate, that will be proposed as the antecedent. Binding constraints belong to the set of filters.

range of phenomena which are syntactically and grammatically rather diverse but semantically closely enough related to form a class of their own" (Löbner, 1987, p.53).

The fact that some natural language terms may express quantification had already been identified in the literature as the fact that they can express a one place second order predicate. It was however with Löbner that the common characteristic of natural language quantifiers was abstracted away and it was noticed that "duality is a fundamental concept in connection with quantification, but has been neglected almost completely in the relevant linguistic literature" (p.54).

The regular pattern of duality stressed by Löbner, and displayed below, had the heuristic value of permitting to unite within the same semantic analysis a set of apparently disparate phenomena.

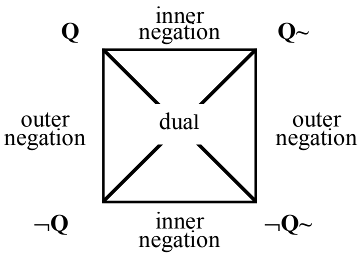

After the contributions of Löbner, as well as other works informed by this goal, several different groups of operators could be related to each other, and it was possible to discover that their semantics comply with the same very simple and abstract pattern. This duality-based perspective on the essence of natural language quantification made it possible to extend quantification beyond the well known cases of nominal quantification supported by the "classic" determiners *all*, *some*, *most*, *many*, etc., namely by covering the realms of temporality and possibility as well. Moreover, items such as *still*/*already*, and others (*enough*/*too*, scaling adjectives, *many*/*few*, etc.) although not lending themselves to straightforward analysis in terms of set quantification, can also be arranged in a square of duality:

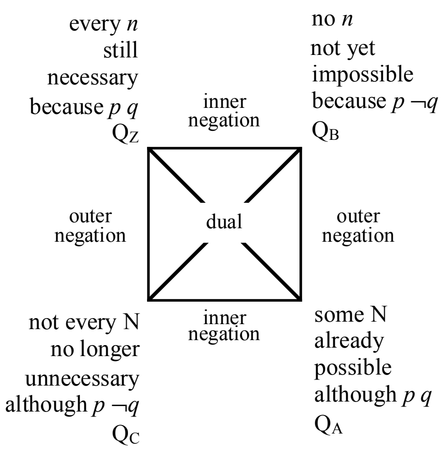

This heuristic value of duality was appropriately stressed in the context of a reflection on semantic universals by van Benthem (1991, p.23), who observed that the duality pattern "serves one further function [...]: it suggests a systematic viewpoint from which to search for comparative facts":

The second relevant abstraction step is to perform a new abstraction over the first one. In this case one should abstract from the different squares of duality involving natural language quantifiers, and notice the regular pattern of the relations between the corners of the square.

The different corners of duality squares, i.e. the different quantifiers, are related to each other by two third order operators, namely internal and external negation. The important point to note is that under an operation of function composition these operators give rise to an algebraic structure with very well known properties (– stands for identity, ¬ for external negation, ~ for internal negation, and # for ¬~):

∘	–	¬	~	#
–	–	¬	~	#
¬	¬	–	#	~
~	~	#	–	¬
#	#	~	¬	–

An algebraic structure of this type is known as a Klein four-group. It is a group (i.e. the relevant binary operation is associative and there is an identity element), which is commutative and has two generators (i.e. every element of the group is the result of iterating the operation over the generators).

Again, as happened with the abstraction relating to duality, this may not merely be a matter of vacuous formalization. This new abstraction degree may also have a critical heuristic impact. It may eventually allow one to recognize that natural language quantification, in general, and binding constraints, in particular, are just one manifestation of a broader class of phenomena, where other phenomena are found to pattern according to the same abstract structure.

Significantly, van Benthem signals that a similar structure was also found in cognitive psychology. In the fourth and last stage of children's cognitive development hypothesized by Piagetian development psychology (vd. Piatelli-Palmarini, 1979), typically attained by fourteen-year olds, the so called stage of formal operational thought[8] involved the maturing of the cognitive operations of identity, negation and converse. Apparently, there is a "Piaget's often repeated observation that [these operations] give rise to a Klein four-group" (van Benthem, 1986, p.206).

Naturally, this similarity of the structure underlying natural language quantification, in general, and binding constraints, in particular, on the one hand, and a certain class of cognitive operations, on the other, is too

[8] Vd. (Gray, 1990) for an overview of formal operational thought.

interesting for new research paths to pass unnoticed. According to a suggestion by van Benthem (1991), "the reason why this is a natural scheme is that further iteration serve no purpose" (p.31). Nevertheless, as this author himself notices in another paper, as interesting as this similarity certainly is, "nobody has ever been able to fit this into some significant theory" (van Benthem, 1986, p.206).

Be that as it may, the highly non-trivial parallelism between underlying abstract structures of so called cognitive formal operations and negation operations involved in linguistic quantification is promising enough, in the light of the results of the present discussion, for it to be seriously taken as an important stimulus for research into the possible cognitive underpinnings of binding constraints and of anaphor resolution

5. Conclusions

In this paper we sought to uncover the primitives underlying the mainstream paradigm of cognition-driven approaches to anaphor resolution. It turned out that a shared assumption, more or less explicitly embodied in the different approaches, is that there is a search optimization rationale behind eventual constraints (as opposed to preferences) on anaphor resolution.

This implies some predictions about the existence of natural classes of anaphors. In particular, it implies that the sets of admissible antecedents for each such natural class bear certain relations among them. These sets are predicted either to be disjoint, or at least to be successively included within each other.

Given such relations are not observed for the sets of admissible antecedents corresponding to binding classes, these natural classes are not offered any principled explanation by that rationale. Moreover, given the fact that these are the only objectively determined natural classes of anaphors, this result casts serious doubts that the search optimization rationale, though being a crucial factor in the set of preferences or heuristics for anaphor resolution, may provide a clear-cut justification for anaphor resolution and its constraints.

When looking at alternatives for some sort of cognitive explanation for binding constraints, recent results on their quantificational nature provide some interesting hints on a yet to explore line of research. Natural language quantification rests on a duality-based logical structure, which, in turn, instantiates a more general algebraic structure known as the Klein four-group. The exciting clue crops up when one notices that this is also the algebraic structure underlying the operations of formal thought, the fourth stage of children's cognitive development hypothesized in development psychology.

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.

van Benthem, J., 1986. *Essays in Logical Semantics*. Dordrecht: D. Reidel.

van Benthem, J., 1991. Linguistic Universals in Logical Semantics. In D. Zaefferrer (ed.), *Semantic Universals and Universal Semantics*. Berlin: Foris.

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. and P. Marrafa, 1999. Long-distance Reflexives and the Binding Square of Opposition. In G. Webelhuth, J. Koening, and A. Kathol (eds.), *Lexical and Constructional Aspects of Linguistics Explanation*. Stanford: CSLI Publications. Chap. 11, 163-177.

Branco, A., 1998a. The Logical Structure of Binding. In *Proceedings, 16th International Conference on Computational Linguistics and 36th Annual Meeting of the Association for Computational Linguistics (COLING-ACL98)*. Montreal: University of Montreal, 181-187.

Branco, A., 1998b, Towards a Lean Constraint-based Implementation of Binding Theory, in New Approaches to Discourse Anaphora: Proceedings of the Second Colloquium on Discourse, Anaphora and Anaphor Resolution. Lancaster: Lancaster University. UCREL Technical Papers, vol. 11, 54-67.

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

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*. Dordrecht: Foris.

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

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

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.

Gray, W., 1990. Formal Operational Thought. In W. Overton. (ed.), *Reasoning, Necessity and Logic: Developmental Perspectives*. Hillsdale: Lawrence Erlbaum, Chap.11, 227-253.

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.), pp.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.

van Hoeck, K., 1997. *Anaphora and Conceptual Structure*. Chicago: The University of Chicago Press.

Lappin, S. and H. Leass, 1994. An Algorithm for Pronominal Anaphora Resolution, *Computational Linguistics*, 20: 535-561.

Löbner, S., 1987. Quantification as a Major Module of Natural Language Semantics. In J. Groenendijk, D. de Jongh and M. Stokhof (eds.), *Studies in DRT and the Theory of Generalized Quantifiers*., Dordrecht: Foris, 53-85.

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.

Mitkov, R., 1998. Robust Pronoun Resolution with Limited Knowledge. In *Proceedings, 36th Annual Meeting of the Association of Computational Linguistics and 17th International Conference on Computational Linguistics (COLING-ACL98)*, 869-875.

Piatelli-Palmarini, M., (ed.), 1979. *Théorie du langage. Théorie de l'Apprentisage. Le débat entre Jean Piaget et Noam Chomsky*. Paris: Éditions du Seuil.

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

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

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 Chldren. *Cognition*, 62:245-290.

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

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), Stanford: CSLI Publications.