

Computation of modifier scope in NP by a language-neutral method

Richard CAMPBELL
Microsoft Research
1 Microsoft Way
Redmond, WA, USA 98052
richcamp@microsoft.com

Abstract

The relative logical scope of multiple modifiers within NP is often semantically significant. This paper proposes a structurally based method for computing the relative scope of such modifiers, based on their order, type, and syntactic complexity. The algorithm is language-neutral, in that it works with minimal errors for a wide range of languages without language-specific stipulations.

Introduction

Noun phrases quite commonly have multiple modifiers, including quantifiers, attributive adjective phrases, relative clauses, possessors, appositives and the like. As frequently noted in the literature (e.g., Shaw and Hatzivassiloglou, 1999), the linear order of modifiers can signify their logical scope (though other factors are involved, too), as in the English examples (1) and (2) (bracketing indicates logical scope):

- (1) my [favorite [new movie]]
- (2) my [new [favorite movie]]

In (1) *favorite* modifies the phrase *new movie*; hence the NP refers to my favorite among the new movies (there may be an old movie I like better); in (2) *new* modifies *favorite movie*; hence the NP refers to my favorite movie, which has just become my favorite.

The computation of the scope of modifiers is of inherent linguistic interest: it is necessary for determining the correct interpretation of NPs like (1) and (2). It follows that it is potentially useful in any application that may depend on such an interpretation. In addition, for multilingual

applications such as transfer-based machine translation (MT) (as discussed for example by Richardson *et al.* (2001)) modifier scope may itself be used as an abstract, language-neutral representation of their surface configuration, including linear order. The generation component of the MT application could then make use of scope information, perhaps in addition to scope-independent ordering conventions (Malouf, 2000), to generate the modifiers in the correct order.

The focus of the current paper is a method for computing the relative scope of modifiers based on structural information, which works independently of any particular language; that is, the same algorithm that computes the scope of the modifiers in the English NP (3) also correctly computes modifier scope in (4), its French translation, even though the two examples do not have exactly parallel surface structures:

- (3) the [twenty-ninth [American state]]
 - (4) le [vingt-neuvième [État américain]]
- the twenty-ninth state American

The proposed algorithm considers several structural factors in addition to linear order, including the type and internal structure of the modifiers themselves. The algorithm described here is currently implemented in the NLPWin system at Microsoft Research (Heidorn, 2000).

The paper is organized as follows: Section 1 examines the various structural factors that determine modifier scope, and a preliminary algorithm for modifier scope assignment is proposed; in Section 2, we compare the predictions of the algorithm to a diverse set of examples from six languages, and propose a revised algorithm; Section 3 considers some related work; and Section 4 is a conclusion.

1 Modifier scope

English examples like (1) and (2) seem to show that linear order is a principal factor in determining the scope of modifiers. Examination of a wider range of examples from a variety of languages makes it clear that matters are not that simple. In this section, we explore some ways in which strict linear order is not sufficient to determine the scope of modifiers.

1.1 Order of postnominal modifiers

Strictly speaking, linear order plays only an indirect role in scope assignment, the relevant ordering factor being distance from the head noun: all else being equal, modifiers that are farther from the head noun in the surface structure have scope over modifiers that are closer to the head.

Linear order plays a different role in determining scope, depending on whether the modifiers precede or follow the head noun; this is illustrated by the Spanish examples (5) and (6):

- (5) una [[moneda americana] falsa]
a coin American counterfeit
'a counterfeit American coin'
- (6) una [[moneda falsa] americana]
'an American counterfeit coin'

The NPs in (5) and (6) differ only in modifier order and indeed this reflects their relative scope; but unlike the prenominal modifiers in (1) and (2), the postnominal modifiers in (5) and (6) have a scope order that is the reverse of their linear order: (5) refers to an American coin that is counterfeit, e.g. a fake half-dollar, perhaps produced in Canada; (6) refers to a counterfeit coin from America, e.g. a fake Canadian dollar produced in the U.S.

The difference in the way that linear order of pre- and postnominal modifiers determines scope makes it clear that in a multilingual application such as MT, it is not sufficient merely to record the order of modifiers, since the order must be reversed going e.g. from Spanish to English. A more straightforward way to record the information stored in the linear order of modifiers in the source language would be to record their relative distance from the head noun. Given the frequent occurrence in many languages of NPs

with both pre- and postnominal modifiers, as in (4), above, this can only be accomplished by recording the logical scope of the modifiers.

1.2 Quantifier-like adjectives

Relative distance from the head is only obvious in cases of multiple prenominal or multiple postnominal modifiers. In many languages, however, it is relatively common for an NP to contain both pre- and postnominal modifiers, as in (4). In such a case, the notion of relative distance from the head is unavailable as a guide to relative scope. The following examples illustrate the problem:

- (7) the [heaviest [isotope found in nature]]
- (8) a [[new domain] devoted to insects and worms]

The prenominal adjective in (7) has wider scope than the postnominal participial phrase; thus the NP refers to the heaviest member of the set of isotopes of some element that are found in nature, and not e.g. to the heaviest isotope overall of that element. In (8) on the other hand, the postnominal participial phrase has wider scope than the prenominal adjective: the NP refers to a new domain or classification (invertebrates), which is devoted to insects and worms, and not, e.g., to a new member of the set of domains so devoted.

The relevant difference between these examples is the prenominal modifier: superlatives, along with comparatives, ordinals, certain other adjectives that are quantifier- or determiner-like, such as *another*, *certain*, *numerous*, and *only* (all the preceding are henceforth referred to as 'q-like adjectives'), along with quantifiers take wider scope than most postnominal modifiers, while everyday modifiers like *new* generally take narrower scope when prenominal. Assuming that the class of q-like adjectives can be identified, a scope assignment algorithm needs to take this into account as well.

1.3 Nonrestrictive postnominal modifiers

Nonrestrictive modifiers generally take very wide scope, typically having wider scope than any restrictive modifiers or quantifiers; compare (7), above, with (9):

- (9) the [[heaviest isotope], which is found in nature]

Unlike (7), this NP does refer to the heaviest isotope (of some element), and not just to the heaviest isotope found in nature, because the nonrestrictive relative clause has wider scope than the prenominal adjective.

In English and other languages, nonrestrictive postnominal relative clauses, adjective phrases, and participial clauses are easily identifiable by the preceding comma or other structural cues. Whether a modifier is restrictive or nonrestrictive is clearly relevant to the computation of its scope relative to other modifiers; however in some languages, notably German and Japanese, there are often no reliable structural cues to whether a relative clause is restrictive or not. This part of the scope computation algorithm can therefore only be expected to work in languages where this information is available in the input.

1.4 Competing principles of scope assignment

In most cases, the various principles of scope assignment outlined in this section are not in conflict. For instance, q-like adjectives, when prenominal, typically precede other prenominal adjectives. In some cases this is not true, however, and in many such cases, relative distance from the head noun is a principle of last resort only. A good illustration of this fact comes from Japanese, where all relative clauses are prenominal, and often precede other prenominal modifiers, such as adjective phrases and quantifiers; this is illustrated by the following examples:

- (10) 被害にあった主要な都市
higai-ni a-tta shuyou-na toshi
 damage-DAT encountered major-ADN cities
 'major cities that were damaged'
- (11) 魚が食べる あらゆる餌
sakana-ga tabe-ru arayuru esa
 fish-NOM eat-PRES all bait
 'all bait which fish eat'

In (10) the relative clause has wider scope than the following adjective, but in (11) the quantifier

arayuru 'all' has wider scope than the preceding relative clause. The principle that such adjectives and quantifiers are assigned wider scope takes precedence over the principle that assigns scope on the basis of relative distance from the head noun.

1.5 Algorithm for computing modifier scope – First pass

Based on these observations, a simple, language-neutral algorithm can be formulated to compute modifier scope based on structural factors. As a first step, we factor all modifiers into three categories: nonrestrictive modifiers, quantifiers and q-like adjectives, and other modifiers. For practical purposes, nonrestrictive modifiers are limited to postnominal relative clauses, adjective phrases, and participial clauses, that have some structural indication of their nonrestrictiveness, such as being preceded by a comma; in principle, however, any nonrestrictive modifier should fall into this category. Q-like adjectives include comparatives, superlatives, ordinals, and modifiers (e.g. *only*) that are marked in the dictionary as being able to occur before a determiner. Also, if a q-like adjective is prenominal, then any other adjective that precedes it is treated as if it were q-like; if the q-like adjective is postnominal, then any other adjective that follows it is treated as if q-like. In this paper, PPs, possessors and appositive NPs are not treated. The algorithm is described in (I):

- I. Computation of modifier scope
 1. nonrestrictive modifiers have wider scope than all other groups;
 2. quantifiers and q-like adjectives have wider scope than other modifiers not covered in (I.1);
 3. within each group, assign wider scope to postnominal modifiers over prenominal modifiers;
 4. among postnominal modifiers in the same group, or among prenominal modifiers in the same group, assign wider scope to modifiers farther from the head noun.

Consider the NP in (12), which, though awkward, will serve as illustration of how (I) works:

- (12) [[two [other [[counterfeit [American coins]] produced here]]], which I saved]

The nonrestrictive (set off by commas) relative clause *which I saved* is assigned widest scope by (I.1); the quantifier *two* and the q-like adjective *other* have wider scope than the remaining modifiers (I.2), and *two* has wider scope than *other* (I.4); the participial clause *produced here* has wider scope than *counterfeit* and *American* (I.3); finally, *counterfeit* has wider scope than *American* (I.4).

For a more realistic illustration, consider again (3) and (4), repeated here:

- (3) the [twenty-ninth [American state]]
 (4) le [vingt-neuvième [État américain]]
 the twenty-ninth state American

In both (3) and (4), the ordinal is assigned wider scope than *American/américain*, regardless of the latter's position in the NP, by principle (I.2). Finally, consider again (5), repeated here:

- (5) una [[moneda americana] falsa]
 a coin American false

By (I.4), *falsa* is assigned wider scope than *americana*.

2 Examination of broader range of examples

The algorithm in (I) is motivated by a small set of examples; a broader range of examples is needed to really determine whether (I), and especially its claim to language-neutrality, is viable. The purpose is not to provide a statistical measure of the accuracy of the algorithm, but simply to provide a set of examples that is larger, more diverse and more realistic than could have been devised by introspection.

2.1 Testing (I)

Initial sets of examples from six languages, Chinese, English, French, German, Japanese and Spanish, were collected; the Chinese examples were taken from *Du Zhe* ('Reader'), for all other languages, the examples were taken from *Encarta*

Encyclopedia. To be selected, sentences had to contain an NP with exactly one prenominal and one postnominal modifier (thus (I.4) was not directly tested). Several hundred such examples were then given to teams of two native speaker linguists per language to determine (a) whether it made any difference which modifier was assigned wider scope, and (b) if so, which one had wider scope; for an NP to make it into the final set of examples, both annotators had to agree on both (a) and (b). Annotators aimed for approximately 100 examples for which the answer to (a) was positive; in practice, between 50 and 100 examples for each language made it into the final set.

The annotated examples were then analyzed using the NLPWin system, which incorporates this algorithm, the results manually compared to the annotations, and differences (errors) examined. In some cases, errors could be attributed to misanalysis or other deficiencies of the system, independently of the scope assignment algorithm; these might include syntactic errors (e.g. *English court favorite* analyzed with *court* as the head noun and *favorite* as a postnominal adjective), or lexical/morphological problems (e.g. Sp. *numeroso* 'numerous' not recognized as a quantifier, failure to treat Fr. *Premier ministre* as a single noun, etc.). Overall, the results were quite good for English (3 errors in 97 examples, including system errors), German (6/81) and Chinese (2/66), mediocre for Spanish (13/60) and Japanese (15/59), and poor for French (46/78).

In addition to system errors, certain patterns emerged. The most obvious one was that, in many languages, prenominal adjectives meaning 'same' or 'different', which should take wider scope than a nonrestrictive postnominal modifier, were assigned narrower scope according to (I.3); a good example from Spanish is given below:

- (13) el [mismo [año que contrajo matrimonio con Carmen Polo]]
 'the same year that he contracted marriage with Carmen Polo'

Adjectives meaning 'same', 'different' and 'other' have much in common with comparatives; for example, English *same*, *different*, and *other* can take *as-* or *than-*complements. Our system does not currently encode this similarity in any

systematic way, but if it did so, examples like (13) would be handled correctly by (I.2).¹

2.2 Syntactic complexity

An unexpected pattern among the errors, however, had to do with the internal structure of postnominal modifiers. The French examples (14) and (15) have the same prenominal adjective:

- (14) de [nouvelles [valeurs culturelles]]
new values cultural
'new cultural values'
(15) un [[nouveau domaine] consacré aux
insectes et aux vers]
'a new domain devoted to insects and
worms'

The relevant difference in this case is in the postnominal modifier: in (14), the postnominal adjective is syntactically simple, containing no syntactic dependents of its own, while in (15) the postnominal participial clause is complex, containing a complement prepositional phrase. (I.3) incorrectly assigns wider scope to the postnominal modifier than to the prenominal modifier in (14); an unexpectedly large proportion of the French and Spanish errors from this example set were of this type.

This leads to a revision of the scope assignment algorithm that treats syntactically simple (unmodified) postnominal modifiers as a special case, getting assigned narrower scope than regular prenominal modifiers:

II. Computation of modifier scope, revised

1. nonrestrictive modifiers have wider scope than all other groups;
2. quantifiers and q-like adjectives have wider scope than other modifiers not covered in (II.1);
3. syntactically complex postnominal modifiers that are not relative clauses have

4. wider scope than other modifiers not covered by (II.1-2);
5. prenominal modifiers not covered by (II.1-3) have wider scope than other modifiers not covered by (II.1-3);
6. otherwise, within each group, assign wider scope to postnominal modifiers over prenominal modifiers;
7. among postnominal modifiers in the same group, or among prenominal modifiers in the same group, assign wider scope to modifiers farther from the head noun.

The difference between (I) and (II) is in (II.3) and (II.4), which ensure that syntactically complex postnominal modifiers have wider scope than non-quantificational prenominal ones, and that prenominal modifiers have wider scope than syntactically simple postnominal ones. In this case, "syntactically complex" means (a) if not a coordinate structure, then there are non-head constituents; and (b) if a coordinate structure, then at least one of the conjuncts is syntactically complex.

Implementing the revised algorithm (II) into the system reduced the number of French and Spanish errors in the example set considerably: French went from 46 errors in 78 examples to 2; Spanish went from 13/60 to 5; all other languages remained essentially the same.

To ensure that the revision from (I) to (II) was not tailored specifically to the example set, a second set of randomly selected examples were annotated as before; the examples were then analyzed using the system incorporating (II), and the results compared to the annotation.

2.3 Discussion

The results bear out the essential correctness of (II), while at the same time highlighting areas in which further refinement is possible. For each language there were just a handful of errors in this example set: Chinese had 3 in 57 total examples, English 10/98, French 4/55, German 4/45, Spanish 8/96 and Japanese 8/55. Since the test is not meant to be a statistical measure of accuracy, it is important to examine the errors. Two English and two Spanish errors were system errors of the kind described in Section 2.1. The problem of identifying nonrestrictive relative clauses arose as well,

¹ The special status of such adjectives has been noted in other contexts: For example, Hawkins (1978) groups *same* together with superlatives into a class of 'unexplanatory' modifiers; Vieira and Poesio (2000), extending this class to include *only* and a few others, make use of them in identifying discourse-new definite descriptions.

accounting for one English and one German error, and possibly all the Japanese errors.

Of the remainder, both the French and Spanish sets contained errors which suggest simple modifications to (II), e.g. where both the prenominal and postnominal modifiers were q-like, and for which (II.5) incorrectly assigns wider scope to the postnominal. Another error that could be handled by a simple refinement of (II) is Sp. *los grandes rebaños de ovejas no estabulados*, ‘the large unstabaled herds of sheep’; in this case the postnominal adjective phrase *no estabulados* is treated as syntactically complex, and incorrectly assigned scope wider than *grandes* by (II.3); it seems simple enough to modify (II.3) to take account of negative morphemes and the like.

Another category of error that suggests a further refinement is suggested by the Spanish example (16) and by English (17):

- (16) una [[quinta cláusula] que no tuvo efecto]
a fifth clause that not took effect
- (17) [[longer poems] written from 1789 on]

In both cases, (II.2) incorrectly assigns wider scope to the prenominal adjective than to the postnominal one. Examples such as these suggest that q-like adjectives, or at least comparatives and ordinals, should be treated as in (II.2) (i.e., as taking wider scope than other modifiers) only in NPs that are definite. One Spanish, one English, and two German examples of this kind occurred in this example set; since there are not many errors of this kind, it is not clear how much would be gained from this modification.

Aside from such systematic errors, the French set contained only one error that is irredeemable; i.e., which (II) could not handle even with perfect input, without being supplemented by lexically specific, and hence language-specific, rules; the Spanish set contained only three irredeemable errors; English had three, and German one.

While it is worth noting that (II) is not without counterexamples, it is significant that true counterexamples are evidently rare enough in actually occurring text, at least when compared to examples for which (II) predicts the correct scope assignment, that (II) appears to be very promising as a means for computing modifier scope in arbitrary languages.

2.4 Why complexity?

It is not clear exactly why the syntactic complexity of a postnominal modifier affects its scope relative to other modifiers in NP. It may be significant that postnominal modifiers that are themselves modified tend to be participial or relative clauses, rather than adjective phrases. Participial and relative clauses are always intersecting in their interpretation, meaning that the denotation of the noun + modifier construction is the intersection of the set denoted by the noun and the set denoted by the modifier; adjectives, on the other hand, are often non-intersecting (Keenan and Faltz, 1985). It is possible then that a deeper principle underlies (II.3), namely that intersecting modifiers take wider scope than non-intersecting ones.

It remains to be explained why prenominal adjectives typically take wider scope than syntactically simple postnominal ones. One possible explanation is that N + unmodified Adj combinations, such as Fr. *valeurs culturelles* in (14), are often analyzed by native speakers as a kind of compound; i.e., as though the Adj were incorporated into the N to form a complex word.

Typically the parts of a complex word cannot be individually syntactically modified. It is to be expected, then, that a prenominal modifier such as *nouvelles* in (14) would be unable to modify *valeurs* by itself, but must modify the whole compound. Moreover, it is to be expected that the adjective in the N + Adj compound could not have modifiers. Consequently, a noun + [adjective + PP] construction, such as Fr. *huiles originaires des régions méditerranéennes* ‘oils originating from mediterranean regions’, could not be analyzed as a compound; instead the postnominal adjective phrase would have to be a true phrasal modifier, taking wider scope than a prenominal modifier such as *célèbres* ‘famous’, according to the general principles in (II).

This account is purely speculative, of course, and thus far untested. Other kinds of explanation for (II.3) are possible as well, but limitations of space preclude substantial discussion of this issue.

3 Related work

Copestake *et al.* (1995) briefly address the issue of adjective scope in an NP such as *a fierce black cat*. Since both adjectives are intersecting in this case,

their relative scope is semantically irrelevant (*a black fierce cat*, though infelicitous, would mean the same thing). Since a translation of this NP might have a different structure, e.g. Spanish *gato feroz y negro* lit. ‘cat fierce and black’, Copestake *et al.* argue that the logical form (LF) for such an NP ought to be flattened (i.e., no scope assigned to the adjectives), so as to ensure that the NP and its translation do not have syntactically different LFs, which is required for system-internal reasons.

If this conclusion is justified, it poses a problem for the approach taken in this paper, since (II) assigns wider scope to *fierce*. Most dictionaries do not mark adjectives according to whether they are intersecting or not, nor do we know of any large corpora that are annotated in this way; therefore, it seems unlikely that (II) could be reliably turned off just in those cases where scope assignment is logically unnecessary.

Copestake *et al.*’s conclusion is based on the system-internal assumption that transfer-based MT requires the LF of the input (e.g. *gato feroz y negro*) to have the same syntactic structure as the LF of the output (e.g. *fierce black cat*). However, a transfer-based MT system in which transfer rules are learned from aligned corpora, such as described by Richardson *et al.* (2001), does not have this requirement; hence the problem Copestake *et al.* discuss does not arise.

4 Conclusion

The fact that (II) works so well across a variety of languages is of inherent linguistic interest, as it suggests that the cross-linguistic variation in word order within NP, while considerable, is nevertheless subject to universal principles that enable the relative scope of modifiers to be recovered. These principles take account primarily of the type of modifier, but also of their placement relative to one another and, at least in the case of postnominal restrictive modifiers, their internal structure, assigning wider scope to modifiers that are themselves modified, perhaps because they are intersecting. Application of the algorithm to more languages would of course be required to fully substantiate this claim.

The success of (II) is of substantial practical interest, as well, since it does not make unreasonable demands on the. For example, although (II.3) may ultimately derive from a

deeper principle that assigns wider scope to intersecting modifiers, it is not necessary to identify modifiers as (non-)intersecting for (II.3) to work, since syntactic complexity works well enough. To work correctly, (II) requires quantifiers to be distinguished from adjectives, adjectives to be identified as superlative, comparative, ordinal or as able to occur before a determiner, and postnominal modifiers to be marked as nonrestrictive. The first two requirements are reasonable things to expect of any parser; the third requirement is not easily met in all languages, but even in those languages where nonrestrictives are not easily identifiable, (II) works reasonable well.

Acknowledgements

I would like to thank my colleagues in the NLP group at MSR, especially Mike Carlson and Hisami Suzuki, for their help.

References

- Copestake A., Flickinger D., Malouf R., Riehmann S. and Sag I. (1995) *Translation using Minimal Recursion Semantics*. In “Proceedings of TMI 6,” Leuven, Belgium.
- Hawkins J. (1978) *Definiteness and Indefiniteness*. Croom Helm, London, 316 p.
- Heidorn G.E. (2000) *Intelligent Writing Assistance*. In “Handbook of Natural Language Processing”, R. Dale et al., ed., Marcel Dekker, New York, pp. 181-207.
- Keenan E.L. and Faltz L.M. (1985) *Boolean Semantics for Natural Language*. D. Reidel Publishing Co., Dordrecht, 387 p.
- Malouf R. (2000) *The order of prenominal adjectives in natural language generation*. In “Proceedings of the 38th Annual Meeting of the ACL,” pp. 85-92.
- Richardson S., Dolan W., Menezes A. and Pinkham J. (2001) *Achieving commercial-quality translation with example-based methods*. In “Proceedings of the VIIIth MT Summit”, Santiago de Compostela, Spain. 293-298.
- Shaw J. and Hatzivassiloglou V. (1999) *Ordering among premodifiers*. In “Proceedings of the 37th Annual Meeting of the ACL,” pp. 135-143.
- Vieira R. and Poesio M. (2000) *An Empirically Based System for Processing Definite Descriptions*. Computational Linguistics 26/4, pp. 539-593.