

Linearization of complex modifiers: Ways of (dis)obeying the Head-Final Filter

June 18, 2021

Abstract

In this paper, we argue that the Head-Final Filter (HFF) in its current formulation is not only left theoretically unexplained, but is also empirically incorrect. For one, languages like Greek or Russian are not subject to it and can have material intervene between prenominal adjectives and modified nouns. In addition, a number of languages, such as Basque, display mirror-HFF effects. These languages bar material intervening between postnominal adjectives and modified nouns, something that is not captured by the HFF.

In place of the HFF, we propose a novel descriptive generalization that captures all the facts above: the Modifier-Noun Adjacency Generalization (MAG). We show that all MAG-related facts follow once it is assumed that direct modification of N by A is impossible and must be mediated through a functional element that is specified for all nominal case and φ -features. Such a functional head may already be part of the extended AP. This is the case when all relevant case and φ -features are present on the predicatively used adjective. Then, there is nothing that would forbid the adjective to select material to its immediate right. If this functional head is part of the extended NP, however, its morpho-phonological properties determine whether it must be adjacent to both the noun and the adjective or only to the noun. In the former case, no material may intervene between A and N; in the latter case, it may. This, we argue, derives MAG.

Key words: Head-Final Filter, adjectives, (nominal) modification, linearization, adjacency, case/ φ -features, rich agreement, affixes vs. clitics

1 Introduction: the Head-Final Filter

1.1 The HFF in English

The earliest formulation, known to us, of the constraint that will later become widely known as the Head-Final Filter is Greenberg (1963)'s Universal 21. This universal describes the following asymmetry in the linearization of nouns, attributive adjectives and their modifiers:

- (1) **Universal 21** (Greenberg, 1963)

If some or all adverbs follow the adjective they modify, then the language is one in which the qualifying adjective follows the noun and the verb precedes the object as its dominant order.

In this universal the asymmetry with respect to the prenominal/postnominal placement of attributive APs is described as a function of their structural complexity and the position of the head. While postnominal APs can be structurally complex independently of the position of their head, prenominal APs can only be structurally complex if their head is adjacent to the modified noun, i.e., if no further adjectival adverbs intervene.¹ This is shown for English in (2)-(3). The distribution of the patterns in Greenberg's sample of 30 languages that gave rise to this generalization is shown in the table below (the data for four languages was not available).²

	A-N	N-A
Adv-A	11	5
A-Adv	0	8
both possible	0	2

Table 1: Distribution of languages with respect to Universal 21 (Greenberg, 1963).

- (2) a. a [smoothly running] meeting (Sadler & Arnold, 1994)
b. *a [running smoothly] meeting
- (3) a. an [extremely original] idea
b. *an [original extremely] idea

¹Greenberg does not mention whether or not the generalization described in Universal 21 is based on some specific semantic class of adverbs.

²Throughout the paper, we use '–' to indicate linear precedence, i.e. 'X–Y' reads as 'X precedes Y'.

The domain of Greenberg’s Universal 21 is relatively narrow: it includes only nouns as modifiees, only adjectives as modifiers, and only adverbs as modifier dependants. However, the relevant generalization does not have to be restricted to these categories. As far as the type of the dependants is concerned, not only adjuncts, like adverbs, but also complements appear to be banned from intervening between prenominal adjectives and their head nouns. Williams (1982), for instance, provides the example with the PP *of his children*, which is a complement of *proud* and not an adjunct.

(4) *the [_{AP} proud of his children] man (Williams, 1982)

Moreover, the syntactic category of the modifier appears to be irrelevant as well. Not only modified adnominal APs, but also PPs and CPs (i.e., relative clauses) cannot appear prenominally, as the examples below show.

(5) *the [_{PP} to Bill] letter (Williams, 1982)

(6) *a [_{CP} which I published in 1991] book (Escribano, 2004)

Unlike adjectives, P and C heads always take a complement. Being head-initial, there will always be some phrase intervening between them and the nouns they modify when they appear prenominally. The impossibility for PPs and CPs to occur prenominally in a language like English is thus even more categorical compared to APs.

In view of these facts, Williams (1982) offers the following formulation of the constraint that applies in the examples above, the *Head-Final Filter* (HFF): an adjacency requirement between the head of a prenominal modifier and the head of the noun phrase that it modifies phrased in terms of the modifier’s head-finality.³

(7) **Head-Final Filter** (Williams, 1982)

English (and German) have a constraint barring post-head material in prenominal modifiers.

Unlike Universal 21, Williams’ HFF is not intended as a typological generalization. It only makes a claim only about English and German, but the claim that it makes for these two languages is actually stronger than Universal 21. Subsequent work has argued that the HFF is possibly even more general than the way it is formulated in Williams (1982), because it also holds for modifiers of verbs and adjectives (Escribano, 2004) as well as for degree constructions (Grosu & Horvath, 2006), and has been taken to be even a universal constraint

³For a predecessor of Williams’ HFF, see Emonds (1976)’s *Surface Recursion Restriction*.

(Philip, 2012; Sheehan, 2017). Within the scope of this paper, we do not concern ourselves with the question of whether the linearization restrictions observed for degree constructions and VP/AP-modifiers can or should be subsumed under the cover of the HFF. However, we will follow the commonly shared view that Williams’ HFF is not restricted to English and German, but is indeed a linguistic (near-)universal. Hence, for the purposes of this paper, we will adopt the following operational definition of the HFF:

(8) **Head-Final Filter** (operational definition)

Prenominal modifiers must not contain any post-head material (either being head-final in general or not taking any dependents to the right).

Below we show that the HFF indeed applies to many languages. We also discuss various “rescue strategies” that languages employ when speakers convey something that would otherwise have to be expressed by means of a non-head-final prenominal modifier. However, before addressing these issues, it should be pointed out that even in a HFF-obeying language like English, the HFF does not apply without exceptions, at least at first sight. At least three apparent English HFF-violating types of constructions can be identified.

The first type of construction concerns right-branching modifiers of the kind exemplified in (9). Such modifiers are, however, lexical formations, as is evidenced by the absence of agreement morphology (cf. **a two-syllables word*), and as such do not constitute real violations of the HFF. Being constructed in the lexicon, they are treated in syntax as single heads.

- (9) a. a [_{QP} two-syllable] (phonological) word
 b. a [_{AP} higher-than-average] (basic) salary
 c. an [_{PP} up-to-date] (linguistic) bibliography (Escribano, 2004)

The same holds for another set of apparent counterexamples to the HFF, namely complex modifiers headed by *tough*-adjectives. Such modifiers can indeed appear in HFF-violating configurations (see (10)), but, as Nanni (1980) points out, they should be considered atomic lexical formations since they cannot contain adverbial modifiers (see (11)), overt experiencers, parasitic gaps or multiple embeddings.

- (10) a. an easy to clean room
 b. a hard to find manuscript (Nanni, 1980)

- (11) a. *an easy to *quickly* clean room
 b. *a hard to find *in the attic* manuscript (Nanni, 1980)

The final class of counterexamples concerns the degree modifier *enough*, which, unlike other degree modifiers in English (like *very* or *too*) occurs after and not before the adjective:

- (12) John is a tall enough guy to play basketball. (Sheehan, 2017)

Note, however, that here *enough* must immediately precede the modified noun in such cases. This evidenced by the fact that a *to*-infinitive cannot intervene and must be placed postnominally instead:

- (13) *John is a [tall enough to play basketball] guy. (Sheehan, 2017)

Unlike the cases with other, preadjectival degree modifiers, the combination of an adjective and *enough* should thus not be considered a matter of lexical formation. This type of construction thus presents a genuine violation of the HFF in English. This means that (8) should be amended such that degree modifiers can be exempt from the HFF if they immediately precede the modified noun. This *enough*-construction will play an important role in our analysis later on.

1.2 The HFF cross-linguistically

We conclude this overview of the domain of application of the HFF with a discussion of its cross-linguistic status. According to the current state of knowledge based on reports in the literature (cf. Cinque, 2010; Sheehan, 2017, for overviews) and our own data collection, the HFF seems to apply in at least the following languages: Dutch, English, German, Icelandic, Swedish (Germanic); French, Italian, Portuguese, Romanian, Spanish (Romance); Czech, Serbo-Croatian, Slovenian (Slavic); Armenian (Indo-European isolate); Estonian, Finnish, Hungarian (Uralic); and Georgian (Kartvelian).

Note that the languages listed above all obey the HFF in a non-trivial manner: all of them are head-initial languages, where A precedes N and where the linear order A–PP in the predicative position is available at least in addition to PP–A; yet the HFF-violating linear order A–PP–N is not possible when adjectives are used attributively. By contrast, in languages with strictly head-final APs preceding their modified nouns (Japanese and Turkish are of this type, cf. Sheehan, 2017), the HFF is obeyed trivially: the A–PP–N order is simply ruled out because adjectives in these languages cannot take their dependents to their right in the first place.

In order to linearize complex attributive APs containing adjectival dependents, non-trivially HFF-obeying languages must make use of various “rescue strategies”, i.e. alternative available word orders that are resorted to instead of the violating A–PP–N order. These strategies include the following three linearization patterns:

- (14) a. PP–A–N (left-branching prenominal AP)
 b. N–A–PP (postnominal AP)
 c. A–N–PP (discontinuous AP)

The possibility of using left-branching prenominal APs is the most commonly attested rescuing strategy. It is available in all the above-mentioned languages except for English and the Romance languages.⁴ Representative examples from Dutch and Hungarian are given below.⁵

- (15) de op zijn vrouw trotse man [Dutch]
 the of his wife proud man
 ‘the man proud of his wife’

- (16) egy a szülei-től független fiú [Hungarian]
 a the parents-POSS.3.SG-from independent boy
 ‘a boy independent of his parents’

Another rescue strategy consists in placing a complex AP in postnominal position, instead of in its regular prenominal position when simplex. This option, arguably involving instances of reduced relative clauses, is less common in our sample, but is still available in a number of languages, mainly in Romance and Slavic ones but also in English and Icelandic. The examples below illustrate this for English, Italian⁶ and Serbo-Croatian.

- (17) a. a person loyal to his friends
 b. a loyal person
 c. *a person loyal

⁴Note that PP–A orders are ruled out for predicatively used adjectives in these languages, though see Cinque (2010) for PP–A–N orders in literary Italian.

⁵We use the following abbreviations in glosses: ABL ablative, ART article, ATTR attributivizer, AV actor voice, CL classifier, COM common (gender), DAT dative, GEN genitive, INSTR instrumental, LNK linker, M masculine, NEUT neuter, NOM nominative, PL plural, PROG progressive, REF referential, SG singular, SUP superlative.

⁶Note that this only holds for a restricted class of prenominal adjectives in Italian; the majority of adjectives in Italian appear postnominally.

- (18) a. una ragazza brava in matematica [Italian]
 a girl good in maths
 ‘a girl good in maths’
 b. una brava ragazza
 c. *una ragazza brava
- (19) a. očevi ponosni na svoje sinove [Serbo-Croatian]
 fathers proud on their sons
 ‘fathers proud of their sons’
 b. ponosni očevi
 c. *očevi ponosni

The third attested rescuing strategy involves discontinuous APs. Here, the adjective itself remains in the prenominal position, but its dependent is placed postnominally. This possibility is not particularly common, but is available, for instance, in English (though not across the board) and Serbo-Croatian (for some speakers), as shown in the following examples.

- (20) a. a hard man (for me) to talk to (Berman, 1973)
 b. a similar car to mine (Escribano, 2005)
 c. a smart kid for her age (Bernstein, 1995)
 d. a more intelligent student than Bill (Hoekstra, 1999)
- (21) ponosan otac na svog sina [Serbo-Croatian]
 proud father on his son
 ‘a father proud of his son’

Finally, it should be noted that languages may exploit more than one rescuing strategy, as is evident from the examples in (19) and (21). These show that Serbo-Croatian can resort to both postnominal and discontinuous APs instead of the HFF-violating A-PP-N order.⁷

⁷Note that left-branching prenominal APs (the PP-A-N order) are also possible in Serbo-Croatian.

2 Challenges for the HFF

In the previous section, we have reviewed data from various languages that suggests that the linearization of prenominal modifiers is subject to the HFF. This, then, requires these modifiers to be head-final. At the same time, the HFF defined as in (8) itself cannot be taken as a principle that bans XP-intervention between nouns and their attributive adjectives throughout. The reasons for this are twofold. First, there are languages in which the HFF-violating A–PP–N order is in fact allowed, which means that the HFF overgenerates. Second, there are languages where postnominal modifiers, which fall outside the domain of application of the HFF, are subject to an analogous restriction that requires them to be head-initial and disallows the mirrored N–PP–A order. In this respect, the HFF undergenerates. We will discuss each of these types of languages in turn.

2.1 HFF-Violating languages

It is well-known in the literature that, despite the cross-linguistic tendency for languages to obey the HFF, some languages do in fact permit HFF-violations. On the one hand, this is the case for a number of languages that have affixal agreement marking on adjectives, such as Bulgarian, Polish, Russian, Ukrainian (Slavic); Latin, Old Romanian (Italic); Lithuanian (Baltic); Modern Greek (Hellenic). In these languages, where APs are generally prenominal and head-initial, the A–PP–N order is actually grammatical and well-attested, as demonstrated by the following corpus examples for Greek and Russian.⁸

(22) **Modern Greek** (Corpus of Modern Greek, <http://www.web-corpora.net/GreekCorpus>)

Afto to [simantiko gia tis arxes] zitima tha paramini ipo sizitisi.
this the important for the authorities issue will remain under discussion

‘This important issue for the authorities will remain under discussion.’

⁸It has been claimed in the literature that the possibility of the A–PP–N order in languages like Greek and Russian is only marginal (cf., e.g., Sheehan, 2017). However, the results of our consultations with native speakers and of the corpus studies we have conducted do not support this conclusion.

- (23) **Russian** (Russian National Corpus, <http://www.ruscorpora.ru/en>)

[Zavisimyje ot eksporta nefti i gaza] strany razrabatyvajut različnyje
dependent from export oil and gas countries develop various
programmy dejstvij.
programs actions

‘Countries that are dependent on oil and gas export develop various programs of actions.’

On the other hand, HFF-violations are also attested in languages in which agreement markers on adjectives are non-affixal, i.e. where they are free words or clitics. This is the case, for instance, in Mandarin Chinese and Tagalog, as the examples below demonstrate. Note that both Mandarin *de* and Tagalog *-ng* are not really markers of agreement with the modified noun, but rather just markers of modification of the noun, often referred to as *attributivizers* or *linkers*; the relation between these two types of markers will be discussed in more detail in Sections 4 and 5.

- (24) **Mandarin Chinese**

yi-ge [duli yu fumu] **de** qingshaonian
one-CL independent from parents LNK teenager

‘a teenager who is independent of his parents’

- (25) **Tagalog** (Schachter & Otones, 1972, 246; gloss ours)

Naghaanap ako ng [bagay (para) sa bata]=**ng** damit.
AV.PROG.search 1.SG.NOM GEN suitable for DAT child=LNK dress

‘I am looking for a dress (that is) suitable for the child.’

The examples above thus show that the HFF as defined in (8) clearly faces several counter examples.

2.2 Mirrored languages

As already mentioned above, it is generally well-known in the literature that a number of languages do not obey the HFF. What, to the best of our knowledge, has however not been yet discussed in connection with the HFF is the fact that some languages seem to display the mirror image of it in postnominal position. A representative example of this

type of languages is Basque. Adjectives in Basque are strictly head-final with respect to their complements/adjuncts, as (26a) shows, and they are postnominal when used attributively, cf. (26b). Crucially, complex attributive APs containing a complement/adjunct of the adjective are ungrammatical (26c) and have to be substituted by (prenominal) relative clauses. In other words, the order N–PP–A, which is the mirror image of the order banned by the HFF, is systematically ruled out in Basque.

(26) **Basque**

(Urtzi Etxeberria, p.c.)

- a. Jon bere gurasoetaz burujabe-a da.
 Jon his parents.INSTR independent-ART is
 ‘John is independent of his parents.’
- b. Jon ume burujabe bat da.
 Jon child independent a is
 ‘John is an independent child.’
- c. *Jon ume bere gurasoetaz burujabe bat da.
 Jon child his parents.INSTR independent a is
 Intended: ‘John is a child who is independent of his parents.’

These facts cannot be explained by the HFF. Naturally, one can try to rule out these examples independently. However, structurally languages like Basque exhibit exactly the same patterns as the apparent HFF-obeying languages: An attributive adjective cannot be selected/modified by a PP (or any other phrase for that matter), if this PP intervenes between the adjective and the noun it modifies.

Note, finally, that also the ban on the N–PP–A order can be lifted in some languages, as it is the case for the A–PP–N order. The following examples from Atong (which has a clitical attributivizer) and Persian (which has a head-marking attributivizer) illustrate that.

(27) **Atong**

(Tibeto-Burman; van Breugel, 2010)

- [naŋʔ=məŋ gore [jal=na rak-khal] =gaba] =aw
 2S=GEN horse run=DAT strong-SUP =ATTR =REF
 ‘your fastest running horse (strongest in running)’

(28) **Persian**

(Zahra Mirrazi, p.c.)

madar-an-e be farzand-an-e xod moftaxar
mother-PL-LNK in child-PL-LNK own proud

‘Mothers proud of their children’

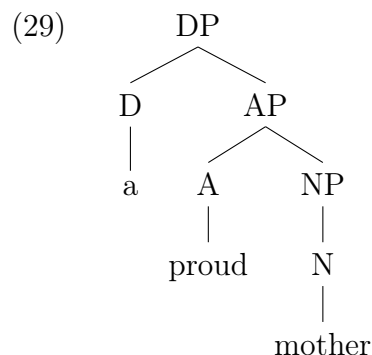
To sum up, these two types of counter examples (though the Basque examples are not counter examples in the strict sense) show that the facts presented above should follow from something else. The HFF both overgeneralizes and undergeneralizes. In this paper we will explain what underlies all these facts. However, first we demonstrate that existing analyses on the ban of A–PP–N configurations cannot explain all discussed facts.

3 Previous accounts

This paper is not the first attempt to account for the ban on A–PP–N configurations. At least three accounts have been formulated to explain the facts, although not all of them aimed at accounting for the attested cross-linguistic differences as well. These accounts are Abney (1987)’s analysis in terms of adjectives occupying the head position in the extended NP (xNP), Escribano (2004)’s analysis in terms of the Linear Correspondence Axiom (LCA) and labeling, and Sheehan (2017)’s analysis in terms of the interaction between the LCA and the head parameter. We’ll discuss each in turn.

3.1 Abney (1987)

In his seminal work on the DP architecture, Abney (1987) argues that attributive adjectives head a position in the extended nominal projection:



Under the assumption that a head can only take one complement (the single complement hypothesis), attributive adjectives cannot select a PP complement because they already have

an NP as their complement. This is different for predicative adjectives, whose complement slot is not occupied by the NP. Furthermore, since the AP contains the NP, an adjunct to an AP can never intervene between A and N. In this way, the A-PP-N pattern can be ruled out as desired, both when the PP is a complement of the adjective and when it is its adjunct. Finally, even though this is not explicitly discussed by Abney (1987), it can account for why postnominal complex APs containing a PP are allowed in English (30d). Following common assumptions (Cinque, 2010), this has to do with the fact that such postnominal APs do not modify nouns directly, but rather occupy the predicative position in a reduced relative clause structure. This together then explains the facts in (30).

- (30) a. the [proud] man
 b. *the [proud of his son] man
 c. *the man [proud]
 d. the man [proud of his son] (Abney, 1987)

However, the strength of Abney’s proposal is also its weakness. The idea that adjectives are heads in the extended projection of the noun has been shown to be problematic for a number of reasons (for a detailed discussion, see Svenonius, 1994; Matushansky, 2002; Escribano, 2006; Pysz, 2006). Here we point out only one major problem with this analysis, namely the fact that Abney’s DP-internal APs simply do not have the distribution of other kinds of APs. If they did, [_{AP} *proud father*] in (29) should be able to occur in other typical adjectival slots, e.g. in the predicative position, contrary to fact:


- (31) *John is [_{AP} proud father].

Moreover, also specifically with respect to the HFF, Abney’s analysis is problematic in being too restrictive in a cross-linguistic perspective, as the only order involving A, N and a PP it can derive is the postnominal N-A-PP order, being a reduced relative clause structure. It cannot explain why the HFF-violating A-PP-N order is in fact possible in certain languages (like Greek or Russian) if attributive adjectives are systematically disallowed to take a complement beyond the NP. Also, it does not explain why the postnominal N-PP-A order is only banned in some languages (such as Basque) but not others (e.g. Atong, Persian). Finally, under Abney’s approach it remains an open question as to how the PP-A-N orders in Dutch or Hungarian, and the A-N-PP orders in English or Serbo-Croatian can be derived (unless in the former case it is assumed that the PP is a specifier of A).

3.2 Escribano (2004)

Another, very different type of approach to the HFF takes it to result from restrictions on (postsyntactic) linearization, rather than from narrow syntactic constraints. One proponent of this approach is Escribano (2004), whose analysis of the HFF relies on specific assumptions concerning the nature of modification and of the labeling and linearization algorithms. More specifically, Escribano assumes that a modifier takes the modified element as its complement, but does not project because what is supposed to project are heads that still have “unsatisfied” features. This is illustrated below for adjectival modification. Given that A is the selecting head in this construction, even though it does not project, Escribano takes it to asymmetrically c-command its complement N and therefore, by Kayne (1994)’s Linear Correspondence Axiom (LCA), to precede N.

(32) a. **syntactic structure:**




b. **linear order:**

A–N

By contrast, if an attributive adjective takes a PP complement first, N serves as its second complement, which corresponds to a *specifier* in X-bar theoretic terms. Since heads asymmetrically c-command complements, but are asymmetrically c-commanded by specifiers in Kayne’s theory, this implies by LCA that the adjective precedes the PP, but follows the noun, as shown below.

(33) a. **syntactic structure:**



b. **linear order:**

N–A–PP

Even though Escribano’s approach to the HFF relies on very different assumptions from the ones that underlie Abney’s analysis, it also only derives the N–A–PP order and, hence, suffers from exactly the same cross-linguistic limitations as Abney’s. Moreover, some of his

specific assumptions about labeling and linearization in modificational structures are also problematic. For instance, it is not clear why the c-command relations are computed on the basis of separate notions of *head*, *complement* and *specifier*, rather than on the basis of the structural relations in the tree (where the noun would count as a head and not as a specifier).

3.3 Sheehan (2017)

A more recent proponent of the LCA-based linearization approach to the HFF has been developed by Sheehan (2017), whose analysis relies on the assumptions (i) that both the head parameter and the LCA are needed for linearization, (ii) that unsuccessful linearization does not necessarily lead to a crash, but may trigger repair operations, and (iii) that attributive APs are base-generated inside postnominal reduced relative clauses.

More specifically, Sheehan assumes that the head parameter is specified for the ordering of heads and their complements, while the LCA operates in the rest of cases as a last resort. Hence, if a language has postnominal adjectives, which for her remain in situ inside a postnominal reduced relative clause, the LCA will linearize the AP (including the PP it contains) as following the head noun, since this head asymmetrically c-commands the AP.⁹ Given that the head parameter of the adjective can be specified as head-initial or as head-final, this yields the following two sets of ordering pairs, leading to the two linearization sequences N–A–PP and N–PP–A:

(34) **Postnominal word order patterns**

<p>a. PP–A (head parameter)</p> <p>N–A (LCA)</p> <p>N–PP (LCA)</p> <p>—————</p> <p>N–PP–A</p>	<p>b. A–PP (head parameter)</p> <p>N–A (LCA)</p> <p>N–PP (LCA)</p> <p>—————</p> <p>N–A–PP</p>
---	---

The situation is different in languages with prenominal adjectives, which for Sheehan are derived by movement from their postnominal base position. Being in a derived position, prenominal adjectives, but crucially not their PP-dependents, asymmetrically c-command the nouns they modify and are therefore linearized by the LCA to their left. By contrast, the adjectival dependents and the modified nouns do not c-command each other, nor do they stand in a c-selectional relation to each other. Therefore no linearization instruction is available for them. This is not a problem when the adjective’s head parameter is specified

⁹Sheehan (2017) assumes Kayne (1994)’s structure of relative clauses, according to which the head noun c-commands the predicative adjective contained in the relative clause.

as head-final: then, the two available linearization instructions PP–A and A–N are sufficient to establish the PP–A–N order by transitivity. Yet, it becomes problematic when the AP is head-initial. Then the ordering between N and PP remains unspecified (the only instruction is that both of them follow A). For those cases, Sheehan assumes that a repair strategy is available where (parts of) higher copies of moved elements are not being spelled out, resulting in the realization of the lower copies. If the lower copy of the PP is spelled out, this leads to the discontinuous A–N–PP order.

(35) **Prenominal word order patterns**

<p>a. PP–A (head parameter)</p> <p>A–N (LCA)</p> <p>N ↔ PP (not specified)</p> <hr style="width: 20%; margin-left: 0;"/> <p>PP–A–N (by transitivity)</p>	<p>b. A–PP (head parameter)</p> <p>A–N (LCA)</p> <p>N ↔ PP (not specified)</p> <hr style="width: 20%; margin-left: 0;"/> <p>A–PP–N–PP (by copy deletion)</p>
--	---

Thus, unlike the other approaches discussed above, Sheehan’s analysis successfully derives four different patterns: PP–A–N and A–N–PP for prenominal adjectives, and N–A–PP and N–PP–A for postnominal adjectives, which covers all the alternative word orders discussed in Section 1.2. However, like the other approaches, her analysis predicts the A–PP–N order to be systematically ruled out and the N–PP–A order to be available across the board, contrary to fact. Moreover, her account relies on a very specific linearization mechanism, the availability of repair strategies, and the movement analysis of prenominal adjectives, all of which may potentially be questioned.

4 A novel generalization

Section 2 has shown that the HFF defined in (8) turns out to be descriptively inadequate to capture the relevant adjacency effects between the heads of modifier phrases and their host nouns. On top of that, the HFF is clearly also analytically inadequate, since it is merely a descriptive statement and does not explain *why* head-finality (or head-initiality) should be a requirement for modifier phrases. Hence, the question arises what determines whether a language allows some material to intervene between modifying adjectives and their head nouns or not. In order to address this question, we have conducted a large survey of approximately 50 languages to see whether they obey or violate the HFF or obey, or violate what can be called the “reverse HFF” or the “Head-Initial Filter”, i.e. the mirror image of the HFF in the postnominal position. As discussions of (im)possible inclusion of material

between nouns and their modifying adjectives in (descriptive) grammars are rather rare, we have not been able to investigate this for every language in our sample. The languages for which we obtained results are listed below.¹⁰

(36) Languages obeying the (reverse) HFF:	Languages violating the (reverse) HFF:
<ul style="list-style-type: none"> - Armenian (Indo-European isolate) - Basque^R (language isolate) - Czech, Serbo-Croatian, Slovenian (Slavic) - Dutch, English, German, Icelandic, Swedish (Germanic) - Estonian, Finnish, Hungarian (Uralic) - French, Italian, Portuguese, Romanian, Spanish (Romance) - Georgian (Kartvelian) 	<ul style="list-style-type: none"> - Atong^R, Mandarin Chinese (Sino-Tibetan) - Bulgarian, Polish, Russian, Ukrainian (Slavic) - Latin, Old Romanian (Italic) - Lithuanian (Baltic) - Modern Greek (Hellenic) - Persian^R (Indo-Iranian) - Tagalog (Austronesian)

On the basis of this survey, we can observe that three relevant factors determine the behavior of languages in this respect: (i) the presence or absence of an overt attributive marker, (ii) the affixal or non-affixal nature of this attributive marker, and (iii) in the affixal case the types of nominal features encoded in it. More concretely, the conditions under which intervening material is permitted between an attributive adjective and its head noun can be stated in terms of the following descriptive generalization that we call the *Modifier-Noun Adjacency Generalization* (MAG). Note that even though MAG in (37) is formulated specifically for modifying adjectives, we assume that it may hold for nominal modifiers of any category.

(37) **Modifier-Noun Adjacency Generalization (MAG)**

The linear order A-XP-N (or N-XP-A, if adjectives are postnominal) is possible iff:

- a. the linear order A-XP (or XP-A) is available in the predicative position, *and*
- b. the adjective modifying the noun exhibits:
 - (i) an overt attributive marker, which is a morpho-phonologically independent

¹⁰Languages of the “reverse” type are marked with the superscript ^R in (36).

clitic/free word form, *or*

- (ii) an agreement marker that is specified for all φ and case (κ) features of N *and* is also present on predicative adjectives.

Condition (37a) of MAG ensures that the trivial prerequisite for the A–XP–N/N–XP–A order is met, namely that right-branching/left-branching APs are allowed in the language at hand to begin with. If A–XP is not a possible word order in a language with prenominal adjectives, the HFF is obeyed trivially (as, e.g., in Japanese or Turkish, see the discussion in Section 1.2). Similarly, languages with postnominal adjectives where the XP–A order is not available, like Italian, trivially obey the reverse HFF (*N–XP–A).

Let us now discuss in more detail the empirical motivation for the disjunctive condition (37b), which is the central part of MAG. We have already seen in Sections 2.1 and 2.2 that the affixal vs. non-affixal nature of the attributive/agreement marker plays a role with respect to whether the language allows for the non-adjacency of attributive adjectives and modified nouns. In particular, we have seen that languages with morpho-phonologically independent attributive markers, i.e. with clitics or free word forms, systematically violate the HFF (for instance, Mandarin) and the reverse HFF (for instance, Persian). The effects originally attributed to the HFF, thus, reduce to properties of morpho-phonologically dependent, i.e. affixal, attributive markers. This motivates part (i) of (37b).

Moreover, we have seen that also some languages with affixal attributive markers can violate the (reverse) HFF, but the conditions under which this is possible have not been clear yet. Based on our sample of languages, we argue that two conditions must be met in this case, as stated in part (ii) of (37b). First, the attributive adjective must be specified for the κ and φ -features that are active in the DP. An illustrative example in this connection is Latin, whose adjectives (as well as nouns) are marked not only for φ -, but also for κ features. In this respect, Latin differs from most modern Romance languages, which lack overt case marking on adjectives (and nouns), as shown below for Italian. As predicted by MAG, it is this difference that underlies the distinction between these languages with respect to the possibility for attributive adjectives not to be noun-adjacent.

(38) **Latin** (Cic. Fam. 1.9; Chiara Gianollo, p.c.)

in [praestantibus in re publica gubernanda] viris
in excellent.ABL.M.PL in thing public to.be.governed man.ABL.M.PL
‘in men who excel in the government of the republic’

(39) **Italian**

*nei [bravi in matematica] studenti
in.the.M.PL good.M.PL in maths student.M.PL

Intended: ‘in the students (who are) good in maths’

Second, these adjectives must be fully marked for all nominal (i.e. κ and φ) features not only when they are used as attributive modifiers, but also when they act as independent predicates. This part of MAG is motivated by the contrast between languages like Greek (HFF-violating) and German (HFF-obeying). In Greek, attributive and predicative adjectives are fully marked for such features, and the HFF does not apply. German, by contrast, has fully marked attributive adjectives, but this is not sufficient to violate the HFF, as predicative adjectives in German are bare. Note that the reverse holds as well: Hungarian, whose predicative adjectives show feature agreement, but whose attributive adjectives are bare, is also subject to the HFF. The examples below illustrate.¹¹

(41) **Greek**

- | | | | |
|----|------------------------------------|----|-------------------------------------|
| a. | (Aftos) einai perifan*(-os). | b. | (Aftoi) einai perifan*(-oi). |
| | he is proud-M.SG.NOM | | they are proud-M.PL.NOM |
| c. | o perifan*(-os) pateras | d. | oi perifan*(-oi) goneis |
| | the.M.SG.NOM proud-M.SG.NOM father | | the.M.PL.NOM proud-M.PL.NOM parents |

¹¹One might question the validity of this part of MAG in light of the distinction between so-called long-form and short-form adjectives in Russian. Unlike attributive adjectives, predicative adjectives in Russian can have morphologically shorter forms, which are equipped with the same (full) set of φ -features as long forms but are not marked for κ (they are only able to occur in nominative/caseless environments). Yet Russian is an HFF-violating language.

(40) **Russian**

- | | | | |
|----|------------------------------------|----|--------------------------------|
| a. | Ona umn-aja/umn-a. | b. | umn-aja/*umn-a |
| | she smart-LONG.F.SG.NOM/SHORT.F.SG | | smart-LONG.F.SG.NOM/SHORT.F.SG |
| | ‘She is smart.’ | | d’evocka |
| | | | girl |
| | | | ‘a/the smart girl’ |

However, these data do not speak against the generalization in (37b-ii) as such, since only the long forms may be used in attributively, as shown by the examples above.

In this connection, it should be pointed out, however, that MAG faces a problem from Slavic HFF-violating languages, such as Czech and Serbo-Croatian, which have richly inflected adjectival forms in the predicative and attributive positions. Our tentative assumption in this respect is that their attributive forms do in fact mark for more features than their predicative counterparts, in particular for specificity (cf. Arsenijević & Stanković, 2009).

(42) **German**

- | | | | |
|----|---------------------------------|----|------------------------------|
| a. | Er ist stolz(*-er). | b. | Sie sind stolz(*-e). |
| | he is proud-MASC.SG.STRONG.NOM | | they are proud-PL.STRONG.NOM |
| c. | stolz*(-er) Vater | d. | stolz*(-e) Eltern |
| | proud-MASC.SG.STRONG.NOM father | | proud-PL.STRONG.NOM parents |

(43) **Hungarian**

- | | | | |
|----|------------------|----|--------------------------|
| a. | (Ő) büszke volt. | b. | (Ők) büszké*(-k) voltak. |
| | s/he proud was | | they proud-PL were |
| c. | egy büszke apa | d. | a büszké*(-k) szülők |
| | a proud father | | the proud-PL parents |

To sum up, MAG states that the heads of modifying APs be noun-adjacent, but that there are two cases in which this requirement can be lifted, allowing HFF violations: either APs have a phrasal, morpho-phonologically independent attributive marker; or their exhibit an affixal agreement marker that is specified for all nominal (κ and φ) features and occurs on both predicative adjectives (i.e., it is a general adjectival agreement marker). MAG thus gives rise to clear predictions: languages with no attributivizers must always be HFF-obeying, languages with overt non-affixal attributivizers should never be HFF-obeying, and languages with affixal markings on the attributive adjectives should be a mixed group.


The discussion above makes it clear that the rationale behind MAG has to do with the familiar relation between the rigidity of word order and the richness of morphological marking: a modifying adjective can be placed at a distance from the noun it modifies only if it is sufficiently marked as belonging to this noun. In what follows, we will spell out an analysis that derives this descriptive generalization.

5 Deriving MAG

In the previous section, we have seen that a modifying adjective may be separated from its head noun by a dependent of this adjective only in the presence of some overt marker. Moreover, this marker must be subject to one of the following conditions: either it must be non-affixal, or it must be specified for all φ - and κ -features and be present on the predicative adjective as well. Hence, this generalization, which we have dubbed MAG, makes reference to two factors: (a) the featural specification of both attributively and predicatively used adjectives, and (b) the morpho-phonological status of an adjectival agreement marker

as an affix, a clitic, a free morpheme or a null morpheme. In order to derive MAG, we need to explain how the interaction between these factors affects adjective-noun adjacency.

We start our explanation by laying out our assumptions about the nature of and the relation between predicatively and attributively used adjectives. We assume that the predicative use of adjectives is the basic one, since this is where they appear as independent predicates. The morpho-syntactic form and the featural specification of predicatively used adjectives thus corresponds to the way adjectives enter the syntactic derivation. Their attributive use can be, by contrast, more complex and turn out to require the presence of additional syntactic structure. This complexity has to do with the workings of modification, in which, we argue, the host of a structure is dissociated from being the projecting element in it. In particular, following ideas present in categorial grammar and implemented in minimalist syntax by Escribano (2005) and Zeijlstra (2020), we assume that modifying adjectives carry at the same time both a categorial [N] feature and a selecting uninterpretable [uN] feature: this captures the nominal character of attributive adjectives as well as the fact that they require the presence of a noun.¹² Since adjectives select for nouns through their selecting uninterpretable [uN] feature, this renders them effectively the heads of the resulting A-N combinations. Phrases containing an attributive adjective and a noun are thus nominal (as desired) because they inherit the interpretable categorial [N] feature from the adjective.




Crucially, the source of the nominal feature on top of the A-N merger is the attributive adjective and not the actual head noun. It is this aspect of the system that allows us to model the requirement that attributive adjectives be feature-complete in the way stated in (37b-ii) of MAG. The reason is that nouns and other elements inside a DP are equipped not only with categorial features, but also with κ -features, as well as a set of φ -features active in the language.¹³ To see this, imagine a language where attributive adjectives are not marked

¹²For the purposes of this paper, we remain open with respect to the question of the categorial specification of predicative adjectives. In cases when an attributivizer is present, it is plausible to assume that it is its role to convert [A] expressions into [N, uN] expressions. However, in languages in which predicative and attributive adjectives have the same inflected form, one would likely need to resort to the assumption that the same form may be featurally ambiguous. Note that this is essentially the same issue as the question in semantics of what to do with predicative adjectives if attributive adjectives are assumed to be of type $\langle\langle e, t \rangle, \langle e, t \rangle\rangle$.

¹³We assume that Case is present as a feature on DPs in all languages independently of whether it is

for the κ and φ -features active in the DP – neither in their predicative nor in their attributive forms. This translates itself into a lack of the corresponding feature slots inside the feature matrix of adjectives. The features and feature values specified on the noun will not be able to percolate up to the next nominal projection in the presence of a modifying adjective, as it is the categorial [N] feature of the φ/κ -featureless adjective that gets projected in this case. As a result, the top node of the A-N merger (the “NP”) ends up being void of any φ/κ -features, as shown below.¹⁴




This outcome is, however, problematic for the φ/κ -completeness of the DP, i.e. the requirement that its highest head must carry the values of all φ/κ -features introduced in the different heads inside the DP, including N itself. This ensures their accessibility outside the DP, primarily for the purposes of DP-external agreement (Danon, 2011). Now, the NP that results from the merger of the attributive adjective and the noun as in (45) above is φ/κ -featureless, but is also the highest NP-projection. Assuming that φ - and κ -features are sub-features of the categorial [N] feature (cf. Adger & Svenonius, 2011, for a discussion) and that, accordingly, the D-head probes for [N] and not specifically for φ and κ , D cannot probe beyond this NP given relativized minimality. Consequently, it will not be able to target any φ/κ -features, and, as a result, the φ/κ -features of the noun will become inaccessible outside the DP.

This problem does not arise if the attributive adjective itself also carries the values of all nominal φ/κ -features available in the language. In that case, they will be projected to the highest NP level, where the probing D-head will have no problem in accessing them, as illustrated below.

(46)

morphologically realized or not.


¹⁴Here and below, the (un)interpretability of φ/κ -features is not indicated for simplicity given that the question concerning the locus of interpretable φ/κ -features is not directly relevant for the issues under discussion.




Naturally, the question arises as to how the attributive adjective obtains these nominal φ/κ -features itself. If the attributive adjective already contains these features when merging with the noun, this means that these features must already be introduced in the extended adjectival projection (xAP) and, therefore, must also be present on the predicative form of the adjective (i.e. the basic form).¹⁵ Note that this reflects the requirement stated in (37b-ii) that both the attributive and the predicative form of adjectives must be equally specified for all nominal φ/κ -features.

Nothing hinges in our analysis on where exactly inside the xAP the nominal φ/κ -features are introduced; in fact, they might even be introduced on the lexical A-head itself. However, for simplicity we assume that these features are located on a separate functional head, which represents the adjectival agreement marker and which we will refer to as ‘Agr’. This is depicted in (47a) in featural terms and in (47b) in traditional labels. What was referred to as “adjective” or “adjectival modifier” before thus corresponds to the full xAP in (47b), i.e. to the AgrP and not the sole AP.

(47) a.



b.



¹⁵The question of how the specific *values* of these φ/κ -features end up being on the adjective is not crucial for us here: they either come specified for them from the lexicon or alternatively are first missing and get valued by the noun upon merge. Note that the former case would also require an additional mechanism which ensures that the features of the adjective match with those of the noun.

What is crucial for our analysis is that adjectival dependents are introduced within the AgrP, rather than the AP proper (see Merchant, 2019, for recent arguments that PPs are selected by *a* and not by A). The Agr marker must always be linearly adjacent to A. For the rest, no linearization constraints prevent Agr' from taking a dependent to its left or to its right, as in (48a) and (48b), the latter being the configuration with an intervening PP. Such configurations are found in HFF-violating languages with affixal agreement marking on adjectives, such as Greek.¹⁶

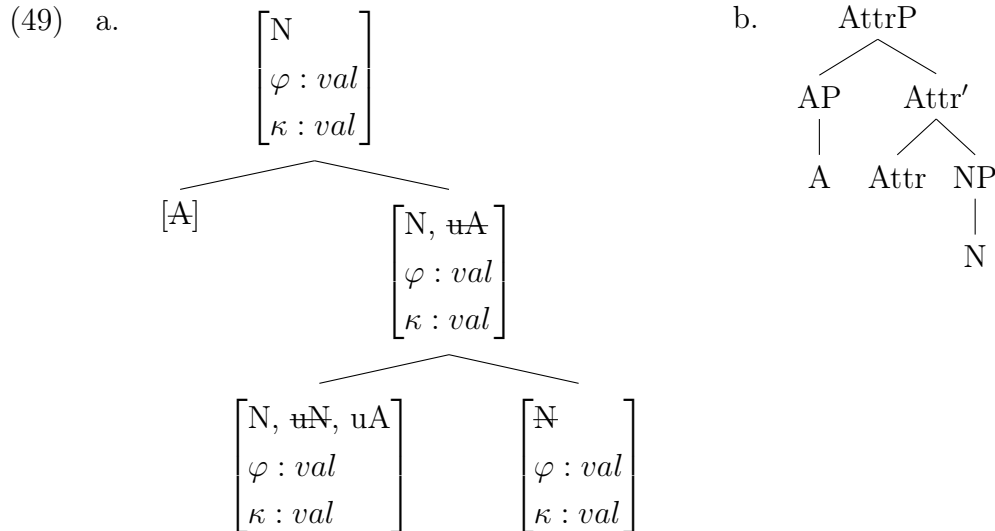


This way, the assumptions outlined above make it possible to capture the generalization in (37b-ii). The featural endowment of attributively used adjectives and the combinatorial principles behind adjectival modification explains why adjectives, despite being “adjuncts”, can act as interveners for Agree: they block φ/κ -feature transmission.


The next question is what happens when predicatively used adjectives are not specified for any or some of the nominal φ/κ -features. As shown in (45), such adjectives should lead to a crash given the lack of φ/κ -completeness. To prevent this crashing from happening, the adjective must be prevented from projecting its nominal features. We argue that this is the role of an attributivizer, i.e. the linker for adnominal modification, which we represent in terms of a separate functional head Attr. Unlike Agr, Attr is part of the xNP and not of the xAP, since attributivizers occur only with attributively used adjectives and are absent in predicatively used ones. Being part of the nominal spine, Attr selects an NP as its first argument. As a result, Attr inherits all nominal φ/κ -features, which get projected together with its categorial [N] feature. Then, it selects the AP as its second argument. This way, the AP no longer acts as an intervener:¹⁷

¹⁶The structures involving prefixal Agr and those for postnominal adjectives with affixal Agr are analogous, though we do not have concrete examples of languages in our sample.

¹⁷Again, we remain agnostic whether these are prespecified or received upon agreement with the noun. Note, moreover, that Attr does not need be overtly marked for *all* these specific features, though of course it can. Crucially, what Attr spells out is its nominality.



Now, the adjective can in principle take a dependent to its left or to its right, just as in the scenario involving Agr discussed before. Unlike in that scenario, however, there is an additional constraining factor. Unlike the case of Agr in the xAP, any dependent of the adjective may linearly separate Attr from the A-head, as in (50b). As a consequence, in (50b) Attr can only morpho-phonologically combine with an adjective, if it is a clitic or an unbound morpheme. Of course, this is not the case in (50a), where the adjective's dependent does not intervene between the adjective and Attr.



Hence, as described in (37b-i), the morpho-phonological status of the attributive marker as an affix, clitic or a null morpheme plays a crucial role in making HFF violations possible or not. We discuss this role in detail in the next section.

6 The morpho-phonological status of attributivizers

6.1 Attr as a clitic

Given the above, in languages in which predicatively used adjectives do not exhibit φ/κ -completeness, an additional Attr-head must be present in the xNP to host the relevant

φ/κ -features. This was shown in (49)-(50). As such, the presence of this Attr head should not systematically ban A-PP-N orders. If every terminal is spelled out independently, there is no reason why it should ban such orderings. This is, in fact, in full accordance with MAG: If the modification of N by A exhibits an overt attributive marker, which is a morpho-phonologically independent clitic or free word, A-XP-N orders are freely available. This is in indeed the case for Mandarin Chinese and Tagalog, and in a mirrored fashion for Atong and Persian.

The fact that A-XP-N orders in languages where Attr has to be art of xNP requires Attr to be morpho-phonologically independent, suggests that the ban on XP-intervention between nouns and their modifying adjectives is due to Attr not being a morpho-phonologically independent. This amounts to two different cases: Cases where Attr is realized by an affix and cases where Attr is realized by a null morpheme. We discuss both in turn.

6.2 Attr as an affix

If the Attr-marker is affixal in nature, its distribution is constrained both syntactically and morpho-phonologically. Following Ackema & Neeleman (2000), we assume that an affix can only attach to the head of the phrase it selects. The following constraint proposed by Ackema & Neeleman (2000), which they call the Input Correspondence Principle (ICP), formulates this restriction concerning possible hosts of affixes.

(51) **Input Correspondence Principle (ICP)** (Ackema & Neeleman, 2000)

If an AFFIX takes a head Y or a projection of Y as its input,
the AFFIX is phonologically realized as /affix/, and
Y is phonologically realized as /y/,
then /affix/ takes /y/ as its input.

This means that Attr, when affixal, must be strictly adjacent to A. However, since Attr first selects an NP and an only then an AP, it does not necessarily end up in a position adjacent to the adjective, as can be seen in (49)-(50). In this way, Attr behaves differently from Agr, which first selects an AP and therefore naturally ends up morpho-phonologically attached to it (which means that there is no objection to modifying AgrP with a PP, either on the left or on the right). In order to ensure that Attr also ends up in a string-adjacent position to the adjective, it must be prevented that further material intervenes between Attr and A.

This may give rise to two kinds of effects. First, in languages in which APs are strictly head-initial and prenominal or strictly head-final and postnominal, attributively used adjectives

tives cannot take any dependents. Second, in languages in which PPs can appear both to the left and to the right of a predicatively used adjective, attributively used adjectives can still take leftward PPs if APs are prenominal (or rightward PPs if APs are postnominal).

The restrictions on the attachment of affixes formulated as the ICP in (51) thus naturally capture the generalization stated in (37b-i) and explain the fact that languages like German obey the HFF.

6.3 Attr as a null morpheme

The third type of language to be discussed are languages in which an Attr-head is predicted to be present, but is not realized overtly. Two sub-types of such languages can further be distinguished: languages in which *not all* instances of Attr are morpho-phonologically realized and languages in which *no* instance of Attr is morpho-phonologically realized. We discuss both varieties in turn, starting with the former one, in which the Attr-head is only sometimes realized overtly.

A language that belongs to this variety is Dutch. Dutch is similar to German in that predicatively used adjectives are always bare forms, as the following examples show.

- (52) a. De man is lang.
 the.COM man is tall
 b. Het kind is lang.
 the.NEUT child is tall
 c. De mannen zijn lang.
 the.PL men are tall
 d. De kinderen zijn lang.
 the.PL children are tall

Dutch differs from German, however, in that nouns do not always trigger overt morphology on attributive adjectives: attributives take a schwa-ending, except when they modify an indefinite neuter singular noun; then the agreement marker is null (see the following table and (53)).

	COMM/PL	NEUT
definite	-e	-e
indefinite	-e	-∅

Table 2: Paradigm of attributive marking on adjectives in Dutch.

- (53) a. de jong-e vrouw
the.COM young-ATTR woman.COM
- b. het jong-e kind
the.NEUT young-ATTR child.NEUT
- c. een jong-e vrouw
a young-ATTR woman.COM
- d. een jong-∅ kind
a young-ATTR child.NEUT

With *overt* attributive morphology (i.e. with the schwa), things work in the same way as in German when it comes to PPs intervening between the noun and the modifying adjective. The schwa must right-attach to the adjective that heads the AP. Therefore attributive adjectives can only take a PP on their left, but not on their right.

- (54) a. een [op haar moeder trots_A]-e vrouw
a of her mother proud-ATTR woman
- b. *een [trots op haar moeder_N]-e vrouw
a proud of her mother-ATTR woman
- c. *een [trots-e op haar moeder] vrouw
a proud-ATTR of her mother woman

This situation is exactly the same when the agreement marker is *covertly* realized, as shown in the examples below (with a neuter noun instead of a common one). Like before, all adjacency requirements are met only in (55a). In (55b), Attr is not adjacent to the head of the AP, and in (55c) it is not adjacent to the NP.

- (55) a. een [op haar vader trots_A]-∅ kind
a of her father proud-ATTR child
- b. *een [trots op haar vader_N]-∅ kind
a proud of her father-ATTR child
- c. *een [trots-∅ op haar vader] kind
a proud-ATTR of her father child


Since uninflected forms of attributive adjectives in Dutch indeed contain a syntactically present but phonologically unrealized Attr, the analogous behavior of complex APs with

overt and covert attributive morphology in (54) and (55) is not surprising at all. Therefore, one might entertain the idea that also null affixes are subject to the ICP in (51), even if Ackema & Neeleman (2000)'s original conception of the ICP was not supposed to cover zero morphology. In this light, it is important to observe, however, that Dutch uninflected attributive forms do in fact behave differently from attributive schwa-forms with respect to certain linearization phenomena. Consider the following paradigm based on observations by van Riemsdijk (1998).

- (56) a. een [groot genoeg_{Deg}]-∅ kind
 a big enough-ATTR child
 b. *een [groot genoeg_{Deg}]-e vrouw
 a big enough-ATTR woman
 c. *een [groot-e genoeg_{Deg}] vrouw
 a big-ATTR enough woman
 d. *een [groot genoeg_{Deg}] vrouw
 a big enough woman

When the attributive agreement marker is overt, the postadjectival degree marker *genoeg* ‘enough’ interrupts the adjacency between the schwa and the adjective, triggering ungrammaticality ((56b)-(56d)). This is not the case, however, when the attributive marker is null (56a). This suggests that zero morphology imposes an adjacency requirement with respect to its host that is less strict compared to that of overt morphology.

The most straightforward way of accounting for the data above offers itself under the assumption that degree markers head a dedicated functional projection inside the xAP (cf. Abney, 1987; Grimshaw, 1990; Kennedy, 1999). In this case, no additional stipulations are necessary for the null version of Attr as in (56a): it simply follows the ICP in that its phonological realization “attaches” to the head of the phrase it selects syntactically, i.e. DegP.¹⁸



¹⁸Note that in order to account for cases with preadjectival degree words, the ICP will probably need to be relativized to extended projections.

However, this implies that more needs to be said about the overt version of Attr in (56b), which cannot attach to the head of the DegP *genoeg* and, therefore, must follow additional restrictions.¹⁹ One obvious additional constraint on overt affixes is that they can only attach to a non-lexical head in an extended projection if this head is an affix itself, which, for instance, disallows that verbal tense/agreement morphology attaches to sentential negation in English (cf. **Mary eat not-ed*). This constraint, which we will call the *Affix Continuity Constraint* (ACC), can be formulated as follows:

(58) **Affix Continuity Constraint (ACC)**

If an AFFIX Z takes as its input a head Y or a projection of Y,
 where Y is part of the extended projection of X,
 and Z is phonologically realized as non-null /z/ and Y as /y/,
 then Y must be an AFFIX as well, such that /z/ takes /y/ as its input, and /y/ takes as its input the phonological realization of the head taken by Y as its input.

Now, the question arises as to how to account for languages like English, where the Attr head never gets morpho-phonologically realized. For our analysis to work, English attributive adjectives need to be modified by a covert Attr head. However, unless the presence of such a null Attr can be evidenced, our explanation for languages like English would end up being circular. We believe that there is actually good evidence for the presence of a zero Attr in English, though. The reason is that the distribution of the English null forms is virtually identical to that of the Dutch ones (whose presence is paradigmatically licensed and therefore undisputed). Just like Dutch, English adjectives may not be modified by a rightward PP, except when they are modified themselves by degree-expression like *enough*-phrases, exactly as was attested in Dutch:


- (59) a. a [proud_A]-∅ child
 b. *a [proud of his mother_N]-∅ child
 c. a [proud enough_{Deg/A}]-∅ child

As, the distributional pattern of English attributive adjectives is identical to that of those Dutch attributive adjectives that exhibit covert agreement morphology, it can safely be concluded that the English Attr head is also realized by a null affix. This fact is not too

¹⁹This raises the more general question of whether it is covert or overt morphology that should be more restricted in its distribution. On the one hand, null morpho-syntactic material is subject to recoverability requirements, which do not affect overt material. On the other hand, however, affixation of overt morphological material is sensitive to morpho-phonological constraints that covert morphology is trivially not subject to. It is this latter type of restrictions that seems to be relevant in the present context.

surprising, given that English reflects fewer inflectional distinctions in general, a result of a higher amount of deflection emerging in the course of time.

A final note has to be made about languages where predicatively used adjectives are inflected for some, but not all features that are active in the DP. Italian is a good example. Even though Italian predicatively used adjectives are inflected for φ -features, they are not for κ -features. This means that in such languages, the presence of an Agr head in the xAP does not exclude the presence of a null Attr head in the xNP. In fact, such an Attr head is even required, since otherwise AgrP would again intervene for κ -Agree between D and N. That means that the structure of an NP modified by an attributive adjective in a language like Italian is as in (60), where both Agr and Attr are present. Consequently, in such a configuration, a PP or other modifier may not appear in between Agr and N, a prediction that is indeed borne out.



7 Concluding remarks

In this paper we have argued that the HFF in its current formulation is not only left theoretically unexplained, but is also empirically incorrect. For one, languages like Russian or Greek are not subject to it and can have material intervene between prenominal adjectives and their nominal modifiees. In addition, a number of languages, such as Basque, display mirror-HFF effects. These languages bar material intervening postnominal adjectives and their nominal modifiees, something that is not captured by the HFF.

On the basis of the language sample we built up, we have argued that the HFF should be replaced by the following generalization that we have dubbed the *Modifier-Noun Adjacency Generalization (MAG)* and that we repeat below:

(61) **Modifier-Noun Adjacency Generalization (MAG)**

The linear order A–XP–N (or N–XP–A, if adjectives are postnominal) is possible iff:

- a. the linear order A–XP (or XP–A) is available in the predicative position, *and*

- b. the adjective modifying the noun exhibits:
 - (i) an overt attributive marker, which is a morpho-phonologically independent clitic/free word form – *or*
 - (ii) an agreement marker that is specified for all φ and κ features of N *and* is also present on predicative adjectives.

The rationale behind MAG is the fact that direct modification of N by A is impossible and must be mediated through a functional element which is specified for all nominal φ and κ -features. Such a functional head may already be part of the xAP. This is the case when all relevant φ and κ -features are present on the predicatively used adjective. Then, there is nothing that would forbid the adjective to select material to its immediate right. If this functional head is part of the xNP, however, its the morpho-phonological properties determine whether it must be adjacent to both the noun and the adjective or only to the noun. In the former case, no material may intervene between A and N; in the latter case, it may. This, we have argued, derives MAG.

As a final note, we emphasize that all the DP-internal linearizations we investigated are fully captured by MAG (and thus also explained). We also note that there may be more types of MAG-compatible grammatical systems than observed and described in this paper. For instance, we have not attested languages with null attributive morphology that allow for both A–PP and PP–A orders. Whether the fact that we have not observed all MAG-compatible grammatical systems (yet) is the result of there being typological gaps due to independent grammatical constraints or the result of the size of our sample is a question we leave for future research.

References

- Abney, S. P. (1987). *The English Noun Phrase in its Sentential Aspect*. MIT: Ph.D. dissertation.
- Ackema, P., & Neeleman, A. (2000). M-selection and phrasal affixation. In C. Iten & A. Neeleman (Eds.), *Ucl working papers in linguistics* (Vol. 12, pp. 307–344). Department of Phonetics and Linguistics, University College London.
- Adger, D., & Svenonius, P. (2011). Features in Minimalist Syntax. In C. Boeckx (Ed.), *The oxford handbook of linguistic minimalism*. doi: 10.1093/oxfordhb/9780199549368.013.0002
- Bernstein, J. (1995). *Adjectives and their complements*. Unpublished LSA paper.

- Cinque, G. (2010). *The Syntax of Adjectives: A Comparative Study*. Cambridge, MA: MIT Press.
- Danon, G. (2011). Agreement and DP-internal feature distribution. *Syntax*, 14(4), 297–317. doi: 10.1111/j.1467-9612.2011.00154.x
- Emonds, J. E. (1976). *A Transformational Approach to English Syntax: Root, Structure-Preserving, and Local Transformations*. New York: Academic Press.
- Escribano, J. L. G. (2004). Head-final effects and the nature of modification. *Journal of Linguistics*, 40(1), 1–43.
- Escribano, J. L. G. (2005). "Discontinuous" APs in English. *Linguistics*, 43(3), 563–610.
- Escribano, J. L. G. (2006). NPs as just NPs. *Language Sciences*, 28(6), 529–579.
- Greenberg, J. H. (1963). Some universals of grammar with particular reference to the order of meaningful elements. In J. H. Greenberg (Ed.), *Universals of language* (pp. 73–113). Cambridge, MA: MIT Press.
- Grimshaw, J. (1990). *Argument Structure*. Cambridge, MA: MIT Press.
- Grosu, A., & Horvath, J. (2006). Reply to Bhatt and Pancheva's "Late Merger of Degree Clauses": The Irrelevance of (Non)conservativity. *Linguistic Inquiry*, 37(3), 457–483. doi: 10.1162/ling.2006.37.3.457
- Kayne, R. (1994). *The antisymmetry of syntax*. Cambridge, MA: MIT Press.
- Kennedy, C. (1999). *Projecting the adjective: The syntax and semantics of gradability and comparison*. New York: Garland.
- Matushansky, O. (2002). *Movement of degree/degree of movement*. MIT: Ph.D. dissertation.
- Merchant, J. (2019). Roots don't select, categorial heads do: Lexical-selection of PPs may vary by category. *The Linguistic Review*, 36(3), 325–341. doi: 10.1515/tlr-2019-2020
- Nanni, D. L. (1980). On the Surface Syntax of Constructions with Easy-Type Adjectives. *Language*, 56(3), 568–581. doi: 10.2307/414450
- Philip, J. (2012). *Subordinating and Coordinating Linkers* (Unpublished doctoral dissertation). University College London.

- Pysz, A. (2006). The structural location of adnominal adjectives: Prospects for Old English. *SKASE Journal of Theoretical Linguistics*, 3(3), 59–85.
- Sadler, L., & Arnold, D. J. (1994). Prenominal adjectives and the phrasal/lexical distinction. *Journal of Linguistics*, 30(1), 187–226.
- Schachter, P., & Otnes, F. T. (1972). *Tagalog Reference Grammar*. Berkeley: University of California Press.
- Sheehan, M. (2017). A labelling-based account of the Head-Final Filter. In L. Bauke & A. Blümel (Eds.), *Labels and roots* (pp. 161–202). Berlin, Boston: De Gruyter Mouton. doi: 10.1515/9781501502118-008
- Svenonius, P. (1994). The Structural Location of the Attributive Adjective. In E. Duncan, D. Farkas, & P. Spaelti (Eds.), *Proceedings of the 12th west coast conference on formal linguistics* (pp. 439–454). Stanford, CA: CSLI Publications.
- van Breugel, S. (2010). No common argument, no extraction, no gap: Attributive clauses in Atong and beyond. *Studies in Language*, 34(3), 493–531. doi: 10.1075/sl.34.3.01van
- van Riemsdijk, H. (1998). Head movement and adjacency. *Natural Language and Linguistic Theory*, 16(3), 633–679. doi: 10.1023/a:1006075919357
- Williams, E. (1982). Another argument that passive is transformational. *Linguistic Inquiry*, 13(1), 160–163.
- Zeijlstra, H. (2020). Labeling, selection, and feature checking. In P. W. Smith, J. Mursell, & K. Hartmann (Eds.), *Agree to agree: Agreement in the minimalist programme* (pp. 31–70). Berlin: Language Science Press.