Adjective order in English: A semantic account with cross-linguistic applications

Enrica Rosato

Linguistics

Department of Philosophy, Carnegie Mellon University

Senior Honors Thesis, May 2013

Advisor: Dr. Thomas Werner

Adjective order in English: A semantic account with cross-linguistic applications

Enrica Rosato

Table of Contents

Abstract

Which is more correct, the "big fat cat" or the "fat big cat?" Why is a particular order preferred? In English, established Phrase Structure rules place no limit on the number of adjectives before a noun. The adjectives, however, cannot occur in just any order, and native speakers of English have very particular intuitions about what order is more correct, even if they have never been explicitly taught ordering rules. In this study, I seek to describe the mechanics of an underlying adjective order in English and explore if the same principles operate cross-linguistically. After outlining some previous work across disciplines on the subject, I prove the existence of a preferred order using the results of searches from the Corpus of Contemporary American English and the British National Corpus. Additionally, I briefly discuss the prosodic differences between a given order and its alternative. Secondly, I develop a semantic theory that describes how pre-nominal adjectives are ordered based on their semantic properties, with adjectives that depict "intrinsic" properties closer to the noun, and adjectives that are "speaker relative" in a more distant position. In the theory, the use of multiple adjectives is described as being equivalent to a sequential series of restrictions placed on the set of properties for a given noun. This allows for a change in adjective order to affect the way in which we conceptualize of a noun, while also establishing an underlying order that is the most cognitively efficient. Lastly, I apply the theory to languages from different families, namely Italian, Sakha, Hebrew, and Welsh. My findings demonstrate that a uniquely semantic theory is successful in describing what native speakers perceive as the "proper" order of adjectives in a diverse group of languages.

1. Introduction

Native speaker intuitions, corpus searches, and prosodic features of ordered adjectives in English all provide substantial evidence that the order of adjectives within a NP is subject to certain restrictions. This idea has been explored by previous work on the subject of adjective order beginning in the 1960s and heavily in the 1990s, and has been reaffirmed in recent years by studies in Cognitive Science. The extent to which an underlying order exists cross-linguistically has also been explored previously; however, linguists have remained pretty consistently within the domain of syntactic analysis in the cross-linguistic context. These syntactic descriptions seek to describe cross-linguistic ordering variation by invoking syntactic movement principles, and have been successful in doing so. However, to my knowledge, a semantic analysis for English adjective order has not been revisited since the 1970s and has not been applied cross-linguistically. My goal is not to simply describe the syntactic variation among adjective order across languages, but to establish a unifying semantic account for how adjectives restrict a noun's set of properties, and the order in which they do so. My theory intends to focus on the mechanics of underlying adjective ordering (AO) rules, and explores how these mechanics operate cross-linguistically.

I have chosen English as the primary language of analysis for many reasons. Among the most prevalent of these reasons is the fact that it is my native language, and I can therefore trust my intuitions as a native speaker and the intuitions of my English-speaking peers to be the basis for proper assessments of "correctness." Secondly, English places very few restrictions on adjectives in addition to the ordering restrictions; for example, adjectives in English are pre-nominal only, with no declension system or gender and number agreement rules. This makes English AO rules more important in terms of the relationship between an adjective's placement and the semantic properties it contributes to the NP, and

there are fewer variables to control for when analyzing the effects of an alternative ordering. With a theory that accounts for AO and only AO, it is possible to see how a semantic account behaves in languages where there are additional restrictions on the adjectives. The additional languages I choose to explore are not thought to be historically related, and on the surface, all treat the syntax of adjectives within a NP differently. Sakha (Yakut) follows a similar structure to English. In Italian, however, adjectives occur both post-nominally and pre-nominally, and in Welsh, all adjectives occur post-nominally. Hebrew is unique in that it embeds the determiner, and the adjectives occur post-nominally. It is also important to note that in this paper, I avoid discussing the role of the determiner or of quantifying adjectives (many, few, etc.) and number. I also choose to refer to a string of adjectives and the noun they modify as a NP, not a DP. This is done so that the emphasis of the discussion is on the semantic value of the noun and *attributive* adjectives rather than on the role of the determiner or any quantifiers.

I begin the discussion by citing previous work on the subject in the domains of semantics, syntax, cognitive science, and from the perspective of functional grammar. In part 3, I begin my own analysis by providing additional evidence for the existence of an underlying AO in English using corpus data and a brief discussion of prosody. In part 4, I begin my semantic description of the underlying order of English, where I establish my own semantic categories for adjectives and place those into even broader categories, which I refer to as "spheres." I then attempt to explain, using a concept I call "cognitive efficiency," why altering an underlying order is inefficient both within a sphere and across sphere boundaries without being given the proper contextual information. I discuss the semantic categories of adjectives in terms of the "intrinsitivity" or "speaker relativity" of the properties they denote, and I establish that adjectives denoting more intrinsic properties must occur closer to the noun while those denoting relative properties occur farther away. Lastly, in part

5, I apply what I establish as the underlying order for English to the aforementioned languages and compare it to previous analyses for AO in those languages. When I apply the theory, I am mostly looking to see if the underlying order that I establish for English has cross-linguistic applications.

2. Previous studies

The work that has been done to study the sequence of pre-nominal adjectives in English was at its most prevalent in the 1960s, when it was first noted by psycholinguists that speakers have a preference for certain orders. At the time, one of the most popular strategies of contemporary psycholinguistics had been to interpret various laboratory phenomena in terms of constructs derived from linguistic theory. At the forefront of this work were three psycholinguists by the names of J.E. Martin, (1969), J.H. Danks, and S. Glucksberg (1971). In his 1969 paper, Martin proposed a semantic rule for explaining adjective ordering that was based on the qualities referred to by the adjectives- definiteness, absoluteness, or intrinsicalness[1]. According to Martin, in comparison to size, color is more definite in meaning, changes less from object to object (absoluteness), and is considered a more intrinsic property of the object. Additionally, Martin claimed that the semantic dimensions of definiteness and absoluteness determine the "accessibility" of adjectives, which in turn determine their ordering. He defined "accessibility" in terms of how quickly participants could produce an adjectival description of a physical stimulus. He found that the response speed was correlated with ordering: The nearer an adjective was preferred to the noun, the more "accessible" it was.

---

[1] In this paper, I attempt to adapt and expand upon Martin's "intrinsicalness" idea; instead I refer to it as "intrinsitivity."

In their study, Danks and Glucksberg (1971) considered violations of adjective ordering constraints by having participants complete a ranking test with six possible permutations of three pre-nominal adjectives. The results showed that the position of the adjective that was most closely related to an intrinsic property of the noun was the primary determinant of acceptability: the closer it was to the noun, the higher the sentence was ranked. This demonstrates a speaker preference to place adjectives that denote intrinsic properties of a noun closer to the noun.

In 1985, M.A.K Halliday published his first edition of *An Introduction to Functional Grammar,* where he provides a theory for adjective order that describes how adjectives "decrease in specifying potential" and become "increasingly permanent as attributes" as they approach the noun. Halliday's concept of permanence is a way of re-framing Martin's idea of "intrinsicalness" from the functional perspective rather than from the perspective of meaning. He also makes the claim that "the more permanent the attributes of a thing, the less likely it is to identify in context." This is similar to the parallels that Martine draws between intrinsicalness and accessibility, although Halliday did not do a psychological test to determine the validity of his assumptions about permanence.

In the 1990s, several studies were done on restrictions on the syntax of adjective ordering cross-linguistically (Sproat and Shih 1991, Svenonius 1993, Bernstein 1993, Cinque 1994, and Bouchard 1998, to name a few). These studies focused primarily on describing the variation of the internal placement of constituents within DPs across languages. Sproat and Shih (1991) looked at adjective ordering in Chinese, and found that it was relatively free in comparison to English when the adjectives were marked by particles. In general, their conclusions lead them to lean against the idea of a universal semantic description for adjective order. Despite this conclusion, they still established a general ordering hierarchy, Quality > Size > Shape > Color > Provenance. Cinque (1994) modified

this order to be slightly different, establishing broader semantic categories and claiming that they exist in the following order: Possesive > Speaker-oriented > Subject-oriented > Manner/Thematic. Cinque focused on the relationship between the ordering restrictions on Germanic and Romance languages, giving a syntactic explanation for the differences. He introduced the idea that in Italian (a language with pre and post-nominal adjectives), it is the noun that moves within the DP, with the adjectives remaining in the same order relative to each other as they occur in English.

The idea of internal movement of the noun within a DP was controversial, and many linguists published critiques of Cinque's work into the late 1990s. Most recently, in 2006, a linguist named David Willis published a critique of N-raising movement theory as it would apply to the Celtic languages, namely Welsh. Willis cited examples of Welsh NPs demonstrating both "mirror image" and "universal" adjective orders, making it nearly impossible for a coherent syntactic explanation to work consistently in the language. In general, the discussion within the linguistics community about AO in recent years has been focused more on cross-linguistic inconsistencies regarding the placement of the noun in multi-adjective NPs, with most of the questions being about the usefulness of movement theories in explaining these inconsistencies. In my own paper, I seek to return to the development of a uniquely semantic theory that explains the underlying ways in which language functions to describe objects in the physical world. With a more complete understanding of these basic concepts, I believe movement theories will be minimally necessary when attempting to explain cross-linguistic variation in adjective use.

In recent years, neurologist David Kemmerer at Purdue University has used more modern methods to explore some of the same questions about the cognition of AO restrictions that were being asked in the early 1970s. In 2000, he tested the supposition that the features of adjective meaning that constrain their linear order reside at a different level of

mental representation than the features that are invisible to syntax. "If this is true," he claimed, "then it is possible that the two components of meaning could be impaired independently of each other by brain damage (Kemmerer 2000)." In his experiment, six of 16 brain-damaged subjects failed a test that assessed their knowledge of the semantic principles that determine pre-nominal adjective order. However, all of the subjects performed within normal limits on a test that assessed their knowledge of the grammatically irrelevant perceptual and conceptual features of the same adjectives that were used in the first test, and all of the subjects performed well on a test that assessed their knowledge of the basic syntactic structure of English NPs. This study therefore provides support for the view that there is an independent level of representation in the mind/brain for grammatical semantics. In 2007, Kemmerer used ERPs (event-related brain potentials) to assess how participants reacted to "proper," reversed, and contradictory adjective pairs when reading aloud. Increased brain activity was recorded for both the reversed and contradictory pairings. Kemmerer believed this to be a reflection of semantic and syntactic aspects of a temporary reanalysis of the adjective order construction, providing further evidence that certain adjective orders require more cognitive effort.

The topic of proper adjective ordering has spanned across disciplines for many decades. All of the work I have mentioned has informed my desire to develop my own semantic account for the proper ordering of adjectives in English, and attempt to apply this across languages. While the work of semanticists in the late 1960's has inspired me to develop my own semantic theory for proper AO, Kemmerer's work has inspired me to use this semantic theory to develop my own ideas for why certain pairings are optimally cognitively efficient. The work of Cinque, Willis, and other linguists in the 1990s and 2000s has inspired me to attempt to apply my theories across languages from multiple language

families. Additionally, I have chosen to incorporate prosodic analysis and some methodology from corpus linguistics as proof for the existence of an underlying order.

3. Proof of an underlying adjective order in English: **"big red" vs. "red big"**

The existence of a preferred underlying adjective order of English can be proved by both corpus analysis and prosodic analysis. Let us, for the sake of simplicity, take a common case of two adjectives depicting size and color, "big" and "red." Case (a.) below shows the adjectives occurring in their more common underlying order and case (b.) demonstrates the inverse, which is perceived by native speakers of English to be less correct. This would make SIZE > COLOR the proper underlying pre-nominal order. Using a corpus, we can determine and compare the ways in which the cases are actually used.

(a). the big red house
(b). # the red big house

My hypothesis for the corpus analysis is that "big red" will occur at a very high frequency, both in written and spoken language. I hypothesize that "red big" will only occur in contexts of spoken language where the speaker is giving an off-the-cuff description of an event or object. Within this description, I would imagine the speaker to pause often and insert discourse makers such as "um" or "like." Additionally, this type of description may elicit such an order because the object being described is being conceptualized or recalled at the same rate at which language is being produced. If the object were to be already conceptualized and simply being described using language, I predict that the underlying order would surface.

Using COCA, the Corpus of Contemporary American English, I searched for occurrences of "big red [noun]" versus occurrences of "red big [noun]." "Big red" yielded over 382 tokens in many contexts, and "red big" yielded 0 tokens. I then moved on to the BNC (British National Corpus). When I searched for "red big" on the BNC, I found only one occurrence of "red big." It was from a speaker categorized as "FLP," which is defined in the corpus as a speaker within a group of 10 Scottish women, who are having a discussion about weddings. The token was recorded on an unknown date and was transcribed as follows:

Example i.
FLP: And their bride has to wear er, all in red and er the bridegroo er groom has to wear a long costume with a **red big** flowers in front i aha and then they get married and there's erm band, the Chinese traditional band with drums and trumpets blowing all the time and er, all the guests have a very nice time.

Much as I expected, the speaker is in the process of describing a wedding. Her utterance is filled with what the transcriber writes as "er," the British equivalent to the modern American "um." In this context of off-the-cuff description, the Scottish woman says that the groom has to wear a "long costume with a red big flowers." Unfortunately, we do not have a recording of this utterance and the transcriber did not annotate it with the prosodic elements that would provide it with some spoken context. Additionally, it is possible that the previous mention of the word "red" in the phrase "all in red" made it contextually reasonable on a discourse-relevance level to list "red" first in the order. Another possible explanation for why this speaker may have ordered her adjectives in this manner is that in her dialect of Scottish English, it is possible that this is acceptable, although no other Scottish speakers in the corpus produced an inverse order and I found no evidence for this fact.

All other occurrences of these adjectives adjacent to one another in the BNC yielded case (a) order. Below are just some of the contexts in which case (a) was recorded. None of

the instances in which "big red" occurred were uttered by a Scottish woman, so there is a possibility that my dialect theory may not be entirely invalid.

> **ii. (2619)** Now why would Kurt be running around Soho with a **big red** woolly jumper on and having intense conversations with lampposts?
> **iii. (303)** Inside were all the things they had asked for, and some they had not- some wine, two chickens, twelve **big red** roses.
> **iv. (901)** Love does not involve giving fancy parcels tied up with **big red** bows.
> **v. (5406)** And it, and it like, it run down me face there and I had a **big red** mark on me face.

Only one of the above cases involves the detailed description of something, and that is example ii. (303). According to the BNC, this example is from a book called *The Railway Children* published in 1993. Because this is written data, it cannot be considered as contextually similar to the Scottish example. Example v. (5406) was uttered by an Irish woman and recorded in 1992. It contains commas, which indicates pause, and in addition, the transcriber decided to include the discourse marker "like." This case of spoken data, although not a case of improvised description, clearly shows the preferred adjective order.

While it was clear that color consistently occurred before size in both COCA and the BNC, with the only example of the inverse as my hypothesis predicted, I decided to investigate the other semantic categories of the adjectives to be sure it wasn't a unique case. I did searches for a variety of categories and their respective inverses, and found that although the inverses of the preferred order do exist in the corpus, they are fewer in number and can be explained in context. Figure 3.1 below displays the occurrences of certain adjective pairs when searched in COCA, and demonstrates that an order is in fact preferred by speakers in both spoken and written contexts, while the inverse occurs consistently fewer than 15 times in the entire corpus. This is evidence that the preferred order is an underlying order, and the inverses likely have a different semantic meaning as a consequence of switching the order. The semantic differences, as well as other consequences, can be

12

explained when the inverse cases are analyzed in their contextual environments. My theory will attempt to explain why an inverse order may produce a certain semantic shift, and why this shift is to be expected, in section 4.

Figure 3-1. Adjective pairs and their occurrences in COCA

Adjectives		COCA occurrences	
color	**material**		*inverse*
red	brick	217	7
white	cotton	262	4
age	**color**		
old	red	114	1
new	white	157	6
quality	**age**		
beautiful	old	101	2
wild	new	24	16
size	**age**		
big	old	247	17
small	young	8	2
size	**quality**		
big	beautiful	33	13
small	fat	3	0

I would now like to return again to the analysis of "big red" vs. "red big," and discuss the prosodic differences between these cases; that is, if asked to read both cases out loud, what prosodic differences would an average English speaker, selected at random, produce? To answer this question, I created sentences with both case (a) and (b), which are designed to be read aloud in sequence without the reader having seen the paragraph before. This eliminates an improvisation variable, but will control for prosody and hopefully illuminate an underlying distinction between the two cases. Below is the paragraph to be read aloud. I asked a participant to read the paragraph aloud into Praat. The spectrogram for her reading is shown on the next page. I will look specifically at intonation and the temporal spaces between words to determine how the participant may have reacted to the adjective orders.

"And in the morning, I walked down the road to the **big red house** on the corner. I knocked on the door, and a woman wearing a **red big hat** entered. She smiled and said good morning."

13

9.924660
Visible part 10.826644 seconds
Total duration 10.826644 seconds
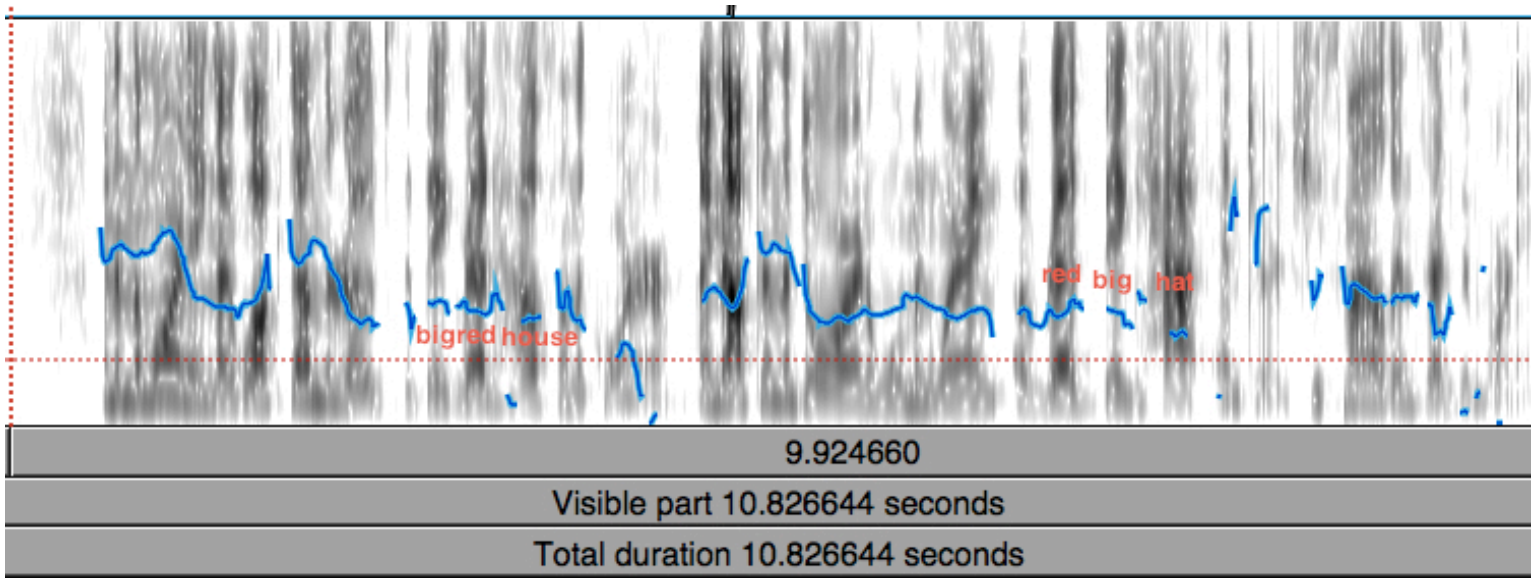
On the above spectrogram, the participant's utterance of "big red house" remains relatively consistent in pitch with each word (pitch is displayed by the blue line). There is a longer pause between red and house than there is between big and red, but it is hard to determine the significance of this. Later in the paragraph, when the participant utters "red big hat," the pitch pattern is quite different. "Red" is at the same basic pitch level that it was at in the utterance of "big red house," which may suggest that certain adjectives have common pitch levels based on their estimated distance from the noun, which would be interesting to investigate. However, in the second case, the adjective "big" drops slightly in pitch when it occurs after "red", creating a pitch contrast. This signifies intent to distinguish the adjectives from one another, which may suggest that the speaker is adjusting for an adjective ordering error by making the adjectives audibly separate entities.

In addition to the creation of pitch contrast, the speaker also pauses for a significantly longer amount of time between adjectives in case (b). This suggests, once again, the speaker's attempt to either adjust for the ordering error or create temporal distance between the two adjectives in an attempt to express their "list-like" nature. In the orthography, we

represent this "list-like" nature using commas, which allow us to freely order adjectives by making them independent syntactic entities that modify the noun equally and simultaneously rather than sequentially according to ordering restrictions.

In conclusion, it is possible to find proof of the prescribed adjective order not only in corpus data of American and British English, but also in the prosodic features of both cases when they are read aloud by an unsuspecting reader. These are two very different supporting cases for the existence of a "correct" order of adjectives in English, which strengthens the argument that a correct order indeed exists.

4. A semantic theory for the description of adjective order in English

A semantic account of adjective ordering requires placing adjectives into semantic categories, which has been done and made available to learners of English by people who have developed grade-school textbooks and instructional materials. Grammarians have worked on this topic before, but adjective order seems to have been taught only in recent years to grade-school students and people acquiring English as a Second Language. This suggests that classifying adjectives into semantic categories is also a rather recent teaching tool for language instruction. The Internet, for example, has numerous sites dedicated to the instruction of English and English grammar which all refer to adjective ordering as "rule-based," claiming that there is a specific order which must be learned in order to have a full grammatical understanding of the English language. According to the British Council[2], a British organization for cultural relations, educational opportunities, and the instruction of

---

[2] http://learnenglish.britishcouncil.org/en/english-grammar/adjectives/order-adjectives

English for children, proper adjective order in English (based on the semantic categories of the adjectives) is as follows, increasing in proximity to the noun from 1-8[3].

1. general opinion
2. specific opinion
3. size
4. shape
5. age
6. color
7. origin (nationality)
8. material
→(noun)

There are issues, however, with the order of the first three categories. The first issue lies in understanding the distinction between "general" and "specific" opinion. What is the difference? According to the British Council, a "general opinion" adjective would be something like "nice" or "interesting," whereas a "specific" opinion would be an adjective like "beautiful" or "curious." What makes "interesting" a more general opinion than "beautiful?" It's also difficult to classify "nice" as a general opinion, when it seems to have semantic variation; it can mean "pleasant", or it can be more specific and mean "kind" or "pretty." The distinction is theoretically difficult to grasp, especially when we attempt to apply this order to real world examples, as I will attempt to do. The second issue is with the relative order of shape and age; in my modified order, I will switch them. Additionally, "size" seems to be incorrectly placed. According to the British Council, the order in 1(a) would be correct, and the order in 1(b) would be incorrect.

1.   (a) the nice beautiful big house
        (b) # the nice big beautiful house

---

[3] This order is modified by the British Council from Grammarian R.M.W. Dixon, who discusses the semantic properties of adjectives and their proper order in his paper, "Where Have all the Adjectives Gone?" *Studies in Language* 1, no. 1 (1977): 19-80.

My intuitions as a native speaker of English (and the intuitions of my English-speaking peers) tell me that 1(b) is actually the correct order. I affirmed these intuitions when I conducted a corpus search for "big beautiful" versus its inverse (see figure 3-1 on page 8), which yielded 33 occurrences of "big beautiful" and only 13 of its inverse. Additionally, "nice" is an interesting case, as it changes its semantic value depending on its relationship to the size adjective. For example, "nice" in "the nice big house" carries the meaning of general appeal, with the house's size is contributing to that general appeal. However, "nice" in "the big nice house" seems to mean something more specific, perhaps in reference to the house's aesthetic appeal. Additionally, "the big nice house" seems to imply that there is a group of aesthetically appealing houses, and the speaker is attempting to identify the big one. Other adjectives, which Cinque (1994) refers to as "operators" behave similarly to adjectives such as "nice." These adjectives include "former" and "alleged," and are able to move more freely within a multiple-adjective NP, depending on the desired semantic outcome. The British Council's order does not include operators.

We therefore have two similar categories, one that is not mentioned by the British Council ("operators") and the other ("specific opinion"), which seems to include a larger variety of adjectives. These two categories seem to have two things in common: a scope-taking quality and an ability to move more freely, making them different from the other categories. With this generalization, as well as the changes I made to the ordering of age, shape, and size, I would like to suggest the following modification to the British Council's order.

1. scope-taking (take all adjectives) (alleged, former, nice)
2. size (big, small, fat, skinny, tall, short)
3. quality (formerly "specific opinion") (beautiful, ugly, silly, *little, *young)
4. age (old, new)
5. shape (square, round, rectangular)
6. color
7. origin (nationality)

8. material
→(noun)

*Adjectives such as "little" and "young" can be a part of multiple categories, although not simltaneously. This will be discussed later.

Let us take this modified order to be the proper underlying order. We then have to explain the semantic shifts that occur when adjectives are switched, as well as account for *why* this order can be taken to be the underlying one. At first glance, the order is seemingly arbitrary, but the above semantic categories can be farther consolidated into larger categories. These larger categories can be thought of as spheres in which the sub-categories lie, with the noun at the core of these spheres. If we begin with the noun and work backwards, we can make observations that can help us to establish our innermost sphere. Firstly, the categories closet to the noun, color, origin, and material, are very rarely found out of order (this can be affirmed by the COCA search). Secondly, these three categories attribute *intrinsic* properties to a noun, that is, they denote properties that are inherent in the object's physical existence (For example, an object cannot exist without being made of something). Lastly, as these three categories increase in distance from the noun, it is apparent that they become less intrinsic. We can therefore take color, origin, and material to make up our innermost sphere containing the adjectives that attribute intrinsic properties to the core noun.

Sphere 1: The intrinsitivity of material and color and the optimization of cognitive efficiency in the restriction of subsets

To explore these observations, let us look at a simple case that makes use of the "material" and "color" categories, "red brick house." Let us first establish that "red brick house" is the preferred order, where "brick red house" is not. If we look at the properties that each of these adjectives attribute to the house, we can see that the material ("brick," in this

case) is a slightly more intrinsic property of a house, or any noun, for that matter. To test the intrinsitivity property of adjectives that attribute "material" to a noun, we can attempt to imagine a house without the property of being made of something. This seems impossible to do, because a physical object (such as a house), by nature, must be made of something. If someone were to ask a group of people to picture a house, one person may picture a wooden house, while another person may picture a brick house, but it would be impossible for either person to picture a material-less house. Therefore, we can consider material an intrinsic property of a physical object. Additionally, both color and material require no additional contextual information to be determined; however, we do require context to determine if something is big, small, beautiful, or ugly.

If we are to be in keeping with the order I have suggested and the observation that the adjectives denote decreasingly intrinsic properties as they get farther from the noun, the question then becomes, "What makes color a *less* intrinsic property than material?" If we use the same diagnostic above and attempt to picture a "colorless" house, some would argue that it is impossible to do so. However, it seems more plausible to imagine a "colorless" house than a "material-less" house. "Colorless" is a word in English that is used to describe something that is lacking color, so it seems that we are in fact capable detaching the property of color from an object successfully without losing our entire concept of it. Detaching the property of material from a physical object seems much less plausible. In fact, if I were to ask you to picture a "colorless house," it is likely you picture a glass house. The house is still made of something, but it does not have a color.

In an attempt to explain *why* to ordering rules are so strict within this first sphere, let us think of the order of the adjectives as the order in which we restrict the noun's properties into subsets, starting from the adjective closest to the noun and working outward. If we do this, then when we conceptualize of the NP "the red brick house," the order in which we

restrict the set of houses is indicative of the order in which we attribute properties to the noun, and thus represents the order of the adjectives. We begin with the noun, which denotes the set of all houses. Then we begin to restrict this set with increasingly non-intrinsic properties. First, we create a subset of houses that are made of brick. Then, we take brick houses of varying colors and isolate the red ones. By doing this, we are being more efficient in how we describe the object we wish to refer to. If we were to first isolate the red houses, we would not only conceptualize houses that are red, but also be cognitively required to assign some material (or materials) to them, since it is impossible to conceptualize of a material-less house. In short, it seems much more plausible for our brains to isolate a set of colorless brick houses than a set of red material-less houses. To be maximally *cognitively efficient*, it makes sense to begin by making a subset of brick houses of no particular color, and then make yet another subset of the red ones. When we reverse the order, the result is a less cognitively efficient concept; this is likely why reversing the order within sphere 1 is so rare.

When the order *is* reversed, however, what are the semantic consequences? Additionally, how can we conceptualize of a "brick red house" without being cognitively inefficient? If we think of a context in which saying "the brick red house" is acceptable, it is required that the discourse context makes it previously clear that we were already discussing a set of red houses. It seems plausible for someone to say, "We passed several red houses on our tour of Brussels. My favorite was the brick red house." Basically, saying the "brick red house" in a grammatically correct fashion simply requires that you begin with a set of red houses before attributing other properties, such as material, to the set. In doing this, you also avoid the problem of cognitive inefficiency. Diagrams 1, 2, and 3 below demonstrate this process. In diagram 1, the context-free underlying order, which is also the most cognitively efficient, is depicted. In diagram 2, the context-free reversed order is

shown, demonstrating the cognitive inefficiency of conceptualizing of the adjectives in that order. Lastly, diagram 3 demonstrates how in context, this problem can be fixed by altering the starting point form which we begin restricting subsets of properties. For the purpose of including another contrastive semantic category from another sphere, the phrase shown below is "big red brick house."

1 MOST EFFICIENT

2 NOT EFFICIENT

3 FIXING 2
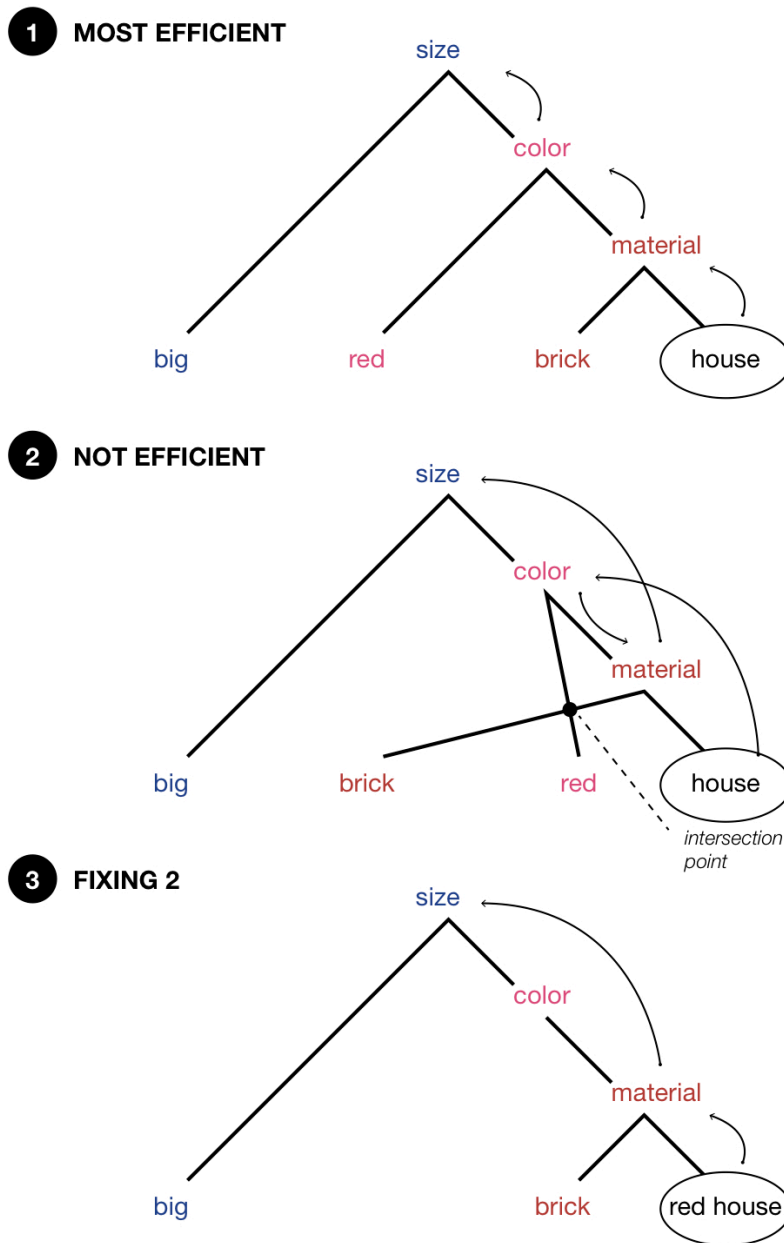
Sphere 2: Age and shape -- Internally comparative adjectives and internal relativity

Color is the final semantic subcategory in Sphere 1, making age and shape the next two semantic categories as we work outward from the noun. Age and shape, in terms of their relationship to sphere 1, are semantically less intrinsic than properties such as material and color, but this is not what places them in a unique sphere. What distinguishes age and shape from the categories in sphere 1 is their ability to be *internally comparative.* Unlike material, origin, and color, age and shape can be comparatives; for example, we are capable of saying that one thing is "rounder" or "older" in comparison to another thing, or to itself at a previous point in time. We are not capable, however, of saying something is "more brick." We can't really say that something is "more French" than something else; if we do, we mean French-like in quality rather than the property of being from the country of France.

Something that distinguishes age and color from the remaining semantic categories is that they are not semantically deictic. This is what I have intended to convey by describing these adjective as *internally* comparative and *internally* relative; such adjectives do not require information from an external context to be determined as true or false, and additionally, they require no information about the speaker's beliefs. For example, if a man were to be alone in an empty room with no other people or things to compare him to, one would still be able to assess his age. Signs of age can exist independent of any context, allowing us to be able to say a man is "old" without needing any comparison class. The same is true for shape. If a table were to be sitting alone in an empty room, we would still be able to determine if it were square, rectangular, or round.

Unlike in sphere 1, sphere 2 adjectives do not possess the same necessary intrinsic properties and it is therefore not obvious that reversing the underlying order within sphere 2 is cognitively inefficient. For example, saying "the old round table" is not necessarily more cognitively efficient than saying the "round old table," it just elicits a different cognitive

process in terms of the sequence of which we denote the set of properties of the table. How do we know then, that age comes before shape in an underlying pre-nominal adjective sequence? Speaker intuitions tell us that "it depends on what you are trying to say," but a COCA search for "old square (noun)" yields 9 occurrences, whereas the inverse yields 1. This shows us that whether speakers are aware of it or not, they are demonstrating a slight preference for saying age before shape, so this must be what feels natural.

How then, do we explain this preference? Because sphere 2 adjectives are characterized by their internal relativity, I believe that it is logical to assume that like the intinsitivity property in sphere 1, the internal relativity property is also scalar. It appears that as we get farther from the noun, the adjectives decrease in intrinsitivity and increase in relativity. Shape is a less relative property than age, because it requires no temporal scale. Age is a property that is impossible to determine without the presence of a time, but it is internally relative (rather than externally relative) because it relates an object to the quantity of time that the object itself has been in existence. Determining shape requires no such information, and is therefore a more intrinsic property, placing it closer to the noun than age.

If we were to cross the sphere boundary, we would see that reversing the order seems much less natural. If we take a sphere 1 adjective, for example, "red," and inverse it with a sphere 2 adjective, "old," the result is a non-preferred NP. Saying the "old red table" is clearly preferred to saying the "red old table," and this is likely due to the fact that the reversed order involves placing a sphere 1 (intrinsic) adjective farther from the noun that a sphere 2 (internally comparative) adjective. The semantic consequences of reversing the sphere order are similar to the semantic consequences of reversing the order of the subcategories within sphere 1. If reversing the order, it is necessary to have a context in which the starting set is pre-established before additional properties restrict it. We have thus far established the following order of adjectives before a noun.

((Sphere 2 AGE > SHAPE) (Sphere 1 COLOR > ORIGIN > MATERIAL) (noun))

Sphere 3: Quality and size -- Externally comparative adjectives and speaker relativity

Continuing outward from the noun, the next two semantic categories are quality (what the British Council referred to as "specific opinion") and size. I made the decision to eliminate matters of opinion; however, adjectives that denote "quality" can, in fact, be words that denote the opinion of the speaker. Unlike the adjectives in the first two spheres, quality words such as "beautiful," "funny," and "interesting," require a speaker who has access to an external comparative class. Words such as this are difficult to define without also understanding their opposites. How can a speaker understand if something is beautiful, funny, or interesting, without first being informed by the context of what makes something beautiful, funny, or interesting and what does not? To a large degree, the definition of quality adjectives are fixed by societal norms and are established a community of speakers; for example, what is "beautiful" to Americans is very different than what is "beautiful" to other cultures. Quality adjectives, to a degree, are a matter of the speaker's relationship to the object, and the speaker requires an external context to determine the quality of that object. This need for an external context is what makes these adjectives *externally comparative* rather than internally comparative, making them different then sphere 1 or sphere 2 adjectives. Additionally, we can say that these sphere 3 adjectives are *speaker relative*; that is, they depend on the context of the speaker and rely on the speaker's observations to be determined as true or false.

Even more externally comparative (and therefore speaker relative) than quality is size. Let us return briefly to the "alone in an empty room" test. If an observer saw an object alone in a colorless, empty room, he would be able to assess its material, color, and shape

without any other information. Age would be more difficult to determine, since the observer would need some reference to time and the object's internal "lifespan." Assessments of quality would require the speaker to contribute some information from his own external context from outside of the empty room, and size would be nearly impossible to assess. Without another object next to it, or an idea of how "big" the room is, the size of that object would not be able to be determined. The speaker would need not only his own information, but any assessment of size would *require* an external comparison class. If that object were an elephant or a mouse, we would have no way of knowing if either creature were big or small without having other creatures to compare it to or the speaker itself. We can therefore say that size, like quality, is also externally comparative and speaker relative. This would result in it being the most distant adjective from the noun in almost every context-free case. According to COCA, size is almost always preferred in the external-most position as possible from a noun. To assess an object's size prior to assessing its other more intrinsic properties would be cognitively inefficient.

This leaves us with a final explanation for the underlying order of pre-nominal adjectives. The noun is preceded by those adjectives that contribute the most intrinsic properties, and as we work outward from the noun, the adjectives become increasingly comparative, first internally (requiring no or minimal outside information), and then externally (requiring information from an external context). Additionally, intrinsitivity decreases as we get farther from the noun. The closer the adjectives are to the noun, the less they have to do with the speaker and his or her external world and the more they have to do with the intrinsic properties of the noun's existence. The diagram on the next page displays this relationship, while also demonstrating the order of the spheres and their subcategories. The final category, the scope-taking adjectives, will be explained in the next section.

1   **UNDERLYING ADJECTIVE ORDER BY SEMANTIC CATEGORY**
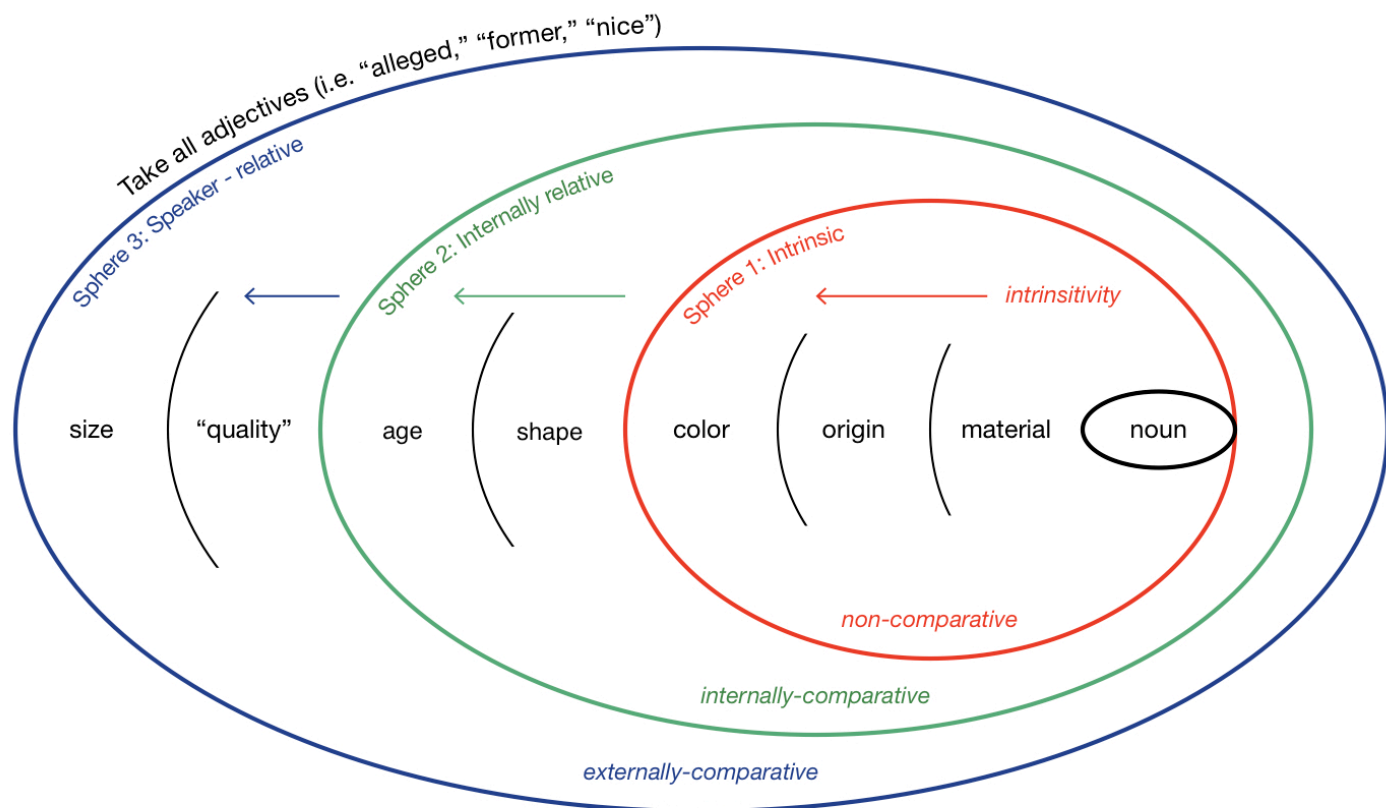
Outside of the spheres: the scope-taking adjectives

The final subcategory, which the British Council referred to as the "general opinion" category, lies outside of the spheres. I have chosen to call this category the "scope-taking" category, because adjectives of this sort have the capability of moving freely between sub-categories depending on the intended meaning. These adjectives, when context-free, can usually be found as far from the noun as possible, since they are capable of modifying the

entire NP and all of the adjectives within it. In some ways, they are as "adverbial" as possible without being actual adverbs. "Former and alleged" fit well into this category but are operators, and have a stricter semantic value. "The former obese mayor" means that he is one in a series of obese mayors, and he is the previous one in the set to which the speaker is referring. However, "the obese former mayor" means that there are a series of former mayors, and he is the obese one. Non-operator scope-taking adjectives operate a similar way, although they have the ability to shift semantically, as demonstrated below.

(1)   the **nice** big long French baguette
(2)   the big **nice** long French baguette
(3)   the big long **nice** French baguette

"Nice," when in the left-most position, almost has an adverbial-like behavior, in that it seems to modify not just the baguette, but also all of its other qualities. Anywhere it goes, it seems to modify all of the adjectives that come after it, making it scope-taking. The farther it gets from the noun, however, the less intrinsic the properties that it denotes become. In its most distant position, "nice" has a more general externally relative meaning, in that it compares the baguette not just to other baguettes, but other objects in the speaker's context. As it approaches the noun, it takes on a more noun-relative scope, assessing its aesthetic quality. Many "quality" adjectives can be removed from their subcategory and be brought outside of the sphere for the purpose of comparing the noun to a broader external comparative class. In some cases, doing this actually changes the semantic value of the adjective. "Interesting," and "curious" are some of these, as demonstrated below, when interacting with the other categories from spheres 1, 2, and 3. In the below examples, I have used size, a sphere 3 category.

(4)   the interesting small paper
(5)   the small interesting paper

In case (4), the paper is small, and is interesting in comparison to other things in the world. In case (5), the paper is interesting for a paper, and is small in size.

(6)   the curious little man
(7)   the little curious man

Between cases (6) and (7), curious has a more distinct semantic difference. In case (6), curious means odd or bizarre, which is an assessment that relates the man to a broader class of things in the world. Case (7) prompts a different understanding of curious, meaning inquisitive or questioning, which are properties that are more person-relative. He is still small in stature in both cases; it is only the semantic value of "curious" that changes, suggesting that "curious" is the one switching semantic categories, rather than the size adjective.

In some cases, an inverse order is actually the convention, because certain adjectives have come to take on a popularized different semantic value depending on the context in which they appear. If searched in a corpus, these conventional adjective pairs would likely occur with more frequency. These cases, on the surface, appear to be exceptions to the underling AO, but in reality, such an inverse order actually results in a semantic shift, making the out-of-place adjective more fitting with a different semantic category. Examples of this usually involve the size adjective "little," because "little" is used to mean more than just small in size; it has the additional meaning of "cutesy" or "child-like," and appears to only occur in the inverse position in cases such as "pretty little liars," or "silly little boy," where it does not entirely mean physically small in size. Another case that demonstrates a semantic shift is the use of "old" in "big old dog." This case, unlike the previous case, is not a surface violation of the underlying AO, but "old" does display a conventional semantic shift from the "age" category to the "quality" category. When it occurs after "big," old does not mean old in age, but tends to retain a more quality-oriented meaning of familiarity and strength. If we contrast this with "small old woman," we see "old" shift back

into the age category. We will be able to explore semantic shifts such as this with regards to placement cross-linguistically when we attempt to apply the theory to Italian.

Tricky cases: using multiple adjectives from the same subcategory

The next question becomes, "Are there ordering restrictions when we use multiple adjectives from the same subcategory?" First it is important to make note of the fact that we can only use multiple adjectives from the same category if they are age, quality, or size adjectives. It does not make sense to say "the red blue car," the brick glass house" or the "rectangular circular table." To summarize, sphere 1 subcategories do not have the ability to be repeated, but sphere 2 and sphere 3 adjectives do. Let us look at the effects of this in the following examples below.

Age:
(8) my new old car
MEANING: The car is old, but I have recently acquired it.
(9) my old new car
MEANING: I once had a car that was new, and now I do not.

(10) the young old man
MEANING: The man is old, but young for an old man.
(11) the old young man
MEANING: The man is young, but old for a young man.

Size:
(12) the big small dog
MEANING: The dog is small, but big for a small dog.
(13) the small big dog
MEANING: The dog is big, but small for a big dog.

Quality:
(14) the lazy crazy man
MEANING: the man is mentally disturbed, and generally lacks motivation
(15) the crazy lazy man
MEANING: the man is by nature unmotivated, and does things that are bizarre, but is not necessarily mentally disturbed.

In the above examples, there does not appear to be a preferred AO, therefore, there are no apparent AO restrictions on the use of two adjectives from the same subcategory.

There are, however, semantic differences between the pairs above, which can be summarized as such: the adjective closest to the noun attributes internal properties to the noun, and the adjective farther away from the noun distinguishes it from a broader external context or set of nouns. This is consistent with my theory that as an adjective gets farther from the noun, intrinsitivity decreases and speaker relativity increases. The adjectives farther from the noun are externally comparative, while the adjectives closer to the noun are internally comparative. In sphere 1, the adjectives are not comparative at all.

5. Application of the theory cross-linguistically

Sakha (Yakut)

While I have very limited resources with regards to the Sakha language, I have spent a few months studying it rather closely and have had interactions with a native speaker. Sakha, or Yakut, is a Turkic language with around 360,000 native speakers, spoken in the Sakha Republic of Russia in northeastern Siberia. It is an agglutinating language with a SVO word order. It demonstrates vowel harmony, and has a very rich case system, but no grammatical gender.[4] In my study of the language, I have seen adjectives primarily in predicative contexts. However, when I asked a native speaker to tell me how to say, "small white clouds" she told me the following.

1. Кыра үрүҥ былыбыпар
   /kɯra yryŋ bɯlɯbɯpar/
   "small white clouds"

The order demonstrated above is the same as it would be in English (SIZE > COLOR > noun). To make sure that this order is actually also preferred in Sakha, I asked the native speaker to tell me what it would mean if you said "white small clouds (in Sakha)," instead. She said it to herself a few times, and immediately concluded that "No, it sounds wrong,"

---

[4] http://en.wikipedia.org/wiki/Sakha_language#cite_ref-Krueger_3-0

just as a native speaker of English would react to such a change. While this is only one native speaker's opinion and one example, it does provide a small amount of evidence that the same semantic restrictions in place in English have the ability to be in place in other languages, even if the language has absolutely no genetic relationship to or history of contact with English. This suggests that these semantic restrictions are underlying.

Italian

Italian is a Romance language, and was the primary language of study (in comparison to English) for Cinque (1994) when he developed his internal N-raising theory. Italian has both pre and post-nominal adjectives, but as we will see, they exhibit similar semantic shifts that English adjectives do depending on our placement. One thing that makes Italian interesting is its "mirror image" post-nominal ordering. For example, I came across the following text on a wine bottle. The bottle was imported from Italy and consequently translated from Italian to English.

The English on the label reads "Red Dry Wine," which demonstrates an incorrect adjective order. In Italian, this order (if read from left to right) is technically correct. Let us compare the correct order in English with the correct order in Italian.

    a.   English: dry red **wine**
    b.   Italian: **vino** rosso secco
English gloss: wine red dry

In correct Italian, we see that "red dry," or "rosso secco," is actually the correct order when literally translated from Italian to English and read from right to left; however, the adjectives follow the noun rather than preceding it. If we then consider the adjectives' placements in relation to the noun in terms of their distances from it, it appears that in this case, we have contrasting headedness rather than contrasting adjective order. More specifically, we see a head-final noun phrase (in English) versus a head-initial noun phrase (in Italian). This difference would result in a "mirrored" order for all words in the noun phrase when translating from English to Italian. We can then conclude that the placement of the adjectives in relation to the noun in a noun phrase is in fact identical in English and Italian, with the word "red" (denoting the sphere 1 color adjective) occurring closest to the noun, and the word for "dry" (denoting the sphere 3 quality adjective) occurring farther away.

If this is indeed the case, the wine bottle above demonstrates how the translator made an error such as this. While the translator seems to understand the difference in noun placement relative to the adjectives, he or she did not understand the difference in headedness, which would imply a "reversal" of every element in the NP, not just the movement of the noun. The fact that the term denoting color ("red," in this case) occurs closest to the noun in both English and Italian supports the idea that N-movement may not be necessary to explain the different in adjective ordering. We do not have to assume an underlying position of the noun to see that the adjectives have the same relationship to it as they would in English. Let us consider another example, from Cique (1994).

a.　　La sola possible invasione romana della Tracia
　　　the only possible　invasion　Roman　of the Tracia
　　　"the only possible Roman invasion of Tracia"

b.　　# La sola possible romana invasione della Tracia
　　　 the　only　possible roman　invasion of the Tracia

In this example, we see that while "sola (only)" and "possible" occur before the noun, while "romana (Roman)" must appear after it. This does not necessarily have to be a due of noun movement. I would like to suggest that it is actually the post-nominal placement of spheres 1, 2, and 3, that make Italian different than English syntactically. For example, the sphere 1 adjective "roman," designating origin, occurs after the noun, while the one that takes all adjectives, "possible," occurs before. The case of "sola" is a little more complicated. Although I did not focus on quantifying adjectives, I believe that "sola" would fall into such a category. Traditionally, quantifiers are analyzed as being in the most distant position, and since they are not attributive, they do not have a place in my theory. If I were to hypothesize, however, I would guess that quantifiers, much like determiners, would remain pre-nominal and distant form the noun. In addition, they would not be part of any particular sphere. We can therefore conclude form our data that in Italian, adjectives in spheres 1, 2, and 3 are post-nominal, where they are analyzed in mirror image order from English, while the determiner, numerals, and scope-taking adjectives are found consistently pre-nominally.

Hebrew

Like Italian, Hebrew exhibits mirror-image order for spheres 1, 2, and 3, as exemplified in the following example from Willis (2006). Hebrew also imbeds the determiner into its nouns and its adjectives, and is a right-to-left language, which makes it unlike any language we have analyzed thus far. Despite these differences, it still demonstrates an order that is in keeping with my theory. The example below Willis's example (example b) was given to me by a native speaker of Hebrew. I was unable to come up with a successful Romanization, since I do not trust my knowledge of transcribing from the Hebrew to theLatin alphabet. However, the native speaker did provide me with a literal translation.

a. ha-mexonit ha-amerika'it ha-aduma
    the-car the-American the-red
    "the red American car"

b. והישן היפה הגדול הבית
    The-house the-big the-beautiful and the-old

With these adjectives, we see the same phenomenon that we see with Italian, where adjectives in the first three spheres are post-nominal and demonstrate mirror image ordering to English. Unfortunately, I do not have data containing adjectives from other spheres, so I am unaware of how Hebrew treats scope-taking adjectives and quantifiers. It would be interesting to investigate, especially considering that Hebrew embeds its determiners.

Welsh

Among all the languages I have examined in the search to find one that does not conform to the ordering rules, Welsh (and other Celtic languages) has been the only one that cannot survive my theory without some form of noun movement. Noun movement in Celtic has been a very established topic in syntax (Willis 2006), however, in his paper, Willis seeks to disprove it. The data in Welsh shows how a simple noun-movement theory is seemingly inefficient; while noun phrases are noun-initial, as in Hebrew, the adjectives alternate between demonstrating a mirror-image order and a non-mirror image order. The example below, from Willis's paper, demonstrates a noun-initial NP that demonstrates English AO rather than mirror-image AO.

a. cwpan mawr gwyrdd Sieineaidd
    cup big green Chinese
    "a large green Chinese cup"

As you can see, the example demonstrates how n-movement would be the easiest way to account for this. It is important to note, however, that the adjectives remain in the same order *relative* to one another as they do in all other languages I have discussed. This further motivates n-movement for this case and makes it the most understandable account. However, in his paper, Willis provides ample data demonstrating that Welsh does in fact behave with mirror image order. One example of this is below.

      b. acen Rwsieg ysgafn.
        accent Russian mild
        "a mild Russian accent"

In this case, Welsh demonstrates a mirror-image order, with origin closer to the noun, and quality further away. If n-movement were to be occurring here, the underlying form prior to movement would violate the semantic restrictions placed on adjective placement as outlined in my theory. It then seems that the only conclusion that we can draw about Welsh and the Celtic languages is that they contain both mirror image ordering and n-movement. The different contexts in which the two occur would be interesting to investigate, and perhaps a topic of future study.

6. Conclusions

In this paper, I have established a uniquely semantic theory for the AO restrictions that are in place in English. I demonstrated that these restrictions can be proved using corpus data, as well as prosodic data and the ways in which native speakers react to inversely ordered adjectives. I discussed the other work that has been done on the topic, citing research from many disciplines spanning across several decades. My theory discusses the "cognitive efficiency" of optimally ordered adjectives, where those adjectives that are most intrinsic occur closest to the noun and those that relate the noun either to

another point in time, another object, or the external world, occur farther away. I discuss the order in which adjectives are placed as the order in which they restrict the set of properties of a given noun.

In my theory, the semantic subcategories of adjectives (shape, color, size, etc.) are further categorized into broader "spheres," where sphere 1 contains adjectives that are decreasing in intrinsitivity as they increase in distance from a noun. Sphere 1 adjectives are also less likely to be found out of order. Sphere 2 adjectives are internally relative, and internally comparative, and when found out of order, result in a semantic shift from the underlying order. Sphere 3 adjectives are externally comparative, and involve the speaker's opinion and perspective. Sphere 3 adjectives are more likely to be found in inverse position, and I discuss the semantic shifts that occur when this is done. Lastly, I discuss the category of scope-taking adjectives, which move more freely and have a more "adverbial" quality. Scope-taking adjectives are capable of taking the noun and all of the adjectives that modify it and modify all of them further. I briefly discussed how adjectives from identical semantic categories behave when they interact and the semantic differences between ordered pairs and their inverses.

I attempted to take my theory about the semantic restrictions of AO and apply it across languages, which was successful in some languages, but not as much so in Welsh. My cross-linguistic analysis is clearly very limited; I did not have extensive data or access to many native speakers, so naturally, it was very difficult to test the semantic effects of inversing the underlying AO in the other languages, as I was able to do in English. In the future, I believe this would be an interesting study to do.

Additionally, I would be interested in exploring in more detail the relationship between scope-taking adjectives and the adjectives within the other spheres; these adjectives seem to behave in an interesting way that is easy to understand in English, but may be having an

affect on how other languages treat their AO. The free nature of these adjectives makes them naturally more likely to vary cross-linguistically, and I would be very interested to analyze this behavior.

In general, however, one thing that is never seen cross-linguistically, despite the controversy about the placement of the noun, is a clear violation of the underlying order I have established. The adjectives consistently remain in the same order in relation to one another, whether in mirror image or not, which is a very robust finding. Having consistent cross-linguistic relative adjective order as predicted by a semantic theory may be evidence that a semantic theory will also suffice in explaining cross-linguistic variation among the syntactic relationships between the placement of a noun and the adjectives that modify it. I hope to have given some insight into this process during the brief application of my theory in Sakha, Hebrew, Italian, and Welsh.

References

Cinque, G. 1994. *On the evidence for partial N movement in the Romance DP*. In G. Cinque, J. Koster, J.-Y. Pollock, L. Rizzi and R. Zanuttini (eds.), Paths Toward Universal Grammar, 85–110. Georgetown: Georgetown University Press.

Danks, J. H., & Glucksberg, S. 1971. *Psychological scaling of adjective orders.* Journal of Verbal Learning & Verbal Behavior, 1.10,63-67.

Halliday, M.A.K. 1994. *An Introduction to Functional Grammar (2nd Ed).* London: Edward Arnold.

Kemmerer, D. 2000. *Selective impairment of knowelsge underlying pronominal adjective order: Evidence for the autonomy of grammatical semantics.* Journal of Neurolinguistics. 13-57-82.

Kemmerer, D., Weber-Fox, C., Price, K., Zdansczyk, C., & Way, H. 2007. *Big brown dog or Brown big dog? An electrophysiological study of semantic constraints on adjective order*. Brain and Language, 100, 238-256

Martin, J. E. 1969. *Semantic determinants of preferred adjective order*. Journal of Verbal Learning & Verbal Behavior, Vol 8(6), 697-704. doi

Sproat R. and C. Shih. 1991. "The Cross-linguistic Distribution of adjective ordering restrictions". In C. Georgopoulos and R. Ishihara (eds.), Interdisciplinary Approaches to Language. Essays in Honor of S.-Y. Kuroda, 565-593. Dordrecht: Kluwer Academic Publishers.

Willis, David. 2006. Against N-raising and NP-raising analyses of Welsh noun phrases (abstract) *Lingua* 116: 1807-39.

Acknowledgements