

On the grammatical source of adjective ordering preferences*

Gregory Scontras
University of California, Irvine

Judith Degen
Stanford University

Noah D. Goodman
Stanford University

Submitted 2018-04-06 / First decision 2018-06-11 / Revision received 2018-09-01 /
Second decision 2018-09-18 / Revision received 2018-11-12 / Accepted 2018-11-14 /
Final version received 2018-11-26 / Published 2019-11-13 / Final typesetting 2023-04-
18

Abstract Scontras, Degen & Goodman (2017) present experimental evidence demonstrating that the best predictor of adjective ordering preferences in the English noun phrase is the subjectivity of the property named by any given adjective: less subjective adjectives are preferred linearly closer to the nouns they modify. The current work builds on this empirical finding by proposing that the reason subjectivity predicts adjective ordering preferences has to do with the hierarchical structure of nominal modification. Adjectives that are linearly closer to the modified noun are often structurally closer, composing with the noun before adjectives that are farther away. Pressures from successful reference resolution dictate that less subjective, more useful adjectives contribute their meaning to the resulting nominal earlier, in an attempt to more effectively limit the reference search space.

Keywords: adjective ordering, subjectivity, hierarchical structure, modification, reference resolution

1 Introduction

Adjective ordering preferences determine the relative order of adjectives in multi-adjective strings. Such preferences dictate that *small brown card-*

* We gratefully acknowledge the feedback we received on earlier drafts from Kjell Johan Sæbø and two anonymous reviewers. Special thanks go to Michael Franke for very helpful discussions.

board box sounds much more natural than *brown cardboard small box*, or any other ordering of the adjectives. These preferences are robustly attested, not only in English, but in a host of unrelated languages (see, among others, Dixon 1982). Remarkably, the *same* preferences surface in each case. Even more remarkably, in post-nominal languages where adjectives follow modified nouns, the preferences are the mirror image of what are found in pre-nominal languages like English; at issue is the relative distance of an adjective from the noun it modifies.

Given their stability within and across languages, a glaring question presents itself: what factors determine these robust preferences? Answers to this question stand to inform not only the preferences, but also the psychological and grammatical systems from which these preferences emerge. For this reason, adjective ordering preferences have been the subject of targeted inquiry since Sweet (1898) wrote about them over a century ago. Hypotheses abound, ranging from the psychological (e.g., Whorf 1945, Martin 1969a) to the grammatical (e.g., Cinque 1994, McNally & Boleda 2004, Truswell 2009). Still, significant progress has proven elusive, owing to the complex empirical work required to test these hypotheses.

Recently, Scontras, Degen & Goodman 2017 brought behavioral and corpus data to bear on the question of adjective ordering. Distilling the proposals that preceded them, Scontras, Degen & Goodman advanced the hypothesis that property subjectivity determines the relative order of adjectives in multi-adjective strings, such that less subjective adjectives occur linearly closer to the nouns they modify (see also Hetzron 1978, Tucker 1998, Hill 2012). In *the small brown cardboard box*, *cardboard* is less subjective than *brown* or *small*, so *cardboard* is preferred closer to the modified noun.

With strong empirical footing for the factor determining ordering preferences, the current paper addresses the question of precisely *why* subjectivity should play such a central role in adjectival modification. Section 2 reviews the empirical methodology and findings from Scontras, Degen & Goodman 2017. Section 3 explores potential explanations of the empirical findings. Section 4 offers a proposal tying subjectivity to reference resolution and the hierarchical structure of adjectival modification. Section 5 concludes.

2 Subjectivity predicts adjective ordering preferences

To identify the factors at play in adjective ordering preferences, we must first determine what the preferences are that need explaining. To that end,

Scontras, Degen & Goodman 2017 established a behavioral measure of ordering preferences. Experimental participants indicated the preferred ordering of adjectives in adjective-adjective-noun strings (e.g., *the small brown chair* vs. *the brown small chair*), which yielded a single preferred-distance measure for each adjective tested. The authors evaluated their behavioral proximity measure against naturalistic productions from English corpora. Finding an extremely strong correlation between the behavioral measure and the corpus counts ($r^2 = .83$), the authors concluded that naïve speakers have reliable and robust adjective ordering preferences, and that the behavioral measure faithfully captured these preferences. The authors then shifted their focus to the aspect of adjective meaning that they hypothesized best predicted ordering preferences: subjectivity.

Inspired by various proposals about aspects of adjective meaning explaining their relative order in multi-adjective strings, Scontras, Degen & Goodman distilled past proposals into the intuitive psychological construct of subjectivity. Crucially, the authors *operationalized* their subjectivity hypothesis as a behavioral measure for the purpose of empirical testing. Adjective subjectivity was measured by asking participants how “subjective” a given adjective was. These raw “subjectivity” scores were evaluated against a potentially more ecologically valid method: faultless disagreement (e.g., Kölbl 2004, MacFarlane 2014). To the extent that two speakers can disagree about a given property for an object without one speaker necessarily being wrong (e.g., disagreeing about whether or not a box counts as small), the property admits that degree of faultless disagreement, which stands proxy for the adjective’s subjectivity. Finding extremely high correlations between the “subjectivity” scores and the faultless disagreement measure ($r^2 = .91$), Scontras, Degen & Goodman concluded that both measures successfully capture adjective subjectivity.

It bears noting that many factors can contribute to the perceived subjectivity of an adjective, including semantic notions typically thought of as vagueness (e.g., red by which standard?), evaluativity (e.g., beautiful according to whom?), or relativity/context dependence (e.g., large compared to what?). Moreover, there exist a variety of semantic theories designed to account for these specific notions (e.g., the supervaluations of Kamp & Partee 1995, the perspectives of Kölbl 2002, or the judges of Lasersohn 2005). There are also various formal notions of the subjective vs. objective distinction defined in terms of judges (Sæbø 2009), counterstances (Kennedy & Willer 2016), outlooks (Coppock 2018), etc. For our purposes and the pur-

poses of Scontras, Degen & Goodman’s study, the semantic source of subjectivity runs orthogonal to the simple fact that speakers have stable estimates of subjectivity operationalized via faultless disagreement. In other words, whatever its source, language users recognize that certain adjectives can lead more frequently to cases of misalignment where people might (faultlessly) disagree about the set of things picked out by a given adjective. It is this notion that we refer to with the label “subjectivity”.

With clear estimates of adjective subjectivity and of the preferences themselves, the authors then tested the predictive power of subjectivity in explaining ordering preferences. Scontras, Degen & Goodman found that an adjective’s semantics does predict its distance from the nouns it modifies, with subjectivity scores accounting for nearly all of the variance in the ordering preference data. Moreover, preference strength increased with the subjectivity differential. To get a clearer picture of the *relative* success of their subjectivity hypothesis, the authors then compared the predictions of subjectivity against operationalizations of competing proposals: adjective inherentness (whereby adjectives with more “essential” meanings occur closer to modified nouns; e.g., [Whorf 1945](#)), intersective vs. subsective modification (whereby intersective modifiers compose first in the hierarchical structure of nominals; [Truswell 2009](#)), and concept formability (whereby adjectives that form complex, idiomatic concepts compose early; e.g. [McNally & Boleda 2004](#)). In each case, subjectivity continued to be a better predictor of ordering preferences.

3 Why subjectivity?

Finding subjectivity to be a reliable and robust predictor, our task now is to explain why subjectivity should determine adjective ordering preferences: why should less subjective adjectives be preferred linearly closer to the nouns they modify? Scontras, Degen & Goodman hint at an answer in the discussion of their results, namely pressure from successful reference resolution. Before reviewing their discussion, it will be useful to first consider the range of possible answers to this *why* question. In the process, we also establish desiderata for successful answers.

To begin, we might propose that the observed subjectivity gradient emerges from a rigid syntax of adjectival modification: adjectives inhabit specialized syntactic projections depending on their semantic class (e.g., Color Phrase for color adjectives, Shape Phrase for shape adjectives, etc.; [Cinque 1994](#), [Scott 2002](#)), and these projections happen to order in a way that

tracks subjectivity. While a cartographic approach along these lines might help to explain the observed behavior, it leaves unanswered the question of *why* subjectivity should matter in the ordering of adjectival projections. Also problematic is the rigidity introduced by a syntax that allows only one ordering for any string of adjectives. This rigidity predicts categorical ordering preferences, yet Scontras, Degen & Goodman observed graded judgments that track differential subjectivity. Thus, a cartographic syntax appears to be a nonstarter for explaining why subjectivity should predict ordering preferences.

Shifting our sights to psychological explanations, we might try to account for subjectivity by appealing to the relative salience of properties for the nouns they modify. Properties that are more salient — more *inherent* to the objects described by the noun (e.g., Whorf 1945) — ought to be more accessible in the construction of nominals. An account based on the accessibility of adjectives during the on-line construction of nominal phrases stands to extend straightforwardly to languages with post-nominal adjectives where the preferences are preserved in the reverse: phrases are built outward from their heads, so more accessible adjectives occur linearly closer to the head of the nominal construction (i.e., the noun). The leap that must be made links subjectivity to inherentness and thereby to accessibility. Unfortunately, this leap appears untenable. Scontras, Degen & Goodman measured adjective inherentness and compared its predictions with those of subjectivity. Whereas subjectivity accounted for at least 75% of the variance in the ordering preferences, inherentness accounted for 0%. While an explanation in terms of adjective accessibility might still prove promising, the implementation via inherentness lacks empirical support.

Rather than inherentness, we might try tying adjective accessibility to adjective frequency, such that more frequent adjectives are more accessible during the hierarchical construction of nominal phrases. To account for the role of subjectivity, we would expect more frequent — and thus more accessible — adjectives to have lower subjectivity scores. Scontras, Degen & Goodman investigated the role of adjective frequency in ordering preferences, finding it to be a significant predictor. However, frequency applies pressure in the direction opposite to what we have been considering: in English, more frequent adjectives occur *farther* from the noun because they occur *linearly* early in multi-adjective strings (cf. Wulff 2003). Moreover, the authors found that subjectivity continued to explain significant variance in the preferences over and above adjective frequency. While frequency likely contributes to the

relative accessibility of adjectives, its contribution is separate from that of subjectivity; the two forces work in tandem and in orthogonal directions.

We find perhaps the most thought-through version of the psychological accessibility hypothesis in [Martin's \(1969b\)](#) experimental investigations. Inspired by prior results demonstrating the predictive power of “definiteness of denotation” in adjective ordering preferences ([Martin 1969a](#)), Martin set out to test the hypothesis that adjectives occurring linearly closer to nouns indeed are more accessible. Participants completed a series of elicited production tasks in which they observed visual arrays of objects and named specific properties of the objects they saw (e.g., size vs. color). There were two versions of the experiments: one for English speakers, and another for speakers of Indonesian, a post-nominal language with the mirror image of the English preferences. By measuring production latencies, Martin discovered that adjectives preferred linearly closer to nouns are produced more quickly. From this he concluded that adjectives closer to the noun are more accessible than adjectives farther away.

Before accepting Martin's results as unambiguous support for an adjective accessibility hypothesis, we must confront two issues. First, production latencies in context likely depend on more than the relative accessibility of words from memory; how can we be sure that the observed differences in latencies did not derive from low-level properties of the visual displays? If the issue at play is truly the *lexical* accessibility of adjectives, then the displays should have controlled for the relative perceptual salience of the properties being named. As things stand, we have no way of teasing apart lexical accessibility from visual salience in Martin's results. Second, if accessibility truly determines adjective ordering preferences, why should adjective frequency apply pressure in the opposite direction? Relative frequency surely determines adjective accessibility, and we saw that adjective frequency is a significant predictor of ordering preferences. However, more frequent adjectives are preferred early in the linear structure of nominal phrases, not closer to the nominal head (at least not in pre-nominal languages like English). Accessibility as measured by adjective frequency seems to be delivering the wrong predictions.

There of course remains the possibility that accessibility is not the primary link between subjectivity and ordering preferences. In an attempt to tie properties of language structure to general principles of cognition, [Bever 1970](#) advanced the hypothesis that structural norms — among them, adjective ordering preferences — emerge from the human perceptual system. To

see how perception could determine the preferred ordering of adjectives, one must appreciate the task of the parser, at least as envisaged by Bever. Upon encountering the speech stream, the parser relies on heuristics to efficiently identify constituents and associate them with an appropriate argument structure. In service of this task, the parser needs an identification mechanism for noun phrases: where do they begin, and where do they end?

According to Bever, the more “nounlike” the adjective, the closer it appears to the modified noun. Returning to the small brown cardboard box, *cardboard* is the most felicitous when used as a noun, *brown* slightly less so, and *small* the least of all (cf. Bever’s examples (68) and (69)). Now, why should ordering adjectives according to their nounlike character ease the burden of the parser in its search for noun phrase boundaries? Bever proposes that a linear parser identifies the beginning of a noun phrase with the presence of a determiner. That same parser identifies the right edge of the noun phrase with the transition from a clearly nounlike element to an item that is “less uniquely a noun” (Bever 1970: p. 323). In other words, the primary cue to the right edge of a noun phrase is a salient decrease in nounlike character. If adjectives were randomly ordered with respect to their nounlike character, the parser might mistakenly identify noun phrase boundaries, as in [*the cardboard*] [*brown box*]. These early errors identifying phrase boundaries would cascade into a total failure for the sentence parse.

Bever’s proposal offers an intuitive explanation of the pressures that determine pre-nominal adjective ordering, and it might extend to handle the mirror-image preferences in post-nominal languages. The proposal even offers a promising connection to subjectivity: perhaps less subjective adjectives yield more-well-defined categories, which are more amenable to naming with nouns. Thus, subjectivity determines nounlike character, and nounlike character determines ordering preferences. Unfortunately, Bever’s proposal suffers a serious flaw: it lacks empirical support. In her corpus analysis of English ordering preferences, Wulff (2003) calculated nounlike character and demonstrated that it does little by way of predicting adjective order. What effect nounlike character does have on the linear order of adjectives applies pressure in the direction opposite of Bever’s hypothesis: more nounlike adjectives are marginally more likely to occur *farther* from the modified noun (Wulff 2003: p. 255).

It would appear that we have arrived at an impasse: accessibility-based accounts struggle to explain the full range of data from both pre- and post-nominal languages; they also face a serious obstacle in the form of lexical

frequencies. Bever’s perception-based account seems well-suited for pre- and post-nominal languages, but Wulff’s facts suggest the proposal is misguided. And all of these proposals lack a clear connection to adjective subjectivity. There remains another strategy, however, which shifts the explanation from how language users use adjectives to what adjectives do for language users. As we shall see, a functional account along these lines stands the best chance of explaining the role of subjectivity in determining ordering preferences.

For Seiler (1978), the clue to understanding ordering preferences lies in the task that adjectives perform: determination. Noun phrases are inherently referential, whether to real-world objects or to well-defined concepts. In either case, determiners in the broad sense — demonstratives, articles, numerals, quantifiers, adjectives, prepositional attributes, relative clauses — contribute to nominal meaning in service of pinning down a referent. With this function in mind, Seiler identifies regularities in the linear order of determiners. First, “the range of head nouns for which a determiner D is potentially applicable increases with the positional distance of that determiner from the head noun N” (Seiler 1978: p. 308). For the purposes of adjective ordering, the more nouns an adjective can felicitously describe, the farther that adjective will appear from the noun. Seiler explicitly links determiner applicability with property inherentness: less broadly-applicable, more special-purpose adjectives name properties that are more inherent to the modified noun. He gives the example of *rote hölzerne Kugeln* ‘red wooden balls’:

The semantic structure of *Kugeln* qua solid objects naturally implies material constitution of some sort; it implies — with a lesser degree of naturalness — some property in the color spectrum. To this gradient decrease in natural semantic implication corresponds the normal word order in which the ‘determiner’ with the strongly implied property is closer to the head noun than the ‘determiner’ with the less strongly implied property. (Seiler 1978: p. 309)

Properties implied by the nominal have meanings that are at least partially contained already within the nominal meaning, hence the implication (cf. the notion of mutual informativity; Futrell 2017). Determiners that are less implied by the nominal will be more informative (i.e., unexpected) when encountered, with greater informativity leading to a greater potential of pinning down the intended referent. Thus, according to Seiler, “the potential of a determiner D for singling out the object referred to by the head noun

N increases proportionally with the positional distance of D from N” (Seiler 1978: p. 309). The ordering that results presumably follows from the desire to introduce the more informative, more useful elements early in the construction of a nominal.

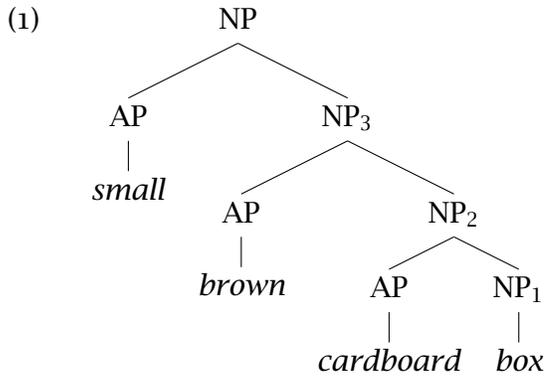
Seiler’s first claim — that adjectives describing a broader set of nouns appear farther from the modified noun — finds empirical support in Wulff’s corpus analysis (Wulff 2003: pp. 266–267). However, the implication of Seiler’s second claim — that speakers introduce more useful elements (for the purpose of determination) earlier — fails in the case of post-nominal languages with mirror-image preferences. This reasoning would hold that speakers in post-nominal languages save the most useful adjectives for last, which stands in direct conflict with the explanation for pre-nominal languages. Still, we should not abandon the functional account of adjective ordering altogether. In the following section, we consider a different proposal that preserves Seiler’s intuition that “in order to fully understand the regularities we must look behind the mere facts and try to see the program and ultimately the purposive functions of which they are manifestations” (Seiler 1978: p. 325).

4 Linking subjectivity to the hierarchical structure of modification

Let us begin as Seiler did, with the observation that adjectives aid in establishing reference. Starting with a noun like *box*, potential referents include every box in the discourse context. Where there are multiple boxes, the listener’s task of establishing reference amounts to a game of chance. We increase our odds of winning this game as we narrow down, or *determine* the set of potential referents. Encountering an adjective like *cardboard*, we now only consider the subset of boxes that are cardboard. Encountering *brown*, we limit ourselves to the cardboard boxes that are brown. Encountering *small*, we further limit ourselves to just those brown cardboard boxes that are small. From the set of all boxes we home in on the small brown cardboard boxes, a much smaller set indeed.

When it comes to the structure of these multi-adjective strings, we treat adjectival modification as syntactic adjunction, as in (1). Semantically, we treat modification as set intersection, where the adjective restricts the set characterized by the nominal denotation to just those elements that hold the specified property. For our purposes, it does not matter whether this intersection proceeds via a special mode of semantic composition (e.g., Predicate Modification; Heim & Kratzer 1998), via functional application with adjectives

of a higher type (e.g., Parsons 1970), or via functional structure (e.g., Scontras & Nicolae 2014). In each case, semantic composition proceeds outward from the noun; adjectives closer to the noun make their semantic contribution earlier than adjectives farther away. The resulting nominal denotation appears in (2), where the full NP characterizes the set of small brown cardboard boxes.



(2) $\llbracket \text{small brown cardboard box} \rrbracket =$

$$\lambda x. \text{box}(x) = \text{cardboard}(x) = \text{brown}(x) = \text{small}(x) = 1$$

To see how we arrive at the denotation in (2), consider the illustration of this process in Figure 1. Each circle corresponds to the denotation of an NP node in (1), with the elements (\square) contained within that circle representing elements of the nominal denotation. The outermost circle represents the denotation of the smallest NP, *box*. In this toy example, there are 59 boxes. Moving inward, we arrive at the next-highest NP denotation, *cardboard box*; there are only 33 such boxes, so the 26 boxes that are not cardboard are pruned from the denotation. Moving inward still, we get the 15 boxes that are both brown and cardboard; the 18 non-brown cardboard boxes have been discarded. Finally, at the innermost circle, we have just those three boxes that are at once cardboard, brown, and small; the 12 brown cardboard boxes that are not small get ignored.

Thinking about the contributions of the adjectives in the example above, we notice that, for the purpose of establishing reference, different adjectives do different amounts of work. Here, “work” gets equated with reference-establishing potential, or potential for information gain. Measured in terms of the number of possible referents considered, more work is done by the adjectives closer to the noun. In Figure 1, *cardboard* operates over the largest

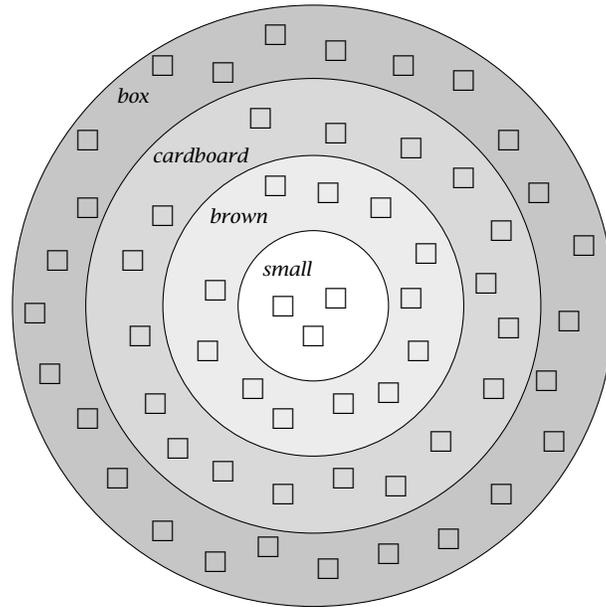


Figure 1 An illustration of restrictive modification in *small brown cardboard box*.

set: for each of the 59 boxes, one must decide whether or not it is cardboard. Thus, there are 59 opportunities to make a mistake in this decision process, wherein a listener might misjudge a box as cardboard or not. The next adjective, *brown*, operates only over the 33 cardboard boxes; thus, there are fewer possibilities for error. The last adjective, *small*, operates over the smallest set, with the smallest chance of error.

Here is the crux of the account, and finally a return to the issue of subjectivity: less subjective content is more useful for effectively communicating about the world (i.e., establishing reference). Encountering a relatively objective adjective like *cardboard*, a listener arrives at a precise concept — one that closely aligns with that of the speaker who uttered the adjective. More subjective adjectives introduce the potential for errors in alignment, as speakers and listeners might (faultlessly) disagree about category boundaries. When it comes to ordering preferences, speakers consolidate the less subjective, more useful content around the modified noun. The claim is that they do so in an attempt to aid the listener in establishing reference by minimizing errors in alignment.

The following subsection walks through concrete examples of how ordering with respect to decreasing subjectivity minimizes alignment errors

and maximizes the probability of successful referent classification. Having demonstrated the utility of subjectivity-based ordering, we then discuss potential worries and further avenues to explore.

4.1 A mathematical demonstration

To get our story off the ground, we must consider the semantics of modification in more detail. To model the potential for faultless disagreement in subjective properties, we introduce noise into the semantics of our adjectives. For each potential referent an adjective classifies, we introduce the potential for misclassification ϵ_{adj} , which stands proxy for the adjective’s subjectivity.¹ On the basis of ϵ_{adj} , each adjective has some probability $p_{adj}(\text{obj})$ of correctly classifying some object obj :

$$(3) \quad p_{adj}(\text{obj}) = 1 - \epsilon_{adj}$$

Stacking a series of noisy adjectives together in a modification structure, there will be a greater chance of incorrect classifications (i.e., misalignments) if the adjectives are not ordered according to subjectivity. As the number of objects to be classified in the nominal denotation (i.e., the cardinality of the NP denotation $|\text{NP}|$) increases, so too does the probability of misclassification $p(\text{error})$:²

$$(4) \quad \begin{aligned} p_{adj}(\text{error}) &= 1 - p_{adj}(\text{no-error}) \\ &= 1 - (1 - \epsilon)^{|\text{NP}|} \end{aligned}$$

For a multi-adjective string, we can calculate the probability of *any* errors in misclassification by multiplying the individual error probabilities from

¹ A general schema for formalizing the potential for misclassification in an adjective’s semantics appears in (i):

$$(i) \quad [\text{ADJ}] = \lambda x. \text{if ADJ}(x) \text{ then flip}(1 - \epsilon_{adj}), \text{ else flip}(\epsilon_{adj})$$

The function $\text{flip}(x)$ returns a sample from a Bernoulli distribution, where a random variable takes the value 1 with probability x ; one can think of this function as simulating the outcome of a weighted coin flip, where heads corresponds to 1 (true) and tails corresponds to 0 (false). The addition of $\text{flip}()$ to an adjective’s semantics introduces noise at the rate ϵ , where ϵ increases with subjectivity.

² For convenience, we identify NP with its extension, a set of objects.

On the grammatical source of adjective ordering preferences

each adjective, as in (5); crucially, $|\text{NP}|$ will decrease as adjectives restrict the nominal denotation.

$$(5) \quad p_{\text{NP}}(\text{error}) = 1 - (1 - \epsilon_{\text{adj}_1})^{|\text{NP}_1|} \cdot \dots \cdot (1 - \epsilon_{\text{adj}_n})^{|\text{NP}_n|}$$

Ordering with respect to subjectivity minimizes the probability of misclassifications for a multi-adjective string by ensuring that $|\text{NP}|$ decreases as ϵ increases. For a concrete example, suppose there are three boxes: box_s is small but not brown, box_b is brown but not small, and box_{sb} is both small and brown. To calculate the probability of any misclassifications for the two adjective orderings *small brown* vs. *brown small*, we will need each adjective's potential for misclassification ϵ ; we use the subjectivity scores from Scontras, Degen & Goodman 2017 to set these values: 0.20 for *brown* and 0.64 for *small*.³ As the following calculations demonstrate, ordering with respect to decreasing subjectivity, (6), results in a lower probability for misclassifications than the reverse order, (7):

$$(6) \quad p_{\text{small brown box}}(\text{error}) = 1 - (1 - 0.20)^3 \cdot (1 - 0.64)^2 = \mathbf{0.93}$$

$$(7) \quad p_{\text{brown small box}}(\text{error}) = 1 - (1 - 0.64)^3 \cdot (1 - 0.20)^2 = \mathbf{0.97}$$

However, minimizing misclassifications and correctly classifying the intended referent, thereby allowing for successful reference resolution, are subtly different notions. If we assume a fixed misclassification potential ϵ for each adjective and a truly intersective semantics for modification, the probability of correctly classifying the intended referent once a noun has been modified by multiple adjectives does not depend on the order of the adjectives (i.e., the order of semantic composition). Given the commutativity of noise in intersective modification, a nominal with two adjectives will correctly classify the intended referent *ref* with probability $p_{\text{adj}_1}(\text{ref}) \cdot p_{\text{adj}_2}(\text{ref})$. In *small brown box*, the probability that the intended referent will remain in the full nominal denotation is equal to $p_{\text{brown}}(\text{ref}) \cdot p_{\text{small}}(\text{ref})$, irrespective of order. If the pressure for subjectivity-based ordering preferences arises out of pressures toward successful reference reso-

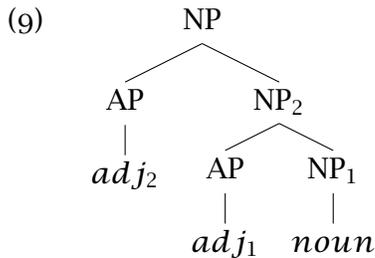
³ These values are adopted for illustrative purposes only. Subjectivity scores are likely an inflated estimate of an adjective's potential for misclassification ϵ . Still, any values for ϵ_{small} and ϵ_{brown} would do so long as $\epsilon_{\text{small}} > \epsilon_{\text{brown}}$.

lution, the commutativity of noise in intersective modification will not deliver subjectivity-based preferences.⁴

We break the commutativity of noise and thus the order-independence of modification once we recognize that ϵ_{adj} is not a fixed value, but rather varies with the size of the set to be classified, $|\text{NP}|$ (i.e., the number of objects under consideration). Each classification that must be made takes some computational processing. If we posit a fixed processing budget, more classifications will necessarily mean making each with fewer resources. Making a stochastic classification with less computation can be done at the expense of precision (i.e., with more noise): classification noise will monotonically increase with the size of the set to be classified. Thus, as $|\text{NP}|$ increases, the precision of each individual classification decreases and so the potential for misclassification grows. We model this tendency by revising $p_{adj}(\text{ref})$ so that ϵ_{adj} depends on the size of the NP denotation that adj restricts, with the constraint that $\epsilon_{adj}(|\text{NP}|) \leq \epsilon_{adj}(|\text{NP}| + 1)$:

$$(8) \quad p_{adj}(\text{ref}, \text{NP}) = 1 - \epsilon_{adj}(|\text{NP}|)$$

In a case with two adjectives, as in (9), the probability that the full multi-adjective NP correctly classifies the intended referent ref , $p_{\text{NP}}(\text{ref})$, takes into account the probability that adj_1 correctly classifies ref in NP_1 , $p_{adj_1}(\text{ref}, \text{NP}_1)$, and that adj_2 correctly classifies ref in NP_2 , $p_{adj_2}(\text{ref}, \text{NP}_2)$.⁵ However, $p_{adj_2}(\text{ref}, \text{NP}_2)$ depends on the size of NP_2 , which itself depends on potential classification errors from adj_1 . We must therefore consider all possible values for NP_2 .



⁴ Not all modification is intersective (Kamp & Partee 1995). Instead, many cases of modification are subjective, such that the interpretation of the modifier (i.e., the adjective) depends crucially on the denotation of its complement (i.e., the modified nominal). The account developed below applies to all cases of restrictive modification, including subjective ones.

⁵ For simplicity, we assume a fixed (i.e., noiseless) extension for the noun.

To calculate the probability of each possible NP_2 , $P(\text{obj} \in NP_2 | NP_1, adj_i)$ in (10) looks up the probability that adj_i (in)correctly classifies each potential element obj ; $V(\text{obj}, adj_i)$ serves as our ground truth from the speaker’s perspective, returning `true` just in case obj actually holds the property named by adj_i . The first case in (10) corresponds to the probability of correctly including an element in NP_2 that holds the property named by adj_i ; the second case corresponds to the probability of correctly excluding elements that do not hold the relevant property. The final two cases correspond to the probabilities associated with misclassifications.

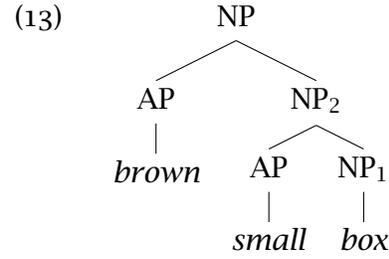
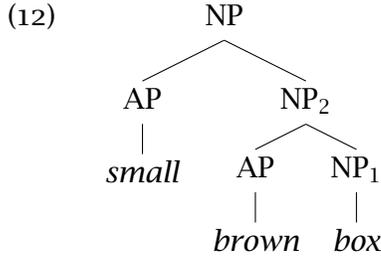
$$(10) \quad P(\text{obj} \in NP_2 | NP_1, adj_i) = \begin{cases} 1 - \epsilon_i(|NP_1|) & \text{if } \text{obj} \in NP_2 \ \& \ V(\text{obj}, adj_i) \\ 1 - \epsilon_i(|NP_1|) & \text{if } \text{obj} \notin NP_2 \ \& \ \neg V(\text{obj}, adj_i) \\ \epsilon_i(|NP_1|) & \text{if } \text{obj} \notin NP_2 \ \& \ V(\text{obj}, adj_i) \\ \epsilon_i(|NP_1|) & \text{if } \text{obj} \in NP_2 \ \& \ \neg V(\text{obj}, adj_i) \end{cases}$$

Our aim is the probability of successful referent classification at the level of NP, $p_{NP}(\text{ref})$. In other words, we want to know how probable it is that all adjectives correctly classify the intended referent. To calculate this probability, we sum over possible values of NP_2 where ref was correctly classified by adj_1 (i.e., where $\text{ref} \in NP_2$). For each potential NP_2 , we then find the probability of successful classification by adj_2 , $p_{adj_2}(\text{ref}, NP_2)$, and multiply it by the probability of having arrived at that NP_2 :

$$(11) \quad p_{NP}(\text{ref}) = \sum_{\substack{NP_2 \subseteq NP_1 \\ \text{where } \text{ref} \in NP_2}} p_{adj_2}(\text{ref}, NP_2) \cdot \prod_{x \in NP_1} P(x \in NP_2 | NP_1, adj_1)$$

Adjectives closer to the noun will compose earlier semantically, and so the number of potential referents they must classify will be larger (cf. the example in Fig. 1). In (11), this fact ensures that $|NP_2| \leq |NP_1|$. Because adjectives that compose earlier classify a larger set, we maximize $p_{NP}(\text{ref})$ by ensuring that adjectives with lower subjectivity (i.e., with a lower ϵ) compose earlier.

To see the role of subjectivity-based ordering in maximizing the probability of successful referent classification, consider the choice between *small brown box* vs. *brown small box* in (12) vs. (13).



Suppose once again that there are three boxes: box_s , box_b , and box_{sb} . To calculate the probability of successfully classifying the referent for the two adjective orderings, we will need each adjective’s potential for misclassification; here, we use the subjectivity scores from Scontras, Degen & Goodman 2017 to set the lower bound of these values (i.e., when $|NP| = 1$), and assume that ϵ increases by 0.04 with each increase in $|NP|$.⁶ We adopt the following values:

$$(14) \quad \epsilon_{small}(|NP|) = \begin{cases} 0.72 & \text{if } |NP| = 3 \\ 0.68 & \text{if } |NP| = 2 \\ 0.64 & \text{if } |NP| = 1 \end{cases}$$

$$(15) \quad \epsilon_{brown}(|NP|) = \begin{cases} 0.28 & \text{if } |NP| = 3 \\ 0.24 & \text{if } |NP| = 2 \\ 0.20 & \text{if } |NP| = 1 \end{cases}$$

In the *small brown box* vs. *brown small box* example, the first adjective to compose will operate over the set of three boxes (i.e., $|NP_1| = 3$). The second adjective will operate over a set that has been restricted by the first adjective, so we ensure that $|NP_2| \leq |NP_1|$. Starting with the preferred order in (12), there are four possibilities for NP_2 (i.e., for $\llbracket \text{brown box} \rrbracket$) that include the intended referent. In (16), we list the possible extensions of NP_2 , together with the probability of each box’s classification in parentheses (i.e., $P(\text{obj} \in NP_2 | NP_1, \text{brown})$ for each box obj); (16) also lists the probability of correctly classifying ref in NP_2 ($p_{small}(\text{ref}, NP_2)$). Multiplying across the rows, we arrive at the probability p of correctly classifying ref for each possible NP_2 ; summing over the values of p , we arrive at the probability of correctly classifying ref for the full NP: 0.23.

⁶ As before, these values are adopted for illustrative purposes only.

On the grammatical source of adjective ordering preferences

		[[brown box]]				
		box_{sb}	box_b	box_s	$p_{small}(\text{ref}, \text{NP}_2)$	p
(16)		T (0.72)	T (0.72)	T (0.28)	0.28	0.041
		T (0.72)	T (0.72)	F (0.72)	0.32	0.119
		T (0.72)	F (0.28)	T (0.28)	0.32	0.018
		T (0.72)	F (0.28)	F (0.72)	0.36	0.052
					$p_{small\ brown\ box}(\text{ref}) =$	0.23

In (17), we perform the same calculations for the dispreferred ordering in (13); with this ordering, the probability of correctly classifying *ref* for the full NP is 0.21.

		[[small box]]				
		box_{sb}	box_b	box_s	$1 - \epsilon_{brown}(\text{NP}_2)$	p
(17)		T (0.28)	T (0.72)	T (0.28)	0.72	0.041
		T (0.28)	T (0.72)	F (0.72)	0.76	0.110
		T (0.28)	F (0.28)	T (0.28)	0.76	0.017
		T (0.28)	F (0.28)	F (0.72)	0.80	0.045
					$p_{brown\ small\ box}(\text{ref}) =$	0.21

As the calculations above demonstrate, ordering with respect to decreasing subjectivity results in a higher probability of successfully classifying the intended referent than the reverse order. Because it is necessarily the case that $|\text{NP}_2| \leq |\text{NP}_1|$ in the presence of restrictive adjectival modification, this pattern holds broadly. We systematically explored the generality of this pattern with a search through the possible parameter space. By varying ϵ_{small} , ϵ_{brown} , and the size and makeup of the initial nominal denotation (i.e., the number of boxes and their properties), we tested 103,740 cases of multi-adjective modification.⁷ Of those cases tested, 93% were such that ordering with respect to decreasing subjectivity resulted in a higher probability of correctly classifying the intended referent.

⁷ Minimum values for the noise parameters varied between 0.01 and 0.77 in steps of 0.04, with the constraint that the minimum value of ϵ_{small} exceeded the minimum value of ϵ_{brown} . The initial nominal denotation varied in cardinality between two and five, with the constraint that only one box was both small and brown (i.e., the intended referent); otherwise, boxes took on all possible combinations of the properties of being small and being brown. For a hands-on look at our parameter exploration, see the online appendix with runnable code at <http://forestdb.org/models/adj-order-appendix.html>.

We thus see how subjectivity-based adjective ordering preferences could emerge once speakers take into account the perspective of their listeners. With the goal of establishing nominal reference, less subjective adjectives are less likely to lead to errors in classification, where a listener could have a diverging opinion about whether or not some objects hold the relevant property. A simple policy of misclassification avoidance delivers subjectivity-based ordering preference. Subjectivity-based ordering preferences likewise maximize the probability of successful classification of the intended referent in cases of resource-bounded computation. Although in English we encounter the reverse order, modification proceeds semantically outward from the noun. Thus, speakers employ the most useful, least subjective adjectives early in this semantic process where there is the greatest potential for misalignment. The proposed account works the same in languages with post-nominal adjectives where we find mirror-image preferences: linear distance corresponds to hierarchical distance, and adjectives that are closer to the noun make their semantic contributions earlier.

4.2 Some potential worries

We saw that the vast majority of cases explored (93%) do maximize the probability of correctly classifying the intended referent. However, there are cases where ordering with respect to decreasing subjectivity yields the opposite result. Those cases fell into one of two classes: 1) there was a larger number of objects from the less subjective category (e.g., one small brown box, two brown boxes, and one small box), ϵ_{brown} was sufficiently small, and the difference in ϵ between the two adjectives was sufficiently small; or 2) there was a larger number of objects from the more subjective category (e.g., one small brown box, one brown box, and two small boxes), ϵ_{brown} was sufficiently large, and the difference in ϵ between the two adjectives was sufficiently small. Both classes of cases rely on classification by the more subjective adjective to shrink the set of objects that will get classified by the less subjective adjective, resulting in a higher overall probability of correctly classifying the intended referent.⁸

Still, most cases of multi-adjective modification are such that ordering with respect to decreasing subjectivity maximizes the probability of correctly classifying the intended referent. From the perspective of language

⁸ For a hands-on exploration of these cases, see our appendix at <http://forestdb.org/models/adj-order-appendix.html>.

evolution, it comes as no surprise that language has regularized this strong trend as stable subjectivity-based ordering preferences (for a demonstration of how language can regularize even the slightest tendency, see Kirby 2017 and the references therein). However, if the reasoning that leads to subjectivity-based ordering preferences is active online as speakers construct noun phrases, then we might expect diverging preferences in the small set of cases that deviate from the general trend. We are unaware of any systematic exploration of grounded adjectival modification that could test these predictions, but such testing could help to resolve the issue of whether our preferences emerge as a result of language evolution, calcifying in the input; whether the pressures driving these preferences continue to be active online as we use language; or whether the preferences rely both on regularities in our input and on active pragmatic reasoning.

The astute reader will recognize that the proposed account of adjective order, which relies on incremental semantic composition, ostensibly stands at odds with the linear nature of sentence processing, specifically with respect to reference resolution. Eberhard et al. (1995) report the results of a visual-world eye-tracking study featuring multi-adjective strings; their results suggest that listeners use information from incoming words to prune the set of potential referents *as that information becomes available*. Sedivy et al. (1999) follow up on this finding by demonstrating incremental reference resolution even for context-dependent adjectives. The empirical picture appears clear: listeners' eye movements narrow in on potential nominal referents as time progresses linearly.⁹ And yet our proposed account assumes that semantic composition proceeds outward from the noun, a direction opposite to the linear uptake of words, at least in pre-nominal languages like English.

The pressures that deliver adjective ordering preferences evidence a case where hierarchical, compositional structure appears to take precedence over linear, incremental processing. The work on predictive looks during incremental processing only serves to increase the interest of this tension. However, the early uptake of semantic information evidenced by predictive looking in eye-tracking studies does not rule out that the semantic composition

⁹ While eye movements might narrow in on the potential nominal referent in visual-world eye-tracking studies, it remains unclear whether the listener's beliefs are similarly narrowed. Recent results from Qing, Lassiter & Degen 2018 suggest that eye movements in reference tasks might be only loosely correlated with the degree to which an object is believed to be the intended referent.

of nominal phrases proceeds outward from the noun; it is this semantic composition process that stands to explain the role of subjectivity in adjective ordering preferences.

A note of caution is in order: this account relies on the assumption that a speaker's goal when using modificational content is to establish reference, yet this might not always be the case. Consider a dialog in which Alex says to Chris, "Do you see what he's wearing?" Chris responds, "What a tacky polyester shirt!" Here, the referent (i.e., the relevant shirt) is already in common ground before Chris's utterance; the adjectives *tacky* and *polyester* are thus unlikely to be employed in service of establishing reference. Instead, these kinds of non-restrictive uses communicate the speaker's stance toward the shirt in question. It remains an empirical question whether speaker goals influence ordering preferences such that subjectivity plays a lesser role in the absence of reference resolution. It might also be the case that an account of non-restrictive adjective use makes the same predictions regarding the role of subjectivity (cf. [Hahn et al. 2018](#)).

Nevertheless, if we are on the right track in assuming that pressures from successful reference resolution, together with awareness of potential disagreement between speakers and listeners, lead to cross-linguistically robust adjective ordering preferences, the question turns next to how these preferences develop and how they get represented. For now we can only gesture toward possible answers. In a recent corpus analysis of child-directed and child-produced speech, [Bar-Sever, Scontras & Pearl 2018](#) documented the emergence of abstract knowledge of ordering preferences by the age of four. But are children engaging in the sophisticated theory-of-mind reasoning described above as they form these preferences? Probably not. A growing body of evidence suggests that children struggle with adult-like subjectivity awareness long after ordering preferences emerge ([Foushee & Srinivasan 2017](#)). It would seem, then, that rather than deploying subjectivity-based heuristics, children are merely tracking and reflecting the statistics of their input, a task they are known to excel at (e.g., [Saffran, Aslin & Newport 1996](#)). Children might categorize the regularities of their input according to semantic classes or adjective function, but the ultimate source of these regularities remains the interaction of property subjectivity with successful reference resolution.

5 Conclusion

Adjective subjectivity predicts adjective ordering preferences, a remarkably stable property of language design. We have offered an answer to the question of why subjectivity should play the role it does in these preferences. Subjective content allows for miscommunication to arise if speakers and listeners arrive at different judgments about a property description. Hence, less subjective content is more useful for communicating about the world. Speakers deploy this more useful content early in the semantic construction of nominals, as reflected in the hierarchical structure of modification: noun phrases are built semantically outward from the noun, and less subjective content enters earlier into this process. This reference-resolution account of subjectivity-based ordering preferences meets the desiderata explored in Section 3: it is predicated upon the findings of Scontras, Degen & Goodman 2017, and so it enjoys firm empirical support. No less important, the current proposal extends seamlessly to cover the mirror-image preferences in post-nominal languages. Perhaps most appealing is the broad applicability of the proposed account: we find the same preferences cross-linguistically because communication is a central goal of language use, so pressures for successful communication apply universally.

References

- Bar-Sever, Galia, Gregory Scontras & Lisa S. Pearl. 2018. Little lexical learners: Quantitatively assessing the development of adjective ordering preferences. In *42nd annual Boston University Conference on Language Development*, 58–71. <https://doi.org/10.1111/j.1423-0410.1990.tb02061.x>.
- Bever, Thomas G. 1970. The cognitive basis for linguistic structures. In John R. Hayes (ed.), *Cognition and language development*, 279–362. New York: Wiley & Sons, Inc. <https://doi.org/10.1093/acprof:oso/9780199677139.003.0001>.
- Cinque, Guglielmo. 1994. On the evidence for partial N-movement in the romance DP. In Richard Kayne, Guglielmo Cinque, Jan Koster, Jean-Yves Pollock, Luigi Rizzi & Raffaella Zanuttini (eds.), *Paths towards Universal Grammar: Studies in honor of Richard S. Kayne*, 85–110. Washington DC: Georgetown University Press. (Republished in Cinque 1996).

- Cinque, Guglielmo. 1996. On the evidence for partial N-movement in the Romance DP. In *Italian syntax and Universal Grammar*, 287–309. Cambridge University Press. <https://doi.org/10.1017/CBO9780511554261.011>.
- Coppock, Elizabeth. 2018. Outlook-based semantics. *Linguistics and Philosophy* 41(2). 125–164. <https://doi.org/10.1007/s10988-017-9222-y>.
- Dixon, Robert M. W. 1982. *Where have all the adjectives gone? And other essays in semantics and syntax*. Berlin: Mouton.
- Eberhard, Kathleen M., Michael J. Spivey-Knowlton, Julie C. Sedivy & Michael K. Tanenhaus. 1995. Eye movements as a window into real-time spoken language comprehension in natural contexts. *Journal of Psycholinguistic Research* 24(6). 409–436. <https://doi.org/10.1007/BF02143160>.
- Foushee, Routhie & Mahesh Srinivasan. 2017. Could both be right? Children's and adult's sensitivity to subjectivity in language. In *39th annual meeting of the Cognitive Science Society*. London: Cognitive Science Society. <https://cogsci.mindmodeling.org/2017/papers/0080/paper0080.pdf>.
- Futrell, Richard. 2017. *Memory and locality in natural language*. Cambridge, MA: Massachusetts Institute of Technology dissertation. <http://hdl.handle.net/1721.1/114075>.
- Hahn, Michael, Judith Degen, Noah D. Goodman, Dan Jurafsky & Richard Futrell. 2018. An information-theoretic explanation of adjective ordering preferences. In *40th annual conference of the Cognitive Science Society*. London: Cognitive Science Society. <https://cogsci.mindmodeling.org/2018/papers/0339/0339.pdf>.
- Heim, Irene & Angelika Kratzer. 1998. *Semantics in generative grammar*. Oxford: Blackwell.
- Hetzron, Robert. 1978. On the relative order of adjectives. In Hansjakob Seiler (ed.), *Language universals*, 165–184. Tübingen: Narr.
- Hill, Felix. 2012. Beauty before age? Applying subjectivity to automatic English adjective ordering. In *The 2012 conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL HLT): Proceedings of the student research workshop*, 11–16. <https://aclanthology.org/N12-2.pdf>.
- Kamp, Hans & Barbara Partee. 1995. Prototype theory and compositionality. *Cognition* 57. 129–191. [https://doi.org/10.1016/0010-0277\(94\)00659-9](https://doi.org/10.1016/0010-0277(94)00659-9).
- Kennedy, Christopher & Malte Willer. 2016. Subjective attitudes and counterstance contingency. In *Semantics and Linguistic Theory (SALT)* 26, 913–933. <https://doi.org/10.3765/salt.v26io.3936>.

- Kirby, Simon. 2017. Culture and biology in the origins of linguistic structure. *Psychonomic Bulletin & Review* 24. 118-137. <https://doi.org/10.3758/S13423-016-1166-7>.
- Kölbel, Max. 2002. *Truth without objectivity*. London & New York: Routledge.
- Kölbel, Max. 2004. Faultless Disagreement. *Proceedings of the Aristotelian Society* 104. 53-73. <https://doi.org/10.1111/j.0066-7373.2004.00081.x>.
- Lasersohn, Peter. 2005. Context dependence, disagreement, and predicates of personal taste. *Linguistics and Philosophy* 28. 643-686. <https://doi.org/10.1007/s10988-005-0596-x>.
- MacFarlane, John. 2014. *Assessment Sensitivity*. Oxford: Clarendon Press.
- Martin, J. E. 1969a. Semantic determinants of preferred adjective order. *Journal of Verbal Learning and Verbal Behavior* 8. 697-704. <https://doi.org/10.1037/e666672011-247>.
- Martin, J. E. 1969b. Some competence-process relationships in noun phrases with prenominal and postnominal adjectives. *Journal of Verbal Learning and Verbal Behavior* 8. 471-480. [https://doi.org/10.1016/S0022-5371\(69\)80091-5](https://doi.org/10.1016/S0022-5371(69)80091-5).
- McNally, Louise & Gemma Boleda. 2004. Relational adjectives as properties of kinds. *Empirical Issues in Formal Syntax and Semantics* 5. 179-196. <https://www.semanticsarchive.net/Archive/jU4MWYwZ/mcnally-boleda.pdf>.
- Parsons, Terence. 1970. Some problems concerning the logic of grammatical modifiers. *Synthese* 21(3-4). 320-324. <https://doi.org/10.1007/BF00484803>.
- Qing, Ciyang, Daniel Lassiter & Judith Degen. 2018. What do eye movements in the visual world reflect? A case study from adjectives. In *40th annual conference of the Cognitive Science Society*. London: Cognitive Science Society. <https://cogsci.mindmodeling.org/2018/papers/0440/0440.pdf>.
- Sæbø, Kjell Johan. 2009. Judgment ascriptions. *Linguistics and Philosophy* 32(4). 327-352. <https://doi.org/10.1007/s10988-009-9063-4>.
- Saffran, Jenny R., Richard N. Aslin & Eli Newport. 1996. Statistical learning by 8-month-old infants. *Science* 274(5294). 1926-1928. <https://doi.org/10.1126/science.274.5294.1926>.
- Scontras, Gregory, Judith Degen & Noah D. Goodman. 2017. Subjectivity predicts adjective ordering preferences. *Open Mind: Discoveries in Cognitive Science* 1(1). 53-65. https://doi.org/10.1162/OPMI_a_00005.
- Scontras, Gregory & Andreea Cristina Nicolae. 2014. Saturating syntax: Linkers and modification in Tagalog. *Lingua* 149. 17-33. <https://doi.org/10.1016/j.lingua.2014.05.005>.

- Scott, Gary-John. 2002. Stacked adjectival modification and the structure of nominal phrases. In Guglielmo Cinque (ed.), *The cartography of syntactic structures, volume 1: Functional structure in the DP and IP*, 91–120. Oxford: Oxford University Press.
- Sedivy, Julie C., Michael K. Tanenhaus, Craig G. Chambers & Gergory N. Carlson. 1999. Achieving incremental semantic interpretation through contextual representation. *Cognition* 71. 109–147. [https://doi.org/10.1016/S0010-0277\(99\)00025-6](https://doi.org/10.1016/S0010-0277(99)00025-6).
- Seiler, Hansjakob. 1978. Determination: A functional dimension for interlanguage comparison. In Hansjakob Seiler (ed.), *Language universals*, 301–328. Tübingen: Narr.
- Sweet, Henry. 1898. *A new English grammar, logical and historical*. London: Clarendon Press.
- Truswell, Robert. 2009. Attributive adjectives and nominal templates. *Linguistic Inquiry* 40. 525–533. <https://doi.org/10.1162/ling.2009.40.3.525>.
- Tucker, Gordon. 1998. *The lexicogrammar of adjectives: A systemic functional approach to lexis*. London: Cassell Academic.
- Whorf, Benjamin Lee. 1945. Grammatical categories. *Language* 21(1). 1–11. <https://doi.org/10.2307/410199>.
- Wulff, Stefanie. 2003. A multifactorial corpus analysis of adjective order in English. *International Journal of Corpus Linguistics* 8(2). 245–282. <https://doi.org/10.1075/ijcl.8.2.04wul>.

Gregory Scontras
3151 Social Sciences Plaza
Irvine, CA 92697
g.scontras@uci.edu

Judith Degen
450 Serra Mall, Bldg. 460
Stanford, CA 94305
jdegen@stanford.edu

Noah D. Goodman
450 Serra Mall, Bldg. 420
Stanford, CA 94305
ngoodman@stanford.edu