

The Deceptively Simple Problem of Cononymy

Stefan DuBois

University of California, Santa Barbara

Abstract: Within the domain of polysemy lies a distinct class of words called cononyms. These words are unique because they have two or more definitions that conflict with one another, and the intended use must typically be derived from context. To the author's knowledge, no attempt has been made to systematically investigate whether one can determine which definition is intended by any means other than contextual clues. The present investigation, however, uses a corpus linguistics approach to argue that the intended definition of one cononym, *deceptively*, can be determined based solely on the adjective it modifies. Analysis of these collocates via distinctive collexeme analysis (Gries and Stefanowitsch 2004 "Extending Collostructional Analysis") revealed that *deceptively* interacts with an adjective's polarity (*i.e.* whether the adjective denotes more or less of a given quantity), shifting the meaning of the adjective toward the positive end of its polar scale. That is, "deceptively small" and "deceptively big" both are likely to mean "bigger than it appears." These conclusions mean that the sense of *deceptively* can be predicted in ambiguous cases or in those where no contextual clues are provided.

Keywords: cononyms, antonyms, distinctive collexeme analysis, corpus linguistics, adjectival polarity.

1. Introduction

Cononymy, a special form of polysemy wherein a lexeme has two directly opposing senses, has received relatively little attention in the field of semantic linguistics. Particularly, no method has been suggested for distinguishing which of the conflicting senses is intended in cases where contextual clues might be insufficient to determine the intended meaning. For example, the phrase, "*The water is deceptively shallow*" permits two opposing interpretations:

- a) *The water may look shallow, but it is actually deep.*
- b) *The water may look deep, but it is actually shallow.*

Strictly speaking, it would be impossible to determine whether (a) or (b) is intended without the help of prosody or more contextual information. This paper, however, uses a corpus-based analysis of the most frequent collocates of *deceptively* to demonstrate that the intended sense of the construction *deceptively X* (where *X* is an adjective) can generally be predicted by assessing the adjective's polarity—that is, where it lies on a scalar gradient as related to a unit of measurement.

The present investigation begins with a summary of previous works on cononymy and adjectival polarity, as well as highlighting the principals of collostructional analysis which will be relevant in the data analysis. In Section 3, a corpus of *deceptively*'s R1 collocates (*i.e.* those words which appear immediately to the right of the adverb) is coded for each of the cononym's two opposing definitions. Section 4 applies an analytical method called *distinctive collexeme analysis* (discussed in Section 2.3) to the corpus data, and uses the results to argue for a single, unifying principle which merges the different senses of *deceptively*. Section 5 summarizes this paper's findings, calling for the investigation of other cononyms via collocate analysis, and for further exploration of the effect of adjectival polarity on the semantic sense of polysemous words in general.

2. Theoretical Background

2.1 Contronyms

Linguists differ on what semantic class contronymy¹ falls under. As summarized by Kijko (2012-2013), contronymy has variously been classified as a form of word-internal antonymy, homonymy, lexical ambiguity (“*Mehrdeutigkeit*”), or even as a type of euphemism similar to the effect sarcastic intonation has on the meaning of words (e.g. “Wonderful!”). Karaman (2008:175) provides the definition which will be used in this paper:

“Contronymy is a form of polysemy that can be defined as sense opposition at the micro-level. This occurs when a minimum of two senses of a polysemous lexical item contradict each other.”

For example, consider the following contronyms in English:

- (1) *The woman wants to **rent** her bike.*
- (2) *The man **dusted** the shelves.*
- (3) *The United Nations **sanctioned** the actions.*

In all three examples, each bold word could have two interpretations, both polar opposites of each other. In (1), the woman could be offering a bicycle to another (i.e. “lending”), or temporarily hiring it instead of buying one (“borrowing”). In (2), the man could be cleaning the shelves (“removing dust”), or dusting for fingerprints by applying a powder (“adding dust”). In (3), the attitude toward the actions could either be positive (“approved”), or negative (“condemned”).

Although relatively infrequent, the words exhibiting this bivalence are more common than one might at first expect, not only in English, but among other languages as well. One of the first identifications of the phenomenon of contronymy was in Carl Abel’s (1884) discussion of Egyptian hieroglyphics. Puzzled by the numerous conflicting senses of certain hieroglyphics, Abel posits that there must have been some way for Egyptians to distinguish between them, and expresses astonishment that communication could have taken place at all with so many opportunities for confusion: “*Da [Aegypten] Tugend und Wissenschaft so früh erwarb, kann es doch in den einfachsten Verstandesoperationen nicht unfähig bis zur Albernheit gewesen und geblieben sein*” (“Considering that Egypt developed virtue and science so early on, it can’t have been or remained so absurdly inept in the most basic of communication tasks.”) (30). He proposed as a solution to this confusion that contentious hieroglyphs were almost always accompanied by a second symbol which clarified the meaning, and that spoken speech most likely emulated this clarification

¹ The naming of the phenomenon is itself a matter of contention. Contronymy has also been called enantiosemy, enantionymy, and antilogy, while the words themselves can be variously referred to as contronyms, antagonyms, auto-antonyms, amphibolous words, enantiodromes, fence-sitters, Janus words, opposonyms, pseudo-opposites, self-antonyms, and self-contradicting words (see Kijko 2012-2013:248). Karaman (2008) chooses to call the words “contronyms” on the grounds that this is the most accurate translation of the German *Gegensinn*, the word used by Abel (1884), whom Karaman credits with first investigating the phenomenon. This paper will follow Karaman’s lead, as he calls for the establishment of a common terminology “both for consistency, and for unambiguous communication” (Karaman 2008, 174).


practice with the use of gesture. Abel presents these auxiliary tools as an intermediate step taken by an immature language towards the development of distinct lexemes, arguing that as a language advances, contronyms disappear by dividing into separate words, each of which embodies one of the two opposing senses.

Kijko (2012-2013) counters the claim that the phenomenon only occurs in ‘primitive’ languages by highlighting contronyms resulting from the speech of modern German youth, and points to authors who have compiled lists of the words for Russian, Ukrainian, and German (248). Karaman (2008) additionally cites examples from Classical Arabic, Turkish, and English, indicating that this phenomenon is likely a relatively common one among the world’s languages.

Still, contronymy has received very little investigation in terms of how one might distinguish between the two opposing senses of a word. Even Lutzeier’s (2007) dictionary of German contronyms, a multi-volume work which provides the contradictory senses of hundreds of words, offers no suggestion for how to determine which sense is intended in ambiguous cases. Other dictionaries acknowledge that distinguishing the intended meaning of a contronym is a difficult task prone to misunderstandings. Oxforddictionaries.com’s definition of *deceptively*, for example, contains the following warning: “To avoid confusion, use with caution (or not at all), unless the context makes clear in what way the thing modified is not what it first appears to be.”² The assumption, then, is that one typically distinguishes between opposing senses based on intuition and context. This paper intends to show that, in the case of the word *deceptively*, one can reliably determine the intended meaning of the speaker by considering the adjectival polarity of the word which it modifies.

2.2 Adjective Polarity

Much has been written about the scalar grading of antonym poles (see, for example, Lehrer & Lehrer 1982 for a thorough investigation), but of particular interest to this paper is the classification of the directional polarity of the pairs themselves. That is, when arranging multiple gradable antonyms in a pairwise fashion one would find the following organization fairly intuitive:



The inter-pair grouping of *small*, *cold*, and *soft* on one side and *large*, *warm*, and *hard* on the other suggests that these words have some inherent property external to a simple relativistic relationship within each antonym pair. That is to say, one could rotate each of these pairs in isolation such that the terms would be in the opposite order without causing any problems. But after grouping all three pairs together, it seems natural to make the above left- and right-groupings, rather than, for instance, grouping *large* with *cold* and *soft*.

This directionality in scalar gradients does, in fact, have a linguistic basis: adjectival polarity—that is, the classification of the members of graded antonym pairs as either *positive* or *negative*. Several authors (Seuren 1978, Kennedy 2001) have written about adjectival polarity; in a more recent work on the matter, Sassoon (2010) builds on the widespread theory that adjectives assign entities a certain numerical

² http://www.oxforddictionaries.com/us/definition/american_english/deceptively

value drawn from the realm of real numbers (Kennedy 1999, 2001). That is, *warm* might have a particular numerical value, which would be smaller than that of *hot*. Sassoon modifies previous analyses in an attempt to create a definition of polarity which is more robust and less prone to exceptions than previous theories. Her arguments are summarized below, with all explanations and examples adapted from his 2010 work.

In the case of paradigmatic antonym pairs (e.g. *tall/short* and *old/young*), adjectives are sorted according to a common unit of measure (*height* and *age*, respectively), which Sassoon calls a *base function*. The “positive” antonym in the pair denotes more of the base function, whereas the “negative” antonym denotes less: more *height* is *tall*, less *height* is *short*; more *age* is *old*, less *age* is *young*. Sassoon argues that, in relation to the base functions *shortness* and *youth*, the above relationships are inverted, leading one to wonder if polarity serves as a useful heuristic at all. Yet, as she points out, *shortness* and *youth* are themselves reversed measures of the unmarked forms *height* and *age*, indicating implicitly that the unmarked form should constitute the base function where possible. Those adjectives identified as positive along this scale exhibit systematically different linguistic behavior than their negative counterparts.

Among these systematic behaviors is the ability of positive adjectives to accept numerical degree modifiers. That is, it is permissible to say something is *30 inches tall*, whereas *30 inches short* is not. This distinction breaks down in the comparative: *30 inches taller* and *30 inches shorter* are both equally felicitous. Furthermore, ratio statements with negative adjectives are typically not permissible: as Sassoon points out, *an ostrich is twice as tall as a chicken* is acceptable, while *a chicken is twice as short as an ostrich* is not. Although the precise categorization of adjectives is an often-difficult process (Sassoon’s article itself is an attempt to construct a more watertight definition for polarity), the exact details of categorization are not necessary for the purposes of the present study and will not be explored here.

As Sassoon is at pains to point out, however, *positive* and *negative* refer to linguistic properties of adjectives, not to evaluative properties. That is, certain adjectival pairs such as *good/bad* have a definitionally-motivated evaluative aspect (see, for example, Paradis *et al.* 2012 for a discussion). Yet, *positive* and *negative* linguistic polarity have nothing to do with this evaluative scale as articulated in the following example from Sassoon (2010:144):

“[W]hether being *old* is regarded as more positive than being *young* or not is orthogonal to the distinction we are after. Linguistically, *old* is positive and *young* is negative: thus it is *old* that combines with numerical degree modifiers, such as *two years*, also in its positive (non-comparative) form.”

To date, it does not seem that adjectival polarity has received attention as a factor which influences the semantic sense of polysemous words, contronyms or otherwise. There is no indication in the literature consulted for this work that words behave differently based on the polarity of the adjective they are paired with, but given that the reverse is true (as stated above, an adjective’s polarity determines its own behavior in terms of permissible syntactic structures such as numerical modifiers and ratio statements), such a conclusion would not be completely unexpected. This paper intends to demonstrate that polarity does play a role in distinguishing between the different senses of *deceptively X*.

2.3 Usage-based Linguistics and Corpus Data

Having explored the theoretical background for the present study in the above discussions of contronymy and polarity, we will now turn to the methodology behind using corpus linguistics to explore the polysemy of *deceptively*. This paper will use a corpus-based (as opposed to corpus-illustrated) approach as defined by Tummers *et al.* (2005), wherein “the empirical evidence and the tendencies found in actual language use constitute the core of the analysis and define the resulting model, rather than just being used to support theoretical assumptions about the language system given in advance” (235).

One idea fundamental to the method of data analysis that will be used in this paper is that of the importance of collocate frequency as a viable means of exploring a word’s semantics. That is to say, the meaning of a given lexeme can be informed upon by those words with which it frequently pairs, a concept summarized by the oft-quoted adage “You shall know a word by the company it keeps” (Firth 1957:11). This theory has formed the basis for a large number of studies and numerous methodological techniques (see, for example, Church *et al.* 1991, Dunning 1993, and Pedersen 1996 for three different approaches).

The analytic method which this work will use is that of *distinctive collexeme analysis*, a specific type of *collostructional analysis*. Developed in a series of papers by Gries and Stefanowitsch (Stefanowitsch and Gries 2003, 2005; Gries and Stefanowitsch 2004 “Extending Collostructional Analysis”, “Co-varying Collexemes in the Into-Causative”), collostructional analysis examines the affinity certain words exhibit in pairing with particular constructions—*e.g.* which words prefer/disprefer a progressive construction (Stefanowitsch and Gries 2003). Rather than simply using raw frequency counts, which can be misleading because they do not take into account how common a particular word is in a corpus, collostructional analysis compares observed frequencies of a word in a construction and compares it to the frequencies one would expect given the total number of occurrences of that word in the corpus as a whole.

Distinctive collexeme analysis is a specific application of collostructional analysis used to contrast two different constructions (*e.g. will vs. be going to*) in terms of the different lexical items which pair with them. It uses the Fisher exact test, advantageous because it does not make any distributional assumptions nor does it require a particular sample size, and therefore particularly useful in the field of corpus linguistics where a given phenomenon might not be frequently attested in a data set (see Stefanowitsch and Gries 2003 for a more in-depth discussion of the advantages of the Fisher exact test). This is used to generate a value called *collostructional strength*, a numerical measure based on expected vs. observed frequencies of how attracted or repulsed a word is to a particular construction. When comparing two different constructions, a high collostructional strength indicates that a given word strongly prefers one construction while dispreferring the other. The analysis of these collostructional strengths and the semantic clusters that result from them can then point to conclusions about the behavior of the constructions in question, as will be done in Section 4 below.

3. Method

This analysis investigates the collocates of *deceptively* in a language corpus in order to search for any observable patterns in its two contradictory definitions. These two definitions are as follows: If we say *A is deceptively X*, where *A* is the subject and *X* is an adjective, the following two definitions are possible:

Definition 1: *A seems X, but is actually $\neg X$*

Definition 2: *A seems $\neg X$, but is actually X*

To repeat the example presented at the beginning of this paper, the phrase *The water is deceptively shallow*, interpreted according to Definition 1, would yield *The water seems shallow, but is actually not shallow (i.e. deep)*. According to Definition 2, it would be construed as *The water seems deep, but is actually shallow*.

The distribution of these definitions was examined by performing a search in the Corpus of Contemporary American English (COCA). Provided by Brigham Young University, the corpus is both large (“450 million words in 189,431 texts” at the time of writing) and contemporary, spanning 1990-2015. Furthermore, it contains entries from a variety of genres (spoken, fiction, newspaper, magazine, academic), all of which featured in this analysis.

A search for *deceptively* and its R1 collocates (the words located immediately to the right of *deceptively* in the corpus), returned a total of 81 types and 421 tokens. The data was manually sorted, eliminating entries which exhibited uses of *deceptively* not under scrutiny. Among these were sentences with verb collocates: “*The product was deceptively marketed*” does not call into question whether or not the marketing took place, merely the manner in which it was done, and therefore constitutes a sense which is not under investigation here. Other verb collocates included *mystifies*, *presented*, *recruited*, and *named*; it should be noted that verbs made up a negligible portion of the data set, and no verb type exhibited more than 2 tokens.

After eliminating non-relevant uses of *deceptively*, the data was then coded for whether the entry corresponded to Definition 1 or Definition 2. The COCA Corpus provides 2-3 sentences on each side of an item, which typically was enough to classify it according to contextual clues. Consider the following examples from the data:

(4)“*In the hot, bright light, dwarfed by his fields, Rick looks **deceptively small**.
Yet he's a big man, 6' 2", 215 lbs...*”

(5)“*In summary, we need to be educating social work students and practitioners
about the value of education as an intervention. It may appear to be a simple
process, but we maintain that it is **deceptively complex**.*”

Sentences like (4) and (5) all but spelled out which definition was appropriate. In (4), Rick appears small, but is actually big (Definition 1: *A seems X, but is actually $\neg X$*). In (5), the process appears simple, but is actually complex (Definition 2: *A seems $\neg X$, but is actually X*).

Not all examples in the data were so obviously identified, however. In order to aid in the coding process, a general heuristic was devised in the question “Does X conceal something beyond the fact that A is X?” If the answer was yes, then the sentence was coded for Definition 1, if not, then it was coded for Definition 2. Note that, applied to (4) and (5), this heuristic gives the same answers: does Rick’s small appearance in (4) conceal something beyond the fact that he looks small? Yes—his small appearance conceals that he is actually a big man, hence Definition 1. In (5), does the purported complexity of the process conceal something beyond the fact that the process is complex? No—the speaker is arguing that the process is complex—more complex than appearances suggest—and therefore this use matches Definition 2.

Although the definitions could be assigned without the aid of this heuristic question in (4) and (5), sentences (6) and (7) were more troublesome:

(6) *“For humans, this is one of the most dangerous sounds in the African bush: The **deceptively gentle** swish of elephants grazing. The world's largest land mammals may appear passive, but their tolerance extends only so far.”*

(7) *“Packing a sleek, **deceptively powerful** netbook sure beats lugging a laptop around.”*

For (6), one might initially be led to Definition 2: *swish* is by nature a word which implies gentleness (“rough swish” is an unnatural collocate pair), and so the second half of Definition 2, “*is actually gentle*” is an attractive choice. Yet, applying the heuristic question, one can see that more is at play: “Does the gentleness of the swishing conceal something beyond the fact that the swish is gentle?” Yes—it conceals the dangerous nature of the elephants who make the noise.

Critically, it should be noted that it is the effect of *deceptively* on a sentence level which is being examined, and not its effect exclusively on the adjective *X*. That is, in sentences such as (6), whether Definition 1 or Definition 2 applies does not change the fact that the swish is gentle, but it does change how one interprets the sentence as a whole: Definition 1 indicates that the gentle swishing belies the danger of the animals, while Definition 2 would emphasize the graceful nature of the elephants despite their dangerous appearance. These two senses differ subtly in their distinctions, but demonstrate two entirely different ways of understanding the sentence: the former emphasizes the danger posed by the elephants, while the latter emphasizes that their dangerous appearance can be deceiving. Applying the heuristic question teases out that subtle distinction in a way first impressions might not.

Continuing to (7), one can ask the same question: “Does the power of the netbook conceal anything beyond the laptop’s power?” No—the fact that the netbook is powerful when one might not expect it to be is the driving force behind the utterance.

Finally, there were some examples in the data which simply could not be determined from the context provided, even with the application of the heuristic question:

(8) *“...according to some of the greatest songwriters of all time: the **deceptively straightforward** Irving Berlin...”*

Without having external information about this particular songwriter, it is impossible to categorize the sense of *deceptively straightforward* according to either definition. Asking whether the straightforwardness of Irving Berlin “conceals something beyond the fact that he was straightforward” is similarly unproductive. Such ambiguous cases were removed from the data set. These deletions for non-

adjective collocates and ambiguous instances left a total of 68 types and 388 tokens.³ Three hundred of these tokens fit Definition 1, while only 88 fit Definition 2, indicating that Definition 1 was the more common, unmarked form.

In order to confirm that the above coding process was carried out as objectively as possible, five additional raters coded 20 randomly-selected samples from the data for the two definitions. They received two example sentences along with their ratings, and it was explained how the heuristic question discussed above could be used to differentiate between cases. After each rater independently coded the 20 samples for Definition 1 or 2, their results were compared to the coding done by the researcher. One of the raters agreed on 18 out of 20 cases, three on 19 out of 20, and one on 20 out of 20. No two raters disagreed on the same rating—that is, in the event that one rater selected a different definition than that selected by the researcher, the other four raters still agreed with the researcher’s coding. Thus, despite some concerns that the coding process would introduce a degree of subjectivity into the study, the method of arriving at those definitions as laid out above seemed sufficiently controlled and replicable so as to validate the analysis of the collected data.

Once the coding for definition was complete, a distinctive collexeme analysis was run using the R script Coll.analysis 3.2a, written by Stefan Th. Gries and available upon request (Gries 2004 “Extending Collostructional Analysis”). This method of statistical analysis was designed for two distinct constructions, yet the present study only contains one: *deceptively X*. Nevertheless, acknowledging that the construction has two contrary, polysemous senses implicitly recognizes that there are in fact two separate constructions: one which invokes Definition 1, and one which invokes Definition 2. Coding the data for these two definitions allowed the distinctive collexeme analysis to be run against the R1 collocates of the adverb. The results of the analysis are presented below.

4. Analysis

4.1 Distinctive Collexeme Analysis

Table 1 shows some example data yielded by the distinctive collexeme analysis:⁴

R1 Collocate	Obs. Freq. Def 1	Obs. Freq. Def 2	Exp. Freq. Def 1	Exp. Freq. Def 2	Pref. Cons	Coll. Str.
SIMPLE	167	24	147.6804124	43.31958763	Def 1	5.73282
CASUAL	9	0	6.958762887	2.041237113	Def 1	1.017416
STRONG	0	10	7.731958763	2.268041237	Def 2	6.623239
DEEP	0	1	0.773195876	0.226804124	Def 2	0.644349

Table 1

The first column lists the adjectives located immediately to the right of *deceptively* in the corpus. The second and third columns show how many times that adjective was observed in cases where the intended meaning was Definition 1 (*A seems X, but is actually ¬X*) and 2 (*A seems ¬X, but is actually X*), respectively. The fourth and fifth columns contain the *expected* frequencies—that is, given the number of times that definition appears in the corpus, how many times one would expect to

³ Deleting types which displayed fewer than a given number of tokens was considered, but in light of the type:token ratio being quite low in general, no types were excluded on the basis of insufficient occurrences.

⁴ Note that some formatting has been done in order to make the table more easily interpretable. Furthermore, the R script outputs delta p values, which are not included above.

see it pairing with that adjective. The differences between these observed and expected frequencies are used to calculate the final two columns: which construction (*i.e.* definition) the adjective prefers, and how strong that preference is. The larger the collostructional strength (Column 7), the more attracted the adjective is to the definition in Column 6, and the more repelled it is by the opposite definition.

Consider the data on *deceptively simple* in Table 1. First of all, one can see that *deceptively simple* occurred more times than expected in Definition 1 (167 vs. 148) and fewer times than expected in Definition 2 (24 vs. 43). That is, sentences such as

(9) “*I like the paradox of it... It's such a **deceptively simple** dance, but really so complex.*”

appeared more commonly than those like

(10) “*Our peppermint cake is delightfully fun and **deceptively simple**. Even children can take part in creating this confection.*”⁵

The construction preference column (column 6) indicates that *deceptively simple* tends to occur in instances where the speaker intends the meaning of Definition 1; that is, if something is *deceptively simple*, statistically speaking, it is more likely to appear simple and be complex than vice versa. The collostructional value indicates that this correlation is highly significant (for collostructional strength greater than 3, $p < 0.001$).

Returning to Table 1, it is evident that *deceptively casual* displays a similar tendency, with all 9 of its observed frequencies fitting Definition 1. Yet, it is here that raw frequency counts might be misleading. One might be tempted to say that since 100% of the tokens paired with Definition 1, it is reasonable to assume that *deceptively casual* will always have the same meaning, even in other data sets. Yet, it is not the pairwise comparisons between columns 2 and 3 that concern us (observed instances of Definition 1 vs Definition 2—9:0), but rather those between columns 2 and 4 (observed frequencies vs. expected frequencies—9:7). This ratio yields a collostructional strength of 1.02, which does not meet the minimum threshold for significance: only values greater than 1.3 can be counted as significant ($p < 0.05$). This means that while *deceptively casual* appears to have a tendency toward Definition 1, one cannot reliably say it will maintain this tendency outside the corpus.

Deceptively strong, on the other hand, shows a preference toward Definition 2 with a much higher collostructional strength (6.62). It was expected to pair with Definition 2 only 2 times, yet did so in all 10 observed instances. A typical example:

(11) “*With impossibly slender and **deceptively strong** arms he helped the stunned Ingrid to her feet.*”

⁵ To clarify how these sentences fit each definition, (9) signals that the dance appears simple, but is actually complex (Definition 1); conversely, (10) suggests that baking a cake should be complex, but is in fact simple enough for a child to participate.

That is, the arms appear weak, but are actually strong (Definition 2). According to the results of the analysis, *deceptively strong* is powerfully attracted to Definition 2, and repelled from Definition 1 (appears strong, but actually weak).

To summarize the above process of interpretation, the most important information gleaned from the analysis is the adjective (Column 1), its preferred definition (Column 6), and its collostructional strength (Column 7). This information is condensed into Tables 2 and 3, which contain the top 13 adjectives associated with each definition (1 and 2, respectively) along with their collostructional strengths. As mentioned above, only collostructional strengths greater than 1.30103 reach the threshold for significance; these are highlighted in dark gray below.

R1 Collocate	Collostruction Strength
Simple	5.73281954
Casual	1.017415973
Gentle	1.017415973
Easy	0.961770292
Peaceful	0.788957872
Soft	0.788957872
Mild	0.67524184
Quiet	0.527498288
Modest	0.448825568
Pleasant	0.448825568
Small	0.37118769
Delicate	0.336121288
Elegant	0.336121288

Table 2 (Definition 1 Preferred)

R1 Collocate	Collostruction Strength
Strong	6.623239239
Complex	5.94222712
Difficult	4.593093553
Quick	3.924844043
Powerful	1.944629619
Bright	1.292540766
Brilliant	1.292540766
Fast	1.292540766
High	1.292540766
Spacious	1.292540766
Tough	1.292540766
Wearable	1.292540766
Large	1.1008934

Table 3 (Definition 2 Preferred)

Beginning with Table 2, only one adjective displayed a statistically significant attraction to Definition 1 (collostruction strength > 1.3). As seen above in Table 1, however, *deceptively simple* appeared a remarkable number of times in the corpus: 191 tokens out of a total 388 (49%). This was far and away the most common collocate—the next highest was *deceptively easy*, which displayed only 15 tokens—and right away suggests that *deceptively simple* is the prototypical, most commonly used form.

Furthermore, given that 167 of these 191 tokens fit Definition 1, this set a very high standard for the number of observed frequencies required for a significant collostructional strength. Given that the total occurrences of other adjectives was considerably smaller (most other types had between 1-5 tokens), it is unsurprising that their collostructional strengths were comparatively low, as there were simply not enough tokens of each type to reach the minimum number required for a significant collostructional strength; the high number of *deceptively simple* tokens had in effect set an unattainable standard. Both *deceptively casual* and *deceptively gentle*, for example, appeared nine times each, with all nine definitions falling under Definition 1, yet were still far from reaching the significance threshold.

In order to counteract this unbalancing effect, the collostructional analysis was repeated with the *deceptively simple* tokens removed from the data set. This second iteration yielded three more adjectives significantly attracted to Definition 1—these are highlighted in light gray above. There was no change in the adjectives attracted to Definition 2, although the collostructional strengths of each were reduced, as one should expect. The output of this second collostructional analysis, along with the new collostructional strengths, is available in Appendix A.

Strictly speaking, these additional adjectives should not be considered in the analysis, as they only reached the significance threshold in an artificially-restricted data set. Yet, given the unbalancing effect of *deceptively simple*, it seems fair to incorporate them in the discussion, and they will be included in references to the significant adjectives of Table 2, although the reader is reminded that their inclusion might be disputable on statistical grounds.

Unlike the high number of tokens necessary to reach significance thresholds in terms of Definition 1, those adjectives that paired with Definition 2 needed comparatively fewer tokens to warrant a high collostructional strength. Five reached the significance threshold, whereas another 7 were just below the 1.30103 cutoff. Note that the second iteration of the collostructional analysis lowered the collostructional strength of all the adjectives in Table 3, although it did not push any of them below the significance threshold. Again, the relevant values are available in Appendix A.

The collostructional strengths of Tables 2 and 3 made for a rather striking set of semantic clustering. Even at first glance, it is immediately apparent that a strong intra-group similarity exists between the two groups of words. Compare the top five adjectives in each table: *simple, casual, gentle, easy, peaceful* (Definition 1), with *strong, complex, difficult, quick, powerful* (Definition 2). Speaking on a purely conceptual level, the first set of words are subdued, passive, low-energy words while the opposite seems true of the second set. More precisely, where a measure term is available, we can say that the words in the first set denote “less” of a given quantity (*complexity, formality, difficulty*) whereas those of the second set denote “more” (*strength, complexity, difficulty, speed, power*). In fact, when comparing these measure terms, it becomes apparent that there are two sets of paradigmatic antonym pairs divided between Table 2 and Table 3: *simple/complex* and *easy/difficult* (*small/large* is present as well, although neither adjective displayed a statistically significant attraction).

In order to hone in on a more precise linguistic terminology, the reader is invited to recall the discussion of adjectival polarity in Section 2. When considered in terms of polarity, a stark difference between the two tables becomes apparent: Table 2 consists primarily of linguistically negative adjectives, while Table 3 is comprised of principally positive terms. Table 4 presents the nine most strongly-attracted adjectives organized with respect to their polarity:

Negative	Base Function	Positive
Simple	Complexity	Complex
Casual	Formality	(Formal)
Gentle	Roughness	(Rough)
Easy	Difficulty	Difficult
(Weak)	Strength	Strong
(Slow)	Speed	Quick
(Weak)	Power	Powerful

Table 4

The words highlighted in blue represent the adjectives taken from Table 2, while the ones in green represent those from Table 3. The words in parentheses represent antonyms which were not attested in the data set, but are included along with the base function in order to make the adjectives’ polarity more transparent.

Adjectives were situated in the table by assigning their polarity according to a common, unmarked base function as outlined in Section 2.2.

Note that while many of the above pairs are fairly intuitive, Sassoon (2010:142) comments that it is “notoriously difficult to define the set of negative as opposed to positive adjectives” (cf. Lehrer 1985:419; Sassoon 2012:33), and judgments about polarity are at times subjective. In fact, according to Sassoon herself (2010:167), *easy/difficult* are reversed from the polarity listed in Table 4—that is, Sassoon places *easy* as the positive antonym of the pair, and we should therefore presumably expect it to be attracted to Definition 2 rather than Definition 1. It is certainly possible that she is correct and the antonym pair simply does not follow the established pattern; after all, she does include *easy/difficult* in a division of antonym pairs which do not exhibit prototypical behavior. One explanation for its attraction to the opposite definition from the expected one therefore is that this is simply an extension of the pair’s erratic conduct.

An alternative method of handling this discrepancy is to recognize that *deceptively* is interacting with something quite similar to, but slightly different from polarity as Sassoon defines it. Whereas Sassoon’s criteria for ambiguous cases of polarity can result in a marked base function (*ease* as opposed to *difficulty* in the case of the pair *easy/difficult*), the present study will adhere exclusively to the unmarked base function. In terms of *easy/difficult*, the two possible base functions are *ease* and *difficulty*, and the latter can easily be proven to be the unmarked form of the two.⁶ Therefore, the polarity of *easy/difficult* in the above table seems reasonable.

Accepting the above judgment of polarity, we are left with the following conclusion: negative adjectives tend to pair with Definition 1, and positive adjectives with Definition 2. Even though this conclusion has been drawn from a very small group of adjectives, an examination of the non-significant adjectives in Tables 2 and 3 lends it further support: *soft*, *quiet*, and *small* (Table 2—Definition 1) are prototypically negative adjectives, while *fast*, *high*, and *large* (Table 3—Definition 2) are prototypically positive. Therefore, the results of the distinctive collexeme analysis suggests a strong correlation between polarity and definition. The implications of this connection are discussed in the following section.

4.2 Discussion

Until now, Definition 1 has been summarized as *A seems X, but is actually* $\neg X$, while Definition 2 has taken the form *A seems* $\neg X$, *but is actually* *X*. Knowing what we do about the distribution of adjectives between the two definitions, let us restate them with reference to adjectival polarity. If we assume *X* is likely to be a negative adjective in Definition 1 as suggested by the collostructional analysis in Section 4.1, let *X* be “*negative*” and $\neg X$ be “*not negative*,” i.e. “*positive*.” This yields the following definition:

Definition 1.2: *A seems negative, but is actually positive*

Similarly, for Definition 2 (*A seems* $\neg X$, *but is actually* *X*), which tends to pair with positive adjectives, let *X* be “*positive*” and $\neg X$ be “*negative*”:

Definition 2.2: *A seems negative, but is actually positive*

⁶ See Lehrer 1985. *Difficult* is neutralized in questions and nominalizations: the sentences “*How difficult is the problem?*” and “*I was surprised by the difficulty of the problem*” make no presuppositions about the problem’s difficulty, whereas “*How easy is the problem?*” and “*I was surprised by the ease of the problem*” do. Additionally, as per Sassoon (2012:13) *easy* decomposes into *little difficult*.

The identical nature of Definitions 1.1 and 1.2 gives us a unifying principle which allows the construction *deceptively X* to be interpreted regardless of which adjective is present. Articulating the two definitions in this form demonstrates that, no matter the polarity of *X*, the adverb *deceptively* attempts to shift its meaning toward the positive end of the scale. Those adjectives which are already positive retain their meaning, while those that are negative change their sense to positive. That is to say,

(12)*The problem is deceptively simple*

and

(13)*The problem is deceptively complex*

are expected to both have the same meaning: the problem looks simple, but is actually complex.

While it is not impossible for (12) and (13) to have different meanings, in the absence of further context, the conclusions of this paper suggest that one could make a reasonable guess about the intended sense of the construction *deceptively X* without much difficulty. According to what we have seen here, no matter what adjective *X* is, looking at where it falls on a gradient scale and moving it toward the positive end will likely reveal the intended meaning.

5. Conclusion

In this paper, an attempt was made to resolve the ambiguity of the contronym *deceptively* when paired with an adjective. A corpus analysis of collocates suggested that context alone is not the only recourse a listener has when deciding which of two definitions is meant. According to an investigation of 388 items, adjectives that pair with *deceptively* tend to lie on one end of a gradable scale. If a listener can successfully ascertain whether that adjective lies on the negative end or the positive end of that scale, they have a strong chance of predicting the intended meaning of the phrase without resorting to context: negative adjectives prefer Definition 1 (*A seems X, but is actually $\neg X$*), whereas positive adjectives words prefer Definition 2 (*A seems $\neg X$, but is actually X*). If we assume adjectives only pair with their expected definition, these two definitions can be summed up under the general principle: *A seems negative, but is actually positive*. That is, *A is deceptively X* pushes the “actual” meaning of *X* towards the positive end of the scale.

Several directions for further study have presented themselves. Most critical to the validity of this study’s results is the elimination of the subjective factors in data coding: matching each data point with Definition 1 or 2 involved at least some contentious instances and was entirely done by a single researcher. While efforts were made to make this coding process as objective as possible, a more rigorous control of the study’s subjectivity is necessary. This control could be achieved by asking a group of subjects to perform the data annotation independently of one another, and only accepting those data points which were coded unanimously. Furthermore, a separate group of subjects could code for adjectival polarity—after coding for both definition and polarity, a statistical test for independence (for example, chi square) could determine to what degree definition and polarity are dependent on one another. As a final note, repeating the very same analysis performed above with a different, larger corpus could potentially yield different results, as the number of available tokens for each type was in fact quite small.

Although the results of this study suggest the discovery of an explanatory factor in the contronymy of the structure under investigation, Tummers *et al.* (2005:243) points out that the confirmation or denial of such an explanatory factor should always fit into a larger framework. Language phenomena are typically far more complex than can be explained by a single factor alone. Such a framework can be found in the recent work of Gries *et al.* in behavior profiling (see Gries & Otani 2010 for an overview), which provides a thorough method of investigating polysemy and antonymy via the interaction of many fine-grained linguistic factors—sentence transitivity, clause level, countability, *etc.*—called ID tags.

While a complete behavioral profile of the construction *deceptively X* is outside the scope of this paper, it can be considered an intermediate step on the road to that goal. The results of this study demonstrate that adjectival polarity is a significant factor in the behavior of the construction, and therefore polarity should be coded as an ID tag if a behavioral profile is constructed. Furthermore, given the evidence that polarity has an effect on the semantic sense of the construction in question, a strong argument can be made in favor of including polarity in behavioral profiles of other polysemous words as well.

Finally, this paper’s conclusions can more generally-speaking be used as further support for the utility of frequency statistics and usage-based linguistics. Collostructional analysis in particular has proven to be an effective tool in shedding light on an otherwise puzzling case of polysemy, and no doubt will be useful in continuing to explore this and other semantic phenomena which until now have received relatively little attention.

Appendix A—Results of Collostructional Analysis Without *Deceptively Simple*

R1 Collocate	Collostruction Strength
Easy	1.720420123
Casual	1.57514177
Gentle	1.57514177
Peaceful	1.217206923
Soft	1.217206923
Mild	1.039977276

Definition 1

R1 Collocate	Collostruction Strength
Strong	5.103724922
Complex	4.569929763
Difficult	3.51877724
Quick	3.001171867
Powerful	1.478844696
Bright	0.981201774

Definition 2

Works Cited

Abel, Carl. *Über den Gegensinn der Urworte*. Leipzig: Verlag von Wilhelm Friedrich, 1884.

Church, Kenneth, et al. “Using Statistics in Lexical Analysis.” *Lexical Acquisition:*

Exploiting On-line Resources to Build a Lexicon 115 (1991): 164.

- Davies, Mark. *Corpus of Contemporary American English (COCA)*,
 corpus.byu.edu/coca/. Accessed 4 March 2015.
- Dunning, Ted. "Accurate Methods for the Statistics of Surprise and
 Coincidence." *Computational Linguistics* 19.1 (1993): 61-74.
- Firth, J. R. "A Synopsis of Linguistic Theory 1930-1955 in Studies in Linguistic
 Analysis, Philological Society." (1957).
- Gries, Stefan Th, and Anatol Stefanowitsch. "Extending Collostructional Analysis: A
 Corpus-based Perspective on Alternations." *International Journal of
 CorpusLinguistics* [e1]9.1 (2004): 97-129.
- . "Co-varying Collexemes in the Into-Causative." *Language, Culture, and Mind*
 (2004): 225-36.
- Gries, Stefan Th, and Naoki Otani. "Behavioral Profiles: A Corpus-based Perspective
 on Synonymy and Antonymy." *ICAME Journal*34 (2010): 121-50.
- Gries, St. Th. "Coll. analysis 3.5. A script for R to compute perform collostructional
 analyses." (2014): n.p.
- Karaman, Burcu I. "On contronymy." *International Journal of Lexicography* 21.2
 (2008): 173-192.
- Kennedy, Christopher. "Gradable Adjectives Denote Measure Functions, not Partial
 Functions." (1999).
- Kennedy, Christopher. "Polar Opposition and the Ontology of Degrees." *Linguistics
 and Philosophy* 24.1 (2001): 33-70.
- Kijko, Svitlana V. "Deutsch-ukrainische zwischensprachliche Enantiosemie." *Das
 Wort. Germanistisches Jahrbuch Russland* (2012-2013): 247-262.
- Lehrer, Adrienne, and Keith Lehrer. "Antonymy." *Linguistics and Philosophy* 5.4
 (1982): 483-501.

- Lehrer, Adrienne. "Markedness and Antonymy." *Journal of Linguistics* 21.2 (1985): 397-429.
- Lutzeier, Peter Rolf. *Wörterbuch des Gegensinns im Deutschen*. de Gruyter, 2007.
- Paradis, Carita, et al. "Evaluative Polarity of Antonyms." *Lingue e Linguaggio* 11.2 (2012): 199-214.
- Pedersen, T. Fishing for Exactness. Proceedings of the SCSUG 96 in Austin, TX (1996), 188–200.
- Sassoon, Galit Weidman. "The Degree Functions of Negative Adjectives." *Natural Language Semantics* 18.2 (2010): 141-181.
- Sassoon, Galit W. "A Typology of Multidimensional Adjectives." *Journal of Semantics* 30.3 (2012): 335-380.
- Seuren, Pieter AM. "The Structure and Selection of Positive and Negative Gradable Adjectives." *Papers from the Parasession on the Lexicon, CLS 14* (1978): 336-346.
- Stefanowitsch, Anatol, and Stefan Th Gries. "Collostructions: Investigating the Interaction of Words and Constructions." *International Journal of Corpus Linguistics* 8.2 (2003): 209-243.
- . "Covarying collexemes." *Corpus Linguistics and Linguistic Theory* 1.1 (2005): 1-43.
- Tummers, José, Kris Heylen, and Dirk Geeraerts. "Usage-based Approaches in Cognitive Linguistics: A Technical State of the Art." *Corpus Linguistics and Linguistic Theory* 1.2 (2005): 225-261.

Recebido para publicação em 07-11-17; aceito em 13-12-17