15

Animacy and Syntactic Structure: Fronted NPs in English

Neal Snider and Annie Zaenen

15.1 Animacy in Natural Language

It has long been known that whether the referent of a nominal is animate or not can be important in determining its syntactic or morphological realization. To describe this effect, researchers have proposed a number of hierarchies. The original hierarchy due to Silverstein (1976) conflates definiteness distinctions, animacy distinctions and person distinctions into one ordering called the 'animacy hierarchy'. We follow Aissen (2003) (based on Croft 1988) in distinguishing separate hierarchies because they refer to different aspects of entity representation within language: the *definiteness dimension* is linked to the status of the entity as already known or not yet known at a particular point in the discourse, the *person hierarchy* depends on the participants within the discourse, and the *animacy status* is an inherent characteristic of the entities referred to. We moreover assume that the traditional definiteness hierarchy, which looks at the morphological marking of the nominal, is in fact a proxy for an ordering according to information status (see below). Each of these three aspects, however, contributes to making entities more or less *salient* or *accessible* at a particular point in the discourse.

As long as one's attention is limited to the distinction between grammatical and ungrammatical sentences, the importance of the animacy hierarchy is mainly relevant for languages with a richer morphology than English. In such languages animacy distinctions can influence

Intelligent Linguistic Architectures: Variations on themes by Ronald M. Kaplan.
Miriam Butt, Mary Dalrymple, and Tracy Holloway King (eds.).

grammaticality of, e.g., case-marking and voice selection. To give just two examples, in Navaho, a *bi*-form is used when the patient is animate and the agent is inanimate, whereas the *yi*-form is used when the agent is animate and the patient is inanimate, as illustrated in (1) (from Comrie 1989:193).

(1) a. At'ééd nímasi bi-díílíd
girl potato burnt
'The potato burnt the girl.'

b. At'ééd nímasi yi-díílíd
girl potato burnt
'The girl burnt the potato.'

In Spanish, animate direct objects are introduced by *a*, whereas inanimates are bare NPs, as illustrated in (2):

(2) a. Vi *el* libro.
'I saw the book.'

b. Vi *al* niño.
'I saw the child.'

As discussed, inter alia, in Aissen (2003), this animacy distinction interacts with definiteness. It is conceptually desirable to distinguish between animacy and definiteness or information status but in practice it is frequently the case that a linguistic phenomenon is conditioned by multiple conceptually independent factors (see e.g. Comrie 1989 for discussion). This needs to be reflected in the way the data is analyzed and modeled. We discuss this later in more detail.

Recent linguistic studies have highlighted the importance of animacy distinctions in languages such as English. For instance, the choice between the Saxon genitive and the *of*-genitive (Leech et al. 1994, Rosenbach 2002, 2003, O'Connor et al. 2004), between the double NP and the prepositional dative (Bresnan et al. 2005), between active and passive (Bock et al. 1992, McDonald et al. 1993) and between pronominal and full noun reference (Dahl and Fraurud 1996, based on Swedish data) have all been shown to be conditioned by the animacy or inanimacy of the referents of the arguments whose realization can vary. In these cases, the difference between animate and inanimate does not lead to a difference between a grammatical or an ungrammatical sentence, as in the cases exemplified above, but to a difference in acceptability.

As some of the references given above indicate, psycholinguists too have investigated the importance of animacy in language. For them it is one of the many factors that play a role in sentence production. The mainstream hypotheses rely on a notion of salience or accessibility that

is an amalgam of different functional factors. Prat-Sala (1998), for instance, lists the following from the literature: predictability, semantic priming, animacy, concreteness, prototypicality, as well as some that are linked to the shape of the words: word length, metrical structure, phonological priming and word frequency. In her own study she adds discourse factors. The lists are not exhaustive but it is clear that it is difficult to study all the listed factors at once. Ideally, though, this should be done as these factors combine to make one entity more salient than another. In experimental studies one typically tries to keep all the factors except one constant. In corpus studies one has to try to tease out the contribution of the various factors statistically. In the study reported on here, we examined the influence and the interactions of weight (represented by the number of words in a constituent), information status and animacy. In this paper we focus on what the results tell us about the role of animacy. Information status is discussed in more detail in Snider (2006).

There are currently two main hypotheses about how animacy influences syntactic realization: the first assumes that the grammatical function realization of a semantic argument follows the 'Syntactic Accessibility Hierarchy' or Grammatical Function hierarchy postulating the following ordering: SUBJ > OBJ > OBJ2 > OBL (henceforth GFH; see, for instance, Bock and Warren 1985). The second hypothesizes that the surface linearization of arguments is directly conditioned by accessibility (henceforth WOH, see for instance Kempen and Harbush 2004).[1]

For English it is difficult to distinguish between these two hypotheses and several sets of data are compatible with both, e.g. the studies done on the choice between passive and active in Prat-Sala (1998), Bock et al. (1992), and McDonald et al. (1993).

15.2 Animacy and Fronting Constructions

One way of distinguishing between the two hypotheses is to look at other languages where word order precedence does not correlate with the Grammatical Function hierarchy. Another way is to look at cases in English where elements come earlier in a sentence without being higher on the Grammatical Function hierarchy. In this paper we pursue the latter and present a comparison of the occurrence of animates and inanimates in Left Dislocation and Topicalization with the occurrence

[1] Hierarchies are generally assumed to be totally ordered, but in fact most studies only look at two adjacent elements, e.g. subjects and objects, or objects and second objects.

of animates and inanimates as in situ arguments.

The two constructions studied are illustrated in (3):

(3) a. Topicalization: 'Brains you're born with. A great body you have to work at.' [Brooke Shields, in health club commercial]

b. Left Dislocation: 'That guy, I met him last week in the grocery store.'

Note that we use *Topicalization* here to refer to a syntactic construction without implying a link to pragmatic topichood. In fact, this syntactic configuration most likely has several different pragmatic uses. We will refer to the sentence initial constituents in these constructions as the *fronted elements* and to the subcategorization relation between the fronted element and the verb as the in-situ grammatical function. For Left Dislocation, the in-situ realization of the fronted element is a pronoun, while for Topicalization there is no overt realization.[2] We did not study other fronting constructions that might occur in the corpus. In examining the role of animacy in these two constructions, we need to extend the hypotheses given above to these cases. As stated, they do not take long distance dependencies into account. Especially for Topicalization, we need to examine the importance of the in-situ grammatical role that the fronted element is linked to. In written text, it is impossible to see whether the subject of the highest clause is 'topicalized' or in-situ. Only the Topicalization of a subject in an embedded clause leads to a marked word order and such cases are very rare (in our corpus, Topicalization from non-subject position is 330 times more frequent than Topicalization from subject position). Subjects tend to be animate, and topicalized NPs overwhelmingly have a gap in a non-subject position. Taking this into consideration, the extension of the predictions of WOH is rather straightforward: given that the hypothesis takes surface word order as the factor that animacy influences, it predicts that, when other factors are controlled for, fronted elements will tend to be animates, regardless of their link to in-situ grammatical functions. Under the assumptions of GFH, it is not so clear what to expect. One could assume topicalized NPs to be animate or non-animate to the same degree as a referent realized in the in-situ position would be. For Left Dislocation, one could reason that the anaphoric binding to a pronoun in-situ might also lead to the same animacy preferences as an in-situ constituent, although this reasoning seems to be weaker in this case than in the case of Topicalization because the identification

[2] We realize that it is not so clear what should go under the term Left Dislocation, but we have limited ourselves to cases where the fronted element can be linked to an in-situ pronoun.

between the fronted constituent and the pronoun can be considered to be weaker (in LFG terms, anaphoric binding versus functional control). Alternatively, one could place the fronted functions on the hierarchy of grammatical functions. From a purely formal point of view, one can propose that they are lower or that they are higher than the subcategorized functions.[3]

To analyze the data we use logistic regression because it gives the researcher great power to determine the factors that influence construction choice, even in rare constructions such as Left Dislocation and Topicalization. In particular, logistic regression provides a means to overcome problems that are inherent in the use of corpus data: correlated variables and multiple speakers. Before going into more details about methods and results we briefly describe the corpus data we used.

15.2.1 The Data

The corpus

We use a part of the Switchboard corpus, a corpus of spoken English that was compiled from telephone dialogues involving speakers from different parts of the United States (Godfrey et al. 1992). This corpus has been annotated over the years to make it more useful for syntactic studies. The Treebank Project (Marcus et al. 1993) of the Linguistic Data Consortium released a version of Switchboard annotated for part of speech and hierarchical syntactic structure. We relied on this annotation to extract the Left Dislocations and Topicalizations from the corpus.

The Switchboard corpus has also been annotated for various semantic and pragmatic features. The Edinburgh-Stanford LINK project on Paraphrase annotated the nominals in subsections of the corpus for animacy and for information status (see Nissim et al. 2004 and Zaenen et al. 2004 for a detailed description of these efforts). The animacy annotation was done for the whole syntactically annotated subpart of the corpus by Stanford, while Edinburgh did the information status annotation for a smaller subpart.

We will briefly describe how the nominals in the corpus were annotated for the various factors and how this information was used in the current study. For a more extensive discussion see Snider and Zaenen (2006).

[3] The psycholinguistic interpretations of these two choices are not equally natural, as we will discuss in the conclusion.

Animacy categories

One major problem with devising an animacy hierarchy is that the linguistically relevant notion of animacy does not directly correspond to biologically-based distinctions. Another is that it is not clear how many distinctions are linguistically relevant nor whether the same distinctions play a role in all languages. Binary animacy distinctions such as human/non-human and animate/inanimate have been proposed as well as more fine-grained ones.

The LINK project opted for a nine-valued scale based on Garretson et al. (2004), distinguishing *humans* (HUM), *organizations* (ORG), *animals* (ANIMAL), *intelligent machines* (MAC), *vehicles* (VEH), other *concrete entities* (CONC), *places* (PLACE), *times* (TIME), and other *non-concrete entities* (NONCONC).[4]

The nominals in the syntactically annotated part of the Switchboard were coded for these categories by three annotators. Evaluation (Zaenen et al. 2004) showed that the interannotator agreement in general was very high. But it also showed that some of the distinctions were not reliably annotated. Moreover, the amount of data that we have available in the corpus for Topicalization and Left Dislocation does not allow us to make a nine-way distinction. There would have been too many variables for the number of facts. Therefore, we collapsed the nine-valued scale into a binary distinction between *animates* and *inanimates*, the animates comprising HUM, ORG, ANIMAL, MAC, and VEH, and the inanimates of CONC, PLACE, TIME, and NONCONC.

The details of the information status annotations are not the focus of this paper. They are described in Nissim et al. (2004) and discussed in Snider (2006). Suffice it to say that a three-way distinction between *old*, *mediated*, and *new*, based on Prince (1981), was used.

15.2.2 Characteristics of the data and analysis techniques

The Switchboard corpus records the speech of a great variety of speakers, but we want to come as close as possible to a general picture of the speech community as a whole. Therefore we need to control for speaker variation. It is also well known that animacy effects are correlated with other factors such as information status (definiteness), as pointed out in the introduction. To overcome these problems we used logistic regression techniques.

[4]The somewhat esoteric categories 'vehicle' and 'intelligent machine' were introduced because it has been claimed that humans treat moving objects like cars and computers as human. There were not enough examples of these categories in the corpus to test this.

Correlated factors

Correlations pervade naturalistic linguistic data. This has often led to reductive theories that attempt to explain the data in terms of one factor. To give just one example, Hawkins (1994) proposes a theory of linear order in sentences. He postulates that shorter expressions occur earlier in sentences to facilitate processing and assumes that because discourse givenness is correlated with shorter, less complex constituents, apparent effects of givenness reduce to the preference to postpone syntactically complex (longer) phrases later. Arnold et al. (2000) and Bresnan et al. (2005) show how regression techniques can be used as a test of such theories. Logistic regression allows one to control for many factors, even correlated ones, simultaneously, so their independent effects can be measured. Bresnan et al. (2005) found that givenness, animacy, and length factors all have independent effects in predicting the dative alternation.

As we said above, the factors interacting with animacy that we are considering in this study are grammatical function, information status and weight. Logistic regression allowed us to determine the independent effects of these factors, and as we will see in more detail in the next section, all the factors do indeed have an independent effect.

Multiple Speakers

Another possible problem with corpora is that the data is pooled from many different speakers. Corpora such as the Switchboard corpus used here explicitly include data from many different speech communities. Newmeyer (2003) claims "There is no way that one can draw conclusions about the grammar of an individual from usage facts about communities, particularly communities from which the individual receives no speech input." Bresnan et al. (2005) point out that this is an empirical question, one that can be answered using modern statistical techniques. They use *bootstrap sampling* to show that data from different speakers does not affect their logistic regression model, which supports the idea that their conclusions about the dative alternation represent generalizations about many English speech communities, whose differences are not significant relative to other substantive factors. Another way to control for different speakers is the one employed here, the mixed model. In a mixed model logistic regression, the speaker is modeled as a random factor, in that each speaker is allowed to have a different base rate of producing the construction in question, assuming only that the inter-speaker variation in rate is normally distributed. Once the presence of different speakers is controlled in this way, the model allows us to draw conclusions about the factors that independently influence

their construction choice.

15.2.3 Interpreting logistic regression models

Regression models are well suited to corpus analysis because they allow one to determine the effects of factors, while controlling for others. They do this by fitting a mathematical model which contains coefficients for all the factors in the data relevant to the researcher's hypotheses. The models used in corpus analyses such as this one are a special type of regression, called logistic regression, which is suited for modeling categorical dependent variables (such as construction choice). These models predict the *odds ratio* of occurrence. Odds ratios should be familiar from their use in horse racing: a bookmaker might put the odds of a horse winning as $1:10$ or 0.1, that is, there is a $\frac{1}{11} = 9.1\%$ chance the horse will win. In a linguistic example, the odds of a particular construction occurring, as opposed to another, might be 50%, with an odds ratio of 1. When a logistic regression models the effects of various factors on the odds ratio, the coefficients associated with the factors are interpreted in terms of how much they increase (or decrease) the odds ratio of the construction's occurrence. If the factor is itself categorical, say animate vs. inanimate, then the regression coefficient is interpreted as how much more likely one value makes the construction to occur over another value of the factor. For example, if the value for the animacy coefficient is 2, then animates make the construction twice as likely to occur as inanimates. If the factor is continuous, like a length in words, then the coefficient represents the increase in odds of the construction for each increment of the factor. For example, if the length coefficient has an odds ratio of 0.5, then the construction is 50% more likely for each one word increase in length. Finally, when using logistic regressions to draw inductive inferences about the general behavior of a construction, and not merely the specific structure of the corpus, one needs to be careful not to build models that "over-fit" the data. This can happen when the model has more degrees of freedom than the data allow. This caveat is relevant to this work in that the topicalization data set is so small, it only allows three degrees of freedom (three independent variables) per model. In the topicalization data below, we report results for more than three independent variables, so all the results were verified by constructing sub-models that were limited to three degrees of freedom. Most importantly, the animacy results are the same in a model that only contains animacy, information status, and grammatical function as predictors.

TABLE 1

Factors	Coefficient	F-value	p-value
(Intercept)	-2.032766	57.61879	$< .0001$
animacy		9.36251	0.0022
inanimate vs animate	1.477782		
status		4.48857	0.0113
old vs mediated	-1.908018		
new vs mediated	-0.199473		
gf		89.80565	$< .0001$
subj vs non-subj	-5.846039		
weight	0.219027	6.23979	0.0125

TABLE 2

Topicalization	Topics	Subjects	Objects
animates	10	8929	2237
inanimates	87	2665	7135

15.3 Results

15.3.1 Topicalization

As predicting factors, we used a binary animate/inanimate distinction, three values for information status, as well as weight and in-situ grammatical function. For Topicalization, the coefficients for the linear logistic regression are as given in table 1.

As the table shows, all the factors are independently significant predictors of Topicalization. In odds-ratio terms: *mediated* nominals are 6.7 times more likely to be topicalized than *old* nominals, and *mediated* nominals are 1.2 times more likely to be topicalized than *new* ones. The first result is not surprising in the light of a theory of Topicalization such as that of Prince (1998). What is more surprising is the second: that new information tends to be in fronted position more often than old. This is discussed further in Snider (2006).

The big surprise, however, is that inanimates are 4.3 times more likely than animates to be in a topic-position. This result contradicts WOH rather directly, but it is also not in agreement with GFH: at first one might be tempted to attribute this result to the fact that most topicalized constituents are linked to non-subject in-situ elements. As animacy correlates strongly with grammatical function, with animates being attracted toward higher levels on the grammatical function hierarchy, one might be tempted to conclude that most topicalized constituents tend to be inanimate by virtue of their link to non-subject

TABLE 3

Factors	Coefficient	F-value	p-value
(Intercept)	-1.212015	360.0704	$< .0001$
animacy		3.8718	0.0492
inanimate vs animate	-0.420939		
status		27.3778	$< .0001$
old vs mediated	-1.961611		
new vs mediated	-0.270017		
gf		157.8188	$< .0001$
subj vs non-subj	-3.209837		
weight	0.401827	144.0611	$< .0001$

grammatical functions and interpret the results as being in favor of one way of amending GFH described in 15.2. But the independent effect of animacy shows that this cannot be the whole explanation for the prevalence of inanimates in topic position: the tests above show that animacy has an independent effect at $p = 0.0022$ and that the effect is that inanimates are favored in topic-position. If we look at the raw numbers given in table 2, we see that these also show that the distribution of animates and inanimates is very different in topic position from what it is in subject position, as we overwhelmingly find inanimates in topic position. The raw data also show that the distribution is different from that of animates and inanimates in object position: the proportion in topicalized position is 1 to 9 whereas for objects it is about 2 to 8.

With respect to weight, the results show that heavier constituents are more likely to be topicalized than light ones, again a result that goes against the grain of purely linear order-based accounts.

15.3.2 Left Dislocation

Information status is significant at $p < .0001$. Thus, it is clear that information status is a significant predictor of Left Dislocation, with *mediated*-coded entities most likely to left-dislocate. A *mediated* NP is 1.3 times more likely to be in a left-dislocated position than a *new* NP, and *mediated* NPs are 7.1 times as likely to be left-dislocated than *old* NPs. These results will be discussed further in Snider (2006). Here we just note that the behavior of *mediated* elements is as expected, but the ratio between *new* and *old* is somewhat surprising given most theories about Left Dislocation. Animacy and grammatical function are also significant factors. And the tests above show that, for this construction too, each of these factors has an independent effect, because each signif-

TABLE 4

Left Dislocation	LD	Subjects	
animates	227	8929	
inanimates	173	2665	$p < .001$

TABLE 5

Factors	Coefficient	F-value	p-value
(Intercept)	-2.40407	1866.1625	$< .0001$
animacy		2.7000	0.1004
inanimate vs animate	0.13064		
gf		516.2203	$< .0001$
subj	-3.03187		
weight	0.41943	1041.6754	$< .0001$

icantly increases the likelihood of the model when added individually. Grammatical function is significant at the $p < .0001$ level, and animacy is significant at the $p < 0.05$ level. In odds-ratio terms, animates are 1.5 times more likely than inanimates to be in a Left Dislocation construction. Here the results are weakly consistent with GFH. When we look at the raw numbers for left-dislocated elements and subjects in table 4, we see that indeed they are more similar, although the proportion of inanimates is higher in Left Dislocation.

The role of weight is similar to that in Topicalization.

15.3.3 Analysis of Larger Animacy Set

In order to further test the animacy effects, we analyzed a larger data set. This was possible because, as mentioned above, the animacy annotation has greater coverage than the information status annotation. In this data set, we were able to use all 399 Left Dislocations and 106 Topicalizations, but the models had fewer factors (only animacy, GF, and weight). Tables 5 and 6 show that using more data and fewer factors, the animacy effect in Left Dislocation disappears ($p > 0.1$) and strengthens for Topicalizations. These differences might just be a measure of the effect of information status that is now lost given that we now have a much poorer model. But they suggest that the anti-animacy effect in Topicalization is not a fluke due to the small dataset.

TABLE 6

Factors	Coefficient	F-value	p-value
(Intercept)	-3.237999	161.16239	$< .0001$
animacy		48.39225	$< .0001$
inanimate vs animate	1.972179		
gf		144.72518	$< .0001$
subj	-6.949687		
weight	0.318986	82.20584	$< .0001$

15.4 Discussion

There have been no previous studies that examined the effects of animacy on Left Dislocation and Topicalization. The above results suggest that there is no effect of animacy on left-dislocated NPs and show a tendency for topicalized NPs to be inanimate. The logistic regression shows that this inanimacy effect for Topicalizations is not merely due to the fact that they are extracted from a non-subject position, where inanimates are preferred, because this factor was included in the regression model.

One should not read too much into the results of one rather small study. We need to do a further analysis of the data to see whether other factors might play a role. For instance, the relative animacy of the subject and the fronted element might be important. Our impression is that in the Switchboard corpus, sentences with fronted elements tend to have animate subjects, but we have not counted them. Another, perhaps more promising hypothesis to pursue is the following: Topicalization involves a long distance dependency and one can hypothesize that as such it carries an extra processing load. It might be that to compensate for this, the argument structure of the sentences where it is used tends to be canonical, i.e. of the animate subject, inanimate object type. If there is such a tendency, it could lead to the results we found. Whatever the exact explanation, if the result stands, the tendency for inanimates to topicalize, documented here, is problematic for theories of production that predict the saliency of referents to directly influence linearization of NPs in the clause (Kempen and Harbush 2004). Such theories would predict that a construction that caused a referent to occur first in the clause would choose the most salient referent. Animate NPs are inherently more salient than inanimates, so these theories would predict that animates should topicalize more. The data in this study show that this is not the case. There is other evidence that a simple 'animate-first' theory is inadequate. For

instance, data from the ordering of temporal adjunct PPs shows that the animacy of the subject does not affect the realization of adjunct PPs before the subject or at the end of the clause (Cueni et al. 2005).

Our results do not support a simple version of the Grammatical Function Hierarchy hypothesis either: according to one version of that hypothesis there would be no effect on Topicalization or Left Dislocation of animacy, but we see that for Topicalization there is a negative effect. As we discussed, this effect cannot be due to the overwhelming influence of information status: even when information status is taken into account, the effect remains. If the explanation is a version of the processing load hypothesis we sketch above, there could be a version of the GFH that is compatible with it.

To make things more complicated, our results contradict WOH, but there are data, from languages other than English, in favor of WOH. Kempen and Harbush (2004) show that in the middle field in German, GFH does not hold but that the hypothesis that animacy is correlated with simple surface word order accounts for the data. One could reconcile the data of our study and those of Kempen and Harbush (2004) by proposing that the linear order effects occur only in the sentence internal domain, the IP domain, where subcategorization plays a direct role, and not in the periphery, the CP domain. This is descriptively adequate but it is puzzling for some psycholinguistic models.

From a psycholinguistic perspective, it is plausible to assume that salient entities are accessed and expressed first. However, here we see that elements sometimes occur first despite the fact that they should be less salient, according to the usual criteria of salience (animacy in this case). At first this may suggest that the influence of salience on ordering might be more construction bound than has been assumed: when a so-called unmarked order is used, the most salient element comes first but marked constructions can be used in which other elements are first. This way of interpreting the data, however, might call into question the very notion of salience that is used: an inanimate in topic position does not strike one as less salient, it rather strikes one as an non-typical salient element. The construction seems to treat as salient an element that is not salient by normal criteria.[5] We do not have a model of sentence production that would allow for such construction-dependent reversals in signaling salience, but it is clear that more attention needs to be paid to variation in syntactic constructions than is done in the production models currently proposed.

[5]Note also that order alone is not what achieves this effect: in Left Dislocation the effect of animacy is not the same as in Topicalization, but both are fronting constructions.

Acknowledgments

The second author thanks Ron for years of interesting discussion and collaboration on various linguistic issues. Thanks also to Tom Wasow and Anna Cueni for initial discussions and to Tom, two anonymous reviewers, and the participants in the Spring 2006 Laboratory Syntax Class for comments. All remaining errors are as usual our own.

References

Aissen, Judith. 2003. Differential object marking: Iconicity vs. economy. *Natural Language and Linguistic Theory* 21:435–483.

Arnold, Jennifer, Thomas Wasow, Anthony Losongco, and Ryan Ginstrom. 2000. Heaviness vs. newness: The effects of structural complexity and discourse status on constituent order. *Language* 76:28–55.

Bock, J. Kathryn, Helga Loebell, and Randal Morey. 1992. From conceptual roles to structural relations: Bridging the syntactic cleft. *Psychological Review* 99:150–171.

Bock, J. Kathryn and Richard K. Warren. 1985. Conceptual accessibility and syntactic structure in sentence formation. *Cognition* 21:47–67.

Bresnan, Joan, Anna Cueni, Tatiana Nikitina, and Harald Baayen. 2005. Predicting the dative alternation. In *Proceedings of the Royal Netherlands Academy of Science, Workshop on Foundations of Interpretation*. Amsterdam, The Netherlands.

Comrie, Bernard. 1989. *Language Universals and Linguistic Typology: Syntax and Morphology*. Chicago, IL: University of Chicago Press.

Croft, William. 1988. Agreement vs. case marking and direct objects. In M. Barlow and C. A. Ferguson, eds., *Agreement in Natural Language: Approaches, Theories, Descriptions*, pages 159–179. Stanford, CA: CSLI Publications.

Cueni, Anna, Neal Snider, and Annie Zaenen. 2005. Boundaries to the influence of animates. Paper presented at the LSA Annual Winter Meeting.

Dahl, Osten and Kari Fraurud. 1996. Animacy in grammar and discourse. In T. Fretheim and J. K. Gundel, eds., *Reference and Referent Accessibility*, pages 47–64. Amsterdam, The Netherlands: John Benjamins.

Garretson, Gregory, Mary C. O'Connor, Barbora Skarabela, and Marjorie Hogan. 2004. *Coding practices used in the project Optimal Typology of Determiner Phrases*. Boston, MA: Boston University.

Godfrey, John, Edward Holliman, and Jane McDaniel. 1992. SWITCHBOARD: Telephone speech corpus for research and development. In *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP'92)*, pages 517–520. San Francisco, CA.

Hawkins, John A. 1994. *A Performance Theory of Order and Constituency*. Cambridge, United Kingdom: Cambridge University Press.

Kempen, Gerard and Karin Harbush. 2004. A corpus study into word order variation in German subordinate clauses: Animacy affects linearization independently of grammatical function assignment. In T. Pechmann and C. Habel, eds., *Multidisciplinary approaches to language production*, pages 173–181. Berlin, Germany: Mouton De Gruyter.

Leech, Geoffrey, Brian Francis, and Xfueng Xu. 1994. The use of computer corpora in the textual demonstrability of gradience in linguistic theories. In C. Fuchs and B. Victorri, eds., *Continuity in linguistic semantics*, pages 57–76. Amsterdam, The Netherlands: John Benjamins.

Marcus, Mitch, Beatrice Santorini, and Mary A. Marcinkiewicz. 1993. Building a large annotated corpus of English: The Penn Treebank. *Computational Linguistics,* 19:313–330.

McDonald, Janet L., J. Kathryn Bock, and Michael H. Kelly. 1993. Word and word order: Semantic, phonological, and metrical determinants of serial position. *Cognitive Psychology* 25:188–230.

Newmeyer, Frederick. 2003. Grammar is grammar and usage is usage. *Language* 79:682–707.

Nissim, Malvina, Shipra Dingare, Jean Carletta, and Mark Steedman. 2004. An annotation scheme for information status in dialogue. In *Proceedings of the 4th Conference on Language Resources and Evaluation (LREC'04)*. Lisbon, Portugal.

O'Connor, Mary C., Arto Anttila, Vivienne Fong, and Joan Maling. 2004. Differential possessor expression in English: Re-evaluating animacy and topicality effects. Paper presented at the Annual Meeting of the Linguistic Society of America, Boston, MA.

Prat-Sala, Mercè. 1998. *The Production of Different Word Orders: A Psycholinguistic and Developmental Approach*. Edinburgh, United Kingdom: Doctoral dissertation, University of Edinburgh.

Prince, Ellen. 1981. Toward a taxonomy of given-new information. In P. Cole, ed., *Radical Pragmatics*, pages 223–256. New York, NY: Academic Press.

Prince, Ellen. 1998. On the limits of syntax, with reference to left-dislocation and topicalization. In P. Culicover and L. McNally, eds., *Syntax and semantics. Vol. 29. The limits of syntax*, pages 281–302. New York, NY: Academic Press.

Rosenbach, Anette. 2002. *Genitive Variation in English. Conceptual Factors in Synchronic and Diachronic Studies*. Berlin, Germany/New York, NY: Walter de Gruyter.

Rosenbach, Anette. 2003. Aspects of iconicity and economy in the choice between the s-genitive and the of-genitive in English. In G. Rohdenburg and B. Mondorf, eds., *Determinants of Grammatical Variation in English*. Berlin, Germany/New York, NY: Walter de Gruyter.

Silverstein, Michael. 1976. Hierarchy of features and ergativity. In R. M. W. Dixon, ed., *Grammatical Categories in Australian Languages*, pages 112–171. Canberra, Australia: Australian Institute of Aboriginal Studies.

Snider, Neal. 2006. Left dislocation and topicalization in English (working title). Unpublished Manuscript.

Snider, Neal and Annie Zaenen. 2006. Animacy and salience. Unpublished Manuscript.

Zaenen, Annie, Jean Carletta, Gregory Garretson, Joan Bresnan, Andrew Koontz-Garboden, Tatiana Nikitina, Mary C. O'Connor, and Tom Wasow. 2004. Animacy encoding in English: Why and how. In *Proceedings of the 42nd Annual Meeting of the Association for Computational Linguistics (ACL'04), Workshop on Discourse Annotation*, pages 118–125. Barcelona, Spain.