

THE SUBJECT OF ROAR IN THE MIND AND IN THE CORPUS: WHAT DIVERGENT RESULTS CAN TEACH US

John Newman
Monash University
University of Alberta

Tamara Sorenson Duncan
Dalhousie University
IWK Health Centre

Abstract

The use of different research methods in linguistics invariably leads to questions about the convergence and divergence of research findings. Aiming for convergence, while understandable, may distort our understanding of language phenomena, if convergence is seen as the only publishable result. We suggest a place for diverging results in furthering our understanding of the data techniques used to investigate linguistic phenomena. We illustrate this point through an experimental and corpus-based investigation of the preferred syntactic subjects of the English verb ROAR and discuss how deeper reflection on these diverging results leads to a better understanding of the different data types.

Key Words: corpus linguistics, convergence, divergence, research methods

1. INTRODUCTION

The increased use of multi-methodological research designs in linguistics, and cognitive linguistics in particular, is a welcome development in the field, providing linguists with a more complete understanding of linguistic phenomena than single-methodological research provides. A combination of lab-based experimental methods with speakers and corpus-based methods is especially favoured as a multi-methodological approach (e.g., Gilquin & Gries 2009, Klaven & Divjak, 2016). Combining research methods such as these can, however, lead to disparate kinds of research findings that are a challenge to reconcile, raising questions as to how to interpret results from different sets of data and different methods. Consequently, questions about convergence vs. divergence of results invariably arise in such research.

Seeking convergence of results is a familiar goal in research, across disciplines. Convergence of research findings has a key role to play in validating and verifying findings, and we do not wish to argue against searching for convergence which, in some real sense, seems fundamental to the scientific enterprise. The “Cambridge Declaration on Consciousness” (Low, Panksepp, Reiss, Edelman, van Swinderen & Koch, 2012) is a case in point. The Declaration may be seen as a compelling instance of how convergence across disciplines provides the basis for advancing knowledge, in this case in the scientific understanding of consciousness. The authors introduce the Declaration in the following terms:

Convergent evidence indicates that non-human animals have the neuroanatomical, neurochemical, and neurophysiological substrates of conscious states along with the capacity to exhibit intentional behaviours. Consequently, the weight of evidence indicates that humans are not unique in possessing the neurological substrates that generate consciousness... (Low, Panksepp, Reiss, Edelman, van Swinderen & Koch, 2012, our emphasis in bold).

This passage illustrates the critical role that convergent evidence has played, and continues to play, in establishing scientific knowledge and examples of this can be found in all disciplines.

To offer an example specifically from linguistics, consider the work of Divjak and Gries (2011). Their study of synonymy of nine Russian words for “try” used both experimental and corpus methods. Using these varied methods, they found converging evidence for three distinct clusters within their large synonym set. For the purposes of the present study, the relevant point from this example is their interpretation of this convergence. They state: “evidence as converging as in the present study, strengthen our account of linguistic phenomena as elusive as lexical semantics” (Divjak & Gries, 2011, p. 210).

We readily acknowledge the unique role that convergence of results plays in sciences and humanities, including linguistics. However, acknowledging the key role of convergence in research requires us also to acknowledge the potential dangers associated with a relentless quest for convergence. One could mention, for example, the danger of a “confirmation bias” whereby a researcher who is determined to establish convergence of results will indeed find convergence (without impugning any malicious intent to the researcher). There is also the danger of a “publication bias”, i.e., the preference for publishing replicated and significant results, as opposed to exploring the full range of methodologies and outcomes that are available in order to ascertain exactly which methodologies are best (cf. Francis 2012). Thus, we caution against casting convergence as the only worthwhile result to aim for when multiple data types and methods are involved.

We believe that there is also a valuable role for divergent results in advancing knowledge. It is through a reflection on divergent results that we may reach a deeper understanding of the characteristics of different data types and their associated methodologies. This applies particularly

in those cases where the data types are relatively new or where the methods are still being developed and refined (as in the case of corpora and new experimental methods). We are not the first to make such claims within the cognitive linguistics tradition, even if this position is relatively underappreciated in the field (cf. the balanced discussions of these issues in Arppe & Järvikivi 2007, Ellis & Simpson-Vlach 2009, and Mollin 2014).

The proliferation of methodologies that now form part of the repertoire of contemporary linguists, and not just cognitive linguists, has not been matched by a fuller appreciation and discussion of how researchers should reconcile results based on different data and methodologies in a principled way. As an illustration of this point, consider the in-depth discussion of alternative research methods (with implied alternative data types) to be found in Podesva and Sharma (2013), covering methods such as surveys, interviews, fieldwork, judgment data, experimental research design, corpora, working with historical texts etc. The issues that arise in reconciling results from employing a selection of these methods are barely acknowledged in the volume, with no more than a relatively brief summary of “mixed-method approaches” in the chapter by Abbuhl, Gass, and Mackey (2013, p. 124-125), concerned mainly with combining quantitative and qualitative methods.

In what follows, we report on divergent research outcomes from a small multi-methodological investigation into preferences for syntactic subjects of the verb ROAR in English, arguing that the divergence of results is instructive as far as deepening our understanding of the data and methods are concerned. This study comes from a larger research project investigating convergence in a set of near-synonyms (*BOWL*, *BELLOW*, *HOLLER*, *HOWL*, *ROAR*, *SCREAM*, *SCREECH*, *SHOUT*, *SHRIEK*, and *YELL*) (Newman & Sorenson Duncan, 2014, 2015; Sorenson Duncan & Newman, 2013). The larger study was inspired by previous work on synonymy (Dabrowska, 2009

and Divjak & Gries 2011). Specifically, Dabrowska (2009) investigated synonyms for bipedal motion (AMBLE, STAGGER, etc.) and offered evidence as to the extent to which words develop unique meanings from their typical collocation patterns. Using her sentence-elicitation task, we sought to build on this work with a different set of synonyms. In expanding on this work, we incorporated an investigation of convergence. As such, following, Divjak & Gries (2011), we also included a corpus study. Through the course of the larger research project we noted divergence in the results between the two methods in one-third of our synonym set (*HOWL*, *ROAR* and *SCREECH*). Each instance of divergence followed the same pattern as described here (Newman & Sorenson Duncan, 2014; Sorenson Duncan & Newman, 2013). In what follows, we offer a detailed description of this divergence, using the data for *ROAR* as an illustration.

2. THE PREFERRED SUBJECT OF ROAR: EXPERIMENTAL STUDY

For the experimental component, a simple sentence-elicitation task was used. Sentences can be expected to contain a verb used in a particular syntactic frame, allowing us to identify subject, objects, etc. Sentence elicitation is thus an effective way of probing speakers' preferences for syntactic subjects of a verb compared with, say, free word associations and, indeed, the method has been used to obtain normative data on the use of English verbs in syntactic frames (Connine, Ferreira, Jones, Clifton, & Frazier 1984; Roland & Jurafsky 2002). Sample sentences containing *ROAR* were elicited from 31 students at *the University of Alberta*. Throughout this paper, we use the convention of small caps (i.e., *ROAR*) to refer to the lemma and lower-case italics for the inflected forms (e.g., *roar*, *roars*, *roared*).

Participants were native-speakers of English and had lived in *Canada* since they were three years of age or younger. Participants had a mean age of 20.55 years (range: 17-57 years).

There were 7 male participants and 24 female. Research participants saw a word, appear on the screen and were asked to provide a sentence using that word. Each participant had three opportunities to construct a sentence with *roar*. Only instances where the target word (in any of its inflected forms) was used as a verb in an active construction were included in the analysis, resulting in 41 sentences. There was no pressure on the participants to complete the task quickly. Sample sentences obtained in this way are provided in (1). As can be seen in (1a) and (1e), the responses could be quite lengthy, with participants constructing quite colourful and involved story lines, utilizing a relatively rich vocabulary of nouns, verbs, adjectives etc. Other participants preferred a more minimal style, as in (1c), consisting of short sentences, a preference of pronouns, and limited vocabulary.

- (1) (a) Upon seeing the test results, the principal **roared** and kicked over his trash can, shouting something about disgrace and lazy children.
 (b) When my teacher is angry, instead of yelling he **roars**.
 (c) I'm a lion. Hear me **roar**.
 (d) The lion **roared** at the lioness in anger.
 (e) The lion **roared** viciously from within its cage, and the spectators backed away quickly save for one little boy who looked up in wonder.
 (f) Natalie pretended to **roar** like an animal

The complete set of syntactic subjects can be seen in Table 1. LION (12 instances, 29%) is clearly the most preferred subject noun in this table, with CROWD a distant second (4 instances). Only occurrences of an inflected form of a noun or pronoun are counted in this table, consistent with how we report on subjects in the corpus study. There are, in fact, additional indirect references to a 'lion' sense in the participants' responses, as in *I'm a lion. Hear me roar*, *The child likes to roar and pretend he's a lion*, *Can more animals roar than just a lion*, and the references to Simba, being a lion character in the film *The Lion King*. In pivot constructions such as *Hear me*

roar, we took *me* to be functioning as a subject of ROAR and counted it as an instance of the lemma I).

Table 1

Preferred subjects from elicitation task

Subject Head	Freq.
LION	12
CROWD	4
HE	3
CAT	2
CHILD	2
I	2
MONSTER	2
SIMBA	2
(YOU)	1
ANIMAL	1
DOG	1
DRAGON	1
IT	1
JET	1
LADY	1
NATALIE	1
PRINCIPAL	1
WAVE	1
WE	1
YOU	1

3. THE PREFERRED SUBJECT OF ROAR: CORPUS-BASED STUDY

We relied on the Corpus of Contemporary American English (COCA, Davies 2008-) for the corpus-based part of our study. COCA is representative of North American usage and includes roughly equal amounts of data from five genres: spoken, fiction, popular magazines, newspapers, and academic journals. At the time of the study (April 2017), the corpus consisted of approximately 520 million words covering the period 1990-2015. We searched on the lemma ROAR (at the time of the searches, this was entered as [roar].vv* in the COCA interface), retrieving as our proxy for subject the noun (entered as nn*) in the position immediately to the left of the verb (“L1” position). One could consider alternative windows of context to search in, e.g.,

2, 3, or 4 positions to the left of the verb. We explored a number of these possibilities and found the results to be similar as far as the ranking of LION is concerned. An issue with the larger window sizes concerns noun phrases in subject position consisting of multiple nouns, such as *Hurricane Katrina*, *jet engine*, and others. Both *Hurricane* and *Katrina* in the case of *Hurricane Katrina* are returned when the window is 2 words to the left of the verb and consequently counted as separate items in the list of collocates which introduces some confusion in tallying the results. Applying a window of just one word to the left of the verb means that we count an instance of the phrase *Hurricane Katrina* just once, as *Katrina*. Sample sentences retrieved in this way are shown in (2).

- (2) (a) As the Joe Louis Arena crowd **roared**, Crosby headed for the dressing room.
 (b) The crowd **roars** approval, boosting his confidence.
 (c) As a spotlight swiped the clearing and a huge helicopter **roared** overhead
 (d) One flick of the switch and the vacuum **roared** to life.
 (e) Will the lion **roar** again?
 (f) You **roar** like a lion.

Table 2 shows the raw frequencies of nouns in the L1 position in the total corpus. The noun collocates were grouped into lemma classes for the purposes of this study. So, for example, the 42 instances of TRUCK in Table 2 include both *truck* (32) and *trucks* (10), just as ROAR finds *roar*, *roared*, *roaring*, and *roars*. LION appears as the seventh most frequent collocate in this table. Appendix A provides a breakdown of top collocates in the sub-genres of COCA. In none of these genres does LION appear as the most frequent collocate and in fact it only appears with frequency >5 in two genres, magazine (where LION is ranked third most frequent with 15 occurrences) and fiction (LION appears at rank 18 with 13 occurrences). As such, we found convergence amongst the varied genres in COCA, suggesting that LION is not a preferred subject of ROAR.

Table 2.

Top L1 collocates of ROAR in all COCA

Rank	L1 Collocate	Freq. in CX
1	CROWD	161
2	ENGINE	140
3	CAR	85
4	FIRE	54
5	WIND	48
6	TRUCK	42
7	LION	39
8	AUDIENCE	38
9	TRAIN	32
10	FLAME	31
11	PLANE	30
12	MOTORCYCLE	27
13	BUS	24
14	JET	23
15	WATER	23
16	MAN	22
17	VOICE	22
18	HELICOPTER	19
19	MOTOR	18
20	FAN	17

NOTE: L1 Collocate is the position immediately left of the verb.

Beyond the raw frequencies of occurrence, it is worthwhile to consider more statistically revealing measures of the attraction of words to constructions. Typically, the preferred measures take some account of the overall frequency of occurrence of the word and/or the construction. Two pertinent measures that do this are Mutual Information (MI) and Reliance, as used by Schmid (2000 [see especially p. 54-55], 2010). Mutual Information, is well known in corpus linguistic studies. It measures the strength of the association of two co-occurring forms based on a logarithmic value of the observed probability of the co-occurrence of the forms with the expected probability of the two forms co-occurring (Church & Hanks 1990). Reliance, also referred to as Relevance in the COCA interface, measures the extent to which a word appears in one particular

pattern (the noun+ROAR pattern) compared with the word's frequency in all patterns in the corpus. Measures of reliance can also be found in several other studies that use corpora (e.g., Gries, Hampe, & Schönefeld, 2005; Janda & Solovyev, 2009; Schmid & Küchenhoff, 2013). Furthermore, the "Faith(fullness)" score returned in Gries's (2007) collocation analyses is based on the same proportional calculation as Schmid's Reliance measure. Both Mutual Information and Reliance are automatically calculated in the COCA interface and the values returned by COCA are shown in Table 3, which includes L1 noun collocates of ROAR occurring with a minimum frequency in the construction of 5. As can be seen, the rankings based on each of these measures are identical, and LION is ranked 9th by these measures. Appendix B provides a breakdown of top collocates in the two sub-genres of COCA with the highest number of LION L1 collocates of ROAR, magazine and fiction. The frequencies of LION L1 collocates are too low (<5) to be considered further in the other genres.

Table 3.
Top L1 collocates of ROAR in COCA, ranked by Reliance (Relevance) and Mutual Information

Rank	L1 Collocate	Freq. as collocate of		Reliance/Relevance	MI
		ROAR in COCA	Freq. in COCA		
1	TYRANNOSAUR	6	163	3.68	11.71
2	BATMOBILE	6	288	2.08	10.89
3	JETLINER	6	529	1.13	10.01
4	MOTORCYCLE	27	4956	0.54	8.96
5	ENGINE	140	30371	0.46	8.72
6	CHOPPER	8	1980	0.4	8.53
7	CROWD	161	46898	0.34	8.29
8	WILDFIRE	6	1817	0.33	8.23
9	LION	39	13309	0.29	8.06
10	FLAME	31	13041	0.24	7.76
11	TORNADO	12	5479	0.22	7.64
12	FURNACE	5	2550	0.2	7.48
13	THUNDER	11	5837	0.19	7.42
14	JEEP	8	4272	0.19	7.42
15	BEAST	15	8177	0.18	7.39
16	HELICOPTER	19	12630	0.15	7.1
17	JET	23	18414	0.12	6.83
18	SAW	5	4430	0.11	6.69
19	DRAGON	8	7373	0.11	6.63
20	CANNON	5	4746	0.11	6.59

NOTE: L1 Collocate is the position immediately left of the verb; MI = mutual information score

A further measure to explore is *collostructional strength* as proposed by Stefanowitsch and Gries (2003). In this case, the measure is based on calculations of expected co-occurrence, given the total frequencies of the L1 noun collocate and total number instances of the noun + ROAR construction in COCA. To estimate the collostructional strength of specific nouns with ROAR, the search term `nn* [roar].vv*` was used in the COCA interface to determine the frequency of the noun + roar construction (=2,162)¹. The search term `vv*` was used to determine the total number of (verbal) constructions in the corpus (=59,997,668). One could imagine somewhat different

¹ Note the syntax for conducting such searches has been updated since the time of our data collection. For example, an underscore is now used in the place of the period. We refer interested readers to the COCA website: <http://corpus.byu.edu/coca/>.

ways of determining the number of relevant constructions in the corpus, given that we are exploring noun + verb combinations (e.g. one could use the total number of noun constructions or the total number of words in the corpus). However, the rankings according to collostructional strength would be the same regardless of the actual number chosen for the total number of constructions in the corpus. The negative log, base 10, of the probability of the co-occurrence was used to calculate the collostructional strength. The results are reported in Table 4, showing LION in the 4th top position. While this is a higher rank than in Tables 2 or 3, the collostructional strength value of LION (59.1) is considerably less than for the top two collocates CROWD (253.6) and ENGINE (238.2).

Table 4.

Top collocates by collostructional strength, for L1 noun collocates of ROAR in all COCA

Rank	L1 Collocate	Freq. a collocate of ROAR in COCA	Freq. in COCA	Expected freq.	Relation	Collostructional Strength
1	CROWD	161	46898	1.7	attraction	253.6
2	ENGINE	140	30371	1.1	attraction	238.2
3	CAR	85	172083	6.2	attraction	64.4
4	LION	39	13309	0.5	attraction	59.1
5	MOTORCYCLE	27	4956	0.2	attraction	48.4
6	WIND	48	64355	2.3	attraction	44.8
7	TRUCK	42	42049	1.5	attraction	44.4
8	FLAME	31	13041	0.5	attraction	44.4
9	FIRE	54	111423	4	attraction	40.7
10	AUDIENCE	38	47667	1.7	attraction	36.7
11	PLANE	30	42886	1.5	attraction	27.5
12	JET	23	18414	0.7	attraction	26.8
13	TRAIN	32	60940	2.2	attraction	25.5
14	HELICOPTER	19	12630	0.5	attraction	23.8
15	BUS	24	34055	1.2	attraction	22.2
16	BEAST	15	8177	0.3	attraction	20.2
17	MOTOR	18	21235	0.8	attraction	18.2
18	TIGER	16	15481	0.6	attraction	17.6
19	TORNADO	12	5479	0.2	attraction	17.2
20	TYRANNOSAUR	6	163	0	attraction	16.3

NOTE: L1 Collocate is the position immediately left of the verb.

4. DISCUSSION

Using an elicitation task and corpus searches, this study investigated the subject preferences for the verb ROAR. In the case of the elicitation data, an overwhelming preference for LION as the subject was found. This result, incidentally, converges with free word association norms which indicate a strong preference for ROAR and LION to be associated (Nelson, McEvoy & Schreiber, 1998, 2004). In the case of the corpus data, the ranking of LION varied, as one might expect, depending on which part of the corpus was searched and which measure was applied. Regardless of these considerations, LION never occurred as the most frequent subject of ROAR or as the most significantly associated subject of ROAR. These findings suggest that, in terms of subject preferences for this verb, the two methodologies do not yield converging results. In the remainder of this section, we discuss the significance of these results and argue that the diverging outcomes reflect quite different, but equally valid, kinds of linguistic realities.

The preferred use of LION as the subject of ROAR in the elicitation data presumably points to a kind of prototypical meaning that speakers have assigned to ROAR. The elicitation tasks required participants to create sentences, out of context, relying on the participants' own imagination. The preference for LION to appear as the subject in their sentences points to a strong conceptual linking of LION to ROAR. Kuperman, Stadthagen-Gonzalez and Brysbaert (2012) suggest that ROAR is an early acquired word, estimating its age of acquisition to be 4.79 years, which suggests that language use in the preschool years may provide valuable insights into the formation of speakers' prototypical knowledge of ROAR. Thus, to throw more light on the use of ROAR in language acquisition, we consulted the CHILDES database to identify the most common subject preferences for ROAR in child-directed speech (MacWhinney, 2000). We constructed a 1,269,389 word corpus from the Bates, Bliss, Bohannon, Brown, Clark and Demetras and Gleason

corpora available through CHILDES. These data represent conversations between children (aged 2-5 years) and a parent. Through these conversations we found 36 instances of ROAR, and in 30 of these instances it was a LION responsible for the ROAR. It is worth noting that because of the limited instances of ROAR in this corpus, we did not restrict our search to verbal instances. Nevertheless, these data illustrate a strong co-occurrence bias for ROAR and LION in child-centered conversations. Apart from this corpus-based support for the association of LION and ROAR in the pre-school years, especially through caregiver-child interactions, there is abundant evidence for this association of words and the concepts behind them in terms of more general cultural influences. These influences include the easy availability of books, films, and videos for children in which animals are featured performing typical animal acts, often intended as allegorical stories intended to illuminate some aspect of human life and society. Lions, in particular, find their way into such stories, as can easily be confirmed through a search of children's literature on the internet. Exposure to such material undoubtedly reinforces associations such as LION and ROAR, DOG and BARK etc. In short, the roaring of lions is something we learn about, typically, as children, indeed quite young children, and it remains a key part of our concept of ROAR, readily invoked when asked to make a sentence with ROAR, as in the elicitation task used here.

The corpus data gave rise to quite different results for the most frequently used or most strongly associated subject nouns of ROAR. LION was not the most highly ranked subject of ROAR. Rather, we find words such as CROWD, FIRE, ENGINE, WIND as the most significant subject nouns. Why should this be the case? The key to understanding this result lies in appreciating just what COCA does and does not represent. The texts which underlie COCA are overwhelmingly texts (including transcripts of spoken language) that come from adult use of North American varieties of English, obtained from books, journals, magazines, and newspapers that are intended for

consumption by adults. The spoken component is based largely on transcripts of TV programs in which adults converse with one another. Although the topics of these texts are wide-ranging, for the most part the texts are about those things that are most familiar to North Americans. Planes, engines, trucks, hurricanes etc. all belong to this world, lions and tigers less so. ROAR in COCA is therefore more likely to be used in its figurative uses to describe the former, rather than being used in its most literal sense to describe the latter. The language use of children or children interacting with their caregivers does enter into COCA but only in the most incidental way, e.g. through the reported speech of children in, say, the magazine genre, the imagined conversations involving children in the fiction genre, etc. For more private, home-based use of language, where LION and ROAR come into play, one has to turn to more “specialized” corpora such as the CHILDES corpora referred to above.

In sum, the two techniques employed in this study yielded divergent results. Through the careful consideration of this divergence, we argue that these methods capture distinct, but equally valid, kinds of linguistic realities (in this case sub-senses of the meaning of ROAR). In the case of the elicitation task, where participants generated written sentences without context, our results reflect a prototypical meaning for the word ROAR, likely established in childhood. In the case of the corpus data, which comes from naturalistic language use where a word is used as part of a larger message the speaker/writer wishes to convey, our results reflect typical usage and figurative extensions that are common in the everyday speech of adult, North American speakers of English.

We are not the first to note that psychological experiments are likely to yield divergence from naturalistic language use. In their discussion of data based on sentence-elicitation vs. corpus frequency, Roland and Jurafsky (2002: 334) comment on an expected divergence of results: “Because of the biases inherent in isolated sentence production, we should not expect results from

such psychological experiments to directly match natural language use.” The biases that Roland and Jurafsky (2002) refer to are ones that, as it happened, played no decisive role with the particular verb investigated here. They refer to differences in preferences such as passivization (higher in the connected discourse of a corpus), and use of first person subjects (higher in elicited sentences than in a written corpus). In order to examine if similar preferences exist in our data, we considered the use of first person subjects. Although there were instances of personal pronouns functioning as the subject of ROAR in our elicited data, these were too infrequent to compete with the overwhelming preference for full nouns as the subjects of the verb. As such, it seems that the divergence in our results is not tied to preferences for one construction type over the other. Instead, our results reveal a divergence that centers around the usage and meaning of the word ROAR. As noted above, each technique yielded results which highlighted one sub-sense of the verb ROAR. Thus, our focus on the different lexical choices for syntactic subjects of ROAR in an active construction highlight a more revealing distinction between experimental and corpus techniques than just what one might expect from eliciting sentences in isolation vs. in context.

6. CONCLUSION

We suggested at the outset that divergent results can serve to deepen our knowledge about data types and their associated methods. We believe this is the case with our investigation into the verb ROAR. Broadly speaking, our results can be seen as evidence of the different realities of sentence-elicitation data and corpus data. Sentence elicitation tasks can tap into deeply entrenched uses that may stretch back into childhood; corpora, such as COCA, reflect current uses by adults in relatively public domains. Neither of these points has been fully appreciated, we believe. The divergence that we found leads us to a finer appreciation of the characteristics of these data types.

References

- Abbuhl, Rebekha, Susan Gass, & Alison Mackey. (2013). Experimental research design. In Robert J. Podesva and Devyani Sharma (eds.), *Research Methods in Linguistics*. 116-134. Cambridge: Cambridge University Press.
- Arppe, Antti. & Järvikivi, Juhani. (2007). Every method counts: Combining corpus-based and experimental evidence in the study of synonymy. *Corpus Linguistics and Linguistic Theory* 3, p. 131-159.
- Bdar, Mario., Gries, Stefan Th. & Fuchs, Milena, Z. (2011). *Cognitive Linguistics: Convergence and Expansion*. Amsterdam: John Benjamins Publishing Co.
- Church, Kenneth W. & Hanks, Patrick (1990). Word association norms, mutual information and lexicography. *Computational Linguistics* 16.1:22-29.
- Connine, Cynthia, Ferreira, Fernanda, Jones, Charlie, Clifton, Charles & Frazier, Lyn (1984). Verb frame preference: Descriptive norms. *Journal of Psycholinguistic Research* 13.4, 307-319.
- Dabrowska, E. (2009). Words as constructions. In V. Evans & S. Pourcel (Eds.), *Human cognitive processing: Vol. 24. New directions in cognitive linguistics*. Amsterdam, Netherlands: John Benjamins Publishing Company, p. 201-223.
- Davies, Mark. (2008-) *The Corpus of Contemporary American English: 425 million words, 1990-present*. Available online at <http://corpus.byu.edu/coca/>.
- Davies, Mark. (2013). Corpus of Global Web-Based English: 1.9 billion words from speakers in 20 countries. URL: <http://corpus2.byu.edu/glowbe>.

- Divjak, D. & Gries, S. Th. (2011). Clusters in the mind? Converging evidence from near synonymy in Russian. *The Mental Lexicon*, 3, 188-213.
- Ellis, Nick C. & Simpson-Vlach, Rita. (2009). Formulaic language in native speakers: Triangulating psycholinguistics, corpus linguistics and education. *Corpus Linguistics and Linguistic Theory*, 5, 61-78.
- Francis, Gregory. (2012). Too good to be true: Publication bias in two prominent studies from experimental psychology. *Psychonomic Bulletin & Review* 19: 151-156.
- Gilquin, Gaëtanelle & Gries, Stehan Th. (2009). Corpora and experimental methods: A state-of-the-art review. *Corpus Linguistics and Linguistic Theory* 5(1), 1–26.
- Gries, Stehan Th. (2007). Collostructional analysis: Computing the degree of association between words and words/constructions. Coll.analysis V 3.2a in R.
- Gries, Stehan Th., Hampe, Beate & Schönefeld, Doris (2005a). Converging evidence: Bringing together experimental and corpus data on the association of verbs and constructions.
- Gries, Stehan Th., Hampe, Beate & Schönefeld, Doris (2005b). Converging evidence II: More on the association of verbs and constructions. In John Newman & Sally Rice (eds.), *Empirical and experimental methods in cognitive/functional research*, 39–72. Stanford: CSLI Publications.
- Klaven, Jane & Divjak, Dagmar (2016). The cognitive plausibility of statistical classification models: Comparing textural and behavioral evidence. *Folia Linguistica*, 20, 35384.
- Kuperman, Victor, Stadthagen-Gonzalez, Hans & Brysbaert, Marc (2012). Age-of-acquisition ratings for 30 thousand English words. *Behav Res Methods*, 44, p. 978-990.
- Low, Philip, Panksepp, Jaak, Reiss, Diana, Edelman, David, van Swinderen, Bruno & Koch, Christof (2012, July). The Cambridge declaration on consciousness. In *Francis Crick*

Memorial Conference, Cambridge, England. url

<http://fcmconference.org/img/CambridgeDeclarationOnConsciousness.pdf>

MacWhinney, Brian (2000). *The CHILDES project: The database* (Vol. 2). Psychology Press.

Mollin, Sandra (2009). Combining corpus linguistic and psychological data on word co-occurrences: Corpus collocates versus word associations. *Corpus Linguistics and Linguistic Theory*, 5(2), 175-200.

Mollin, Sandra (2014). *The (Ir)Reversibility of English Binomials: Corpus, Constraints, Developments*. Amsterdam: John Benjamins Publishing Co.

Nelson, Douglas L., McEvoy, Cathy L., & Schreiber, Thomas A. (1998). *The University of South Florida word association, rhyme, and word fragment norms*.

<http://www.usf.edu/FreeAssociation/>.

Nelson, Douglas L., McEvoy, Cathy L. & Schreiber, Thomas A. (2004). The University of South Florida free association, rhyme, and word fragment norms. *Behavior Research Methods, Instruments, & Computers*. 36(3): 402-7.

Newman, J. & Sorenson Duncan, T. (2014, November). Howling wolves and roaring lions: What speakers think and what a corpus tells us. *Conceptual Structure, Discourse and Language 2014*. University of California Santa Barbara, Santa Barbara, CA.

Newman, J. & Sorenson Duncan, T. (2015, July). Convergence and divergence in cognitive linguistics: Facing up to alternative realities of linguistic categories. *International Cognitive Linguistics Conference*, Northumbria University, Newcastle England.

Podesva, Robert J. & Sharma, Devyani. (2013). *Research Methods in Linguistics*. Cambridge: Cambridge University Press.

- Roland, Douglas & Jurafsky, Daniel (2002). Verb sense and verb subcategorization probabilities. In Suzanne Stevenson and Paolo Merlo (eds.), *The Lexical Basis of Sentence Processing: Formal, computational, and Experimental Issues*, 325-346. Amsterdam: John Benjamins.
- Schmid, Hans-Jörg. (2000). *English abstract nouns as conceptual shells. From corpus to cognition*, Berlin/New York: Mouton de Gruyter.
- Schmid, Hans-Jörg. (2010). Does frequency in text instantiate entrenchment in the cognitive system? In D. Glynn and K. FISCHER (Eds.), *Quantitative methods in cognitive semantics*, 101-133. Berlin and New York: Mouton de Gruyter.
- Schmid, Hans-Jörg and Küchenhoff, Helmut (2013). Collostructional analysis and other ways of measuring lexicogrammatical attraction: Theoretical premises, practical problems and cognitive underpinnings. *Cognitive Linguistics* 24.3: 531-577.
- Schönefeld, Doris (ed). 2011. *Converging Evidence: Methodological and Theoretical Issues for Linguistic Research*. Amsterdam: John Benjamins Publishing Co.
- Sorenson Duncan, T. & Newman, J. (2013, June). English verbs of “yelling”: A quantitative corpus-based exploration of synonymy. *International Cognitive Linguistics Conference*, University of Alberta, Edmonton, AB.
- Stefanowitsch, Anatol. & S. Th. Gries, Stefan Th. (2003). Collostructions: investigating the interaction between words and constructions. *International Journal of Corpus Linguistics* 8.2: 209-43.
- Wulff, Stefanie (2009). Converging Evidence from corpus and experimental data to capture idiomaticity. *Cognitive Linguistics and Linguistic Theory*, 5, 131-159.

Wulff, Stefanie, Ellis, Nick, C., Römer, Ute, Bardovi-Harlig, Kathleen & LeBlanc, Chelsea, J. (2009). The acquisition of tense-aspect: Converging evidence from corpora and telicity ratings. *The Modern Language Journal*, 93, 354-369.

Appendix A

Frequencies of L1 collocates of ROAR in the sub-genres of COCA

Table A.1.

Results for frequencies >12 in the FICTION genre

Rank	L1 Collocate	Freq.
1	ENGINE	110
2	CAR	75
3	CROWD	68
4	FIRE	34
5	TRUCK	33
6	WIND	27
7	TRAIN	22
8	BUS	20
9	MOTORCYCLE	18
10	AUDIENCE	17
11	FLAME	17
12	WATER	17
13	MAN	15
14	MOTOR	15
15	VOICE	15
16	BEAST	13
17	HELICOPTER	13
18	LION	13
19	PLANE	13

Table A.2.
Results for frequencies >3 from the MAGAZINE genre

Rank	L1 Collocate	Freq.
1	CROWD	41
2	WIND	16
3	LION	15
4	ENGINE	12
5	AUDIENCE	10
6	CAR	6
7	FIRE	6
8	FLAME	6
9	MARKET	6
10	PLANE	6
11	TORNADO	6
12	FAN	5
13	TRUCK	5
14	BLOOD	4
15	TRAIN	4

Table A.3.
Results for frequencies >2 in the NEWSPAPER genre

Rank	L1 Collocate	Freq.
1	CROWD	43
2	ENGINE	12
3	AUDIENCE	10
4	PLANE	10
5	FAN	7
6	JET	7
7	MOTORCYCLE	6
8	FIRE	5
9	TRAIN	5
10	LION	4
11	HELICOPTER	4
12	MAN	3
13	STORM	3
14	TRAFFIC	3
15	VOICE	3
16	WILDFIRE	3
17	WIND	3
18	WOMAN	3

Table A.4.
Results for frequencies >1 from the SPOKEN genre

Rank	L1 Collocate	Freq.
1	FIRE	7
2	FLAME	6
3	CROWD	4
4	JET	4
5	MARKET	4
6	LION	3
7	STREET	3
8	ENGINE	2
9	ECONOMY	2
10	HELICOPTER	2
11	TORNADO	2
12	WIND	2
13	WILDFIRE	2
14	WAVE	2

Table A.5.
Results for frequencies >1 in the ACADEMIC genre

Rank	L1 Collocate	Freq.
1	CROWD	5
2	ENGINE	4
3	LION	4
4	PEOPLE	2
5	TORNADO	2
6	FIRE	2

Appendix B

Reliance scores for L1 collocates of ROAR

Table B.1.

Word association measures for L1 collocates (>5 occurrences) of ROAR in the FICTION genre

Rank	L1 Collocate	Freq. as collocate of ROAR in FICTION	Freq. in FICTION	Reliance
1	TYRANNOSAUR	6	124	4.84
2	BATMOBILE	6	243	2.47
3	MOTORCYCLE	18	1250	1.44
4	ENGINE	110	7716	1.43
5	CHOPPER	7	1006	0.7
6	TIGER	12	1865	0.64
7	MOTOR	15	2388	0.63
8	HELICOPTER	13	2373	0.55
9	LION	13	2893	0.45
10	CROWD	68	16868	0.4
11	AUDIENCE	17	4973	0.34
12	JET	9	2799	0.32
13	BEAST	13	4137	0.31
14	THUNDER	9	3095	0.29
15	JEEP	6	2111	0.28
16	FLAME	17	6372	0.27
17	TRUCK	33	15985	0.21
18	PICKUP	6	3029	0.2
19	BUS	20	10471	0.19
20	DRAGON	8	4293	0.19

Table B.2.

Reliance scores for L1 collocates (≥ 5 occurrences) of ROAR in the MAGAZINE genre

Rank	L1 Collocate	Freq. as collocate of ROAR in MAGAZINE	Freq. in MAGAZINE	Reliance
1	TORNADO	6	1165	0.52
2	CROWD	41	8853	0.46
3	LION	15	3456	0.43
4	FLAME	6	2317	0.26
5	AUDIENCE	10	7795	0.13
6	ENGINE	12	11389	0.11
7	WIND	16	17149	0.09
8	TRUCK	5	7559	0.07
9	PLANE	6	9496	0.06
10	FAN	5	11567	0.04
11	FIRE	6	20802	0.03
12	MARKET	6	36172	0.02
13	CAR	6	37172	0.02