Discourse Topics and Metaphors

Beata Beigman Klebanov
Northwestern University
beata@northwestern.edu

Eyal Beigman
Washington University in St. Louis
beigman@wustl.edu

Daniel Diermeier
Northwestern University
d-diermeier@northwestern.edu

Abstract

Using metaphor-annotated material that is sufficiently representative of the topical composition of a similar-length document in a large background corpus, we show that words expressing a discourse-wide topic of discussion are less likely to be metaphorical than other words in a document. Our results suggest that to harvest metaphors more effectively, one is advised to consider words that do not represent a discourse topic.

Traditionally, metaphor detectors use the observation that a metaphorically used item creates a local incongruity because there is a violation of a selectional restriction, such as providing a non-vehicle object to the verb *derail* in *Protesters derailed the conference*. Current state of art in metaphor detection therefore tends to be "localistic" – the distributional profile of the target word in its immediate grammatical or collocational context in a background corpus or a database like WordNet is used to determine metaphoricity (Mason, 2004; Krishnakumaran and Zhu, 2007; Birke and Sarkar, 2006; Gedigian et al., 2006; Fass, 1991).

However, some theories of metaphor postulate certain features of metaphors that connect it to the surrounding text beyond the small grammatical or proximal locality. For example, for Kittay (1987) metaphor is a discourse phenomenon; although the minimal metaphoric unit is a clause, often much larger chunks of text constitute a metaphor. Consider, for example, the TRAIN metaphor in the following excerpt from a Sunday Times article on 20 September 1992:

> Thatcher warned EC leaders to stop their endless round of summits and take notice of their own people. "There is a fear that the European train will thunder forward, laden with its customary cargo of gravy, towards a destination neither wished for nor understood by electorates. But the train can be stopped," she said.

In the example above, the quotation is not in itself a metaphor, as there is no indication that something other than the actual train is being discussed (and so no local incongruities exist). Only when situated in the context prepared by the first sentence (and indeed the rest of the article), the train imagery becomes a metaphor.

According to Kittay, a metaphor occurs when a semantic field is used to discuss a *different* content domain. The theory therefore predicts that a metaphorically used semantic domain would be off-topic in the given document.

Although a single document can have singular, idiosyncratic topics, it is likelier to discuss a mix of topics that are typical of the discourse of which it is part. We therefore derive the following hypothesis: Words in a given document that represent a common topic of discussion in a corpus of relevant documents would be predominantly non-metaphorical. That is, a smaller share of metaphorically used words in a document would fall in such topical words than the share of topical words in the document.

We test this hypothesis in the current article.

Using a large background corpus, we estimate the topical composition of the target documents (section 1) that were annotated for metaphors (section 2). We then report the results of the experiment (section 3) that strongly support the hypothesis, and discuss the findings (section 4). The concluding section provides a summary and outlines the significance of the results for the practice of metaphor detection.

1 Topic identification

1.1 EUI corpus

Our aim was to create a large corpus of British media discourse regarding the emerging European Union institutions, with both Euro-phile and Euro-sceptic camps represented. Our corpus consists of 12,814 articles drawn from three British newspapers: *The Guardian* (34%), *The Times* (38%), and *The Independent* (28%), dating from 1990 to 2000. We used LexisNexis Academic[1] to search for the Subject index term *European Union Institutions* (henceforth, *EUI*).[2] After results are retrieved, we further narrow them down to only documents on the subject *European Union Institutions* in the detailed subject index of the retrieved results.[3,4]

1.2 Identification of discourse topics

We converted all 12,858 documents[5] (henceforth, **EUI+M** corpus) into plain text format and removed

---

[1]http://academic.lexisnexis.com/online-services/academic-features.aspx

[2]In LexisNexis subject index hierarchy: Government and Public Administration/International Organization and Bodies/International Governmental Organizations/European Union Institutions.

[3]In the initial search, an article that scores 72% on the subject would be retrieved, but it would not be classified as being on this subject, and so would not be included in the final dataset. Articles in the final dataset tend to score about 90% on the subject, according to LexisNexis index.

[4]There is a gap in LexisNexis' index coverage of *The Times* during 1996-7 and of *The Independent* during 2000. To avoid under-representation of the newspaper and of the relevant years in the sample, we added articles returned for the search SECTION(Home news) AND (European Union OR Brussels) on *The Times* 01/1996 through 04/1998, and SECTION(News AND NOT Foreign) AND (European Union OR Brussels) on *The Independent* throughout 2000.

[5]12,814 EUI corpus plus 44 documents annotated for metaphors, to be described in section 2.

words from a list of 153 common function words. We then constructed an indexing vocabulary **V** that included all and only words that (a) contained only letters; and (b) appeared at least 6 times in the collection. All documents were indexed using this 21,046 word vocabulary. We will designate all the indexed words in document i as $\mathbf{D}_i$.

To identify the main discourse topics in the EUI+M corpus, we submitted the indexed documents to an unsupervised clustering method Latent Dirichlet Allocation (Blei et al., 2003) (henceforth, **LDA**).[6] The designation of the clusters as topics is supported by findings reported in Blei et al. (2003) that the clusters contain information relevant for topic discrimination. Additionally, Chanen and Patrick (2007) show that LDA achieves significant correlations with humans on a topic characterization task, where humans produced not just a topic classification but also identified phrases they believed were indicative of each class.

Using the default settings of LDA implementation,[7] we analyzed the corpus into 100 topics. Table 1 exemplifies some of the emergent topics.

1.3 Topical words in a text

LDA is a generative model of text. According to its outlook, every text is about a small (typically 5-7) number of topics, and each indexed word in the text belongs to one of these topics. However, in many cases, the relationship between the word and the topic is quite tentative, as the word is not particularly likely given the topic. We therefore use parameter k to control topic assignments – we only take LDA's assignment of word to topic if the word is in the top k most likely words for that topic. For k=25, about 15% of in-vocabulary words in a document are assigned to a topic; for k=400, about half the in-vocabulary words are assigned to some topic. We designate by $\mathbf{T}_i^k$ all indexed words in document i that are assigned to some topic for the given value of k. The ratio $\frac{|T_i^k|}{|D_i|}$ describes the proportion of discourse topical words in the indexed words for the given document.

---

[6]No stemming was performed.

[7]downloaded from http://www.cs.princeton.edu/~blei/lda-c/

Table 1: Examples of topics identified by LDA in the EUI+M corpus. All words are taken from top 25 most likely words given the topic. We boldface one word per cluster, that could provide, in our view, an appropriate label for the cluster.

foreign nato military war russian **defence** soviet piece un kosovo sanctions bosnia moscow
rail tunnel **transport** train pounds channel eurostar ferry trains passengers services paris eurotunnel
countries europe **enlargement** new membership members eastern conference reform voting summit commission foreign join poland negotiations
parliament mep party socialist strasbourg christian vote leader labour conservative right political green democrat elections epp
television commission satellite tv broadcasting tickets film broadcasters bbc programmes **media** industry channel public directive
court article justice member directive treaty question provisions case law regulation judgment interpretation rules order proceedings
social workers **employment** working hours jobs week employers legislation unions employees chapter rights health minimum
bank central euro **monetary** rates currency interest bundesbank markets economic exchange finance inflation dollar german
players **football** clubs uefa league fifa game cup
fishing fish fishermen fisheries quota vessels boats waters sea fleet
racism racist ethnic xenophobia black minorities jury discrimination white relations
drugs patent research human companies genetic scientists health medical **biotechnology** disease
children parents punishment school rights family **childcare** corporal education law father mother
controls **immigration** border asylum checks passport police citizens crime europol
energy nuclear emissions oil electricity gas environment carbon tax pollution fuel global cut
commission fraud commissioners brussels report allegations officials inquiry meps **corruption** mismanagement staff santer

2 Metaphor annotation

Ideally, we should have sampled a small sub-corpus from the EUI corpus for metaphor annotation; however, the choice of the data for annotation predated the construction of the EUI corpus.

Our interest being in the way metaphors used in public discourse help shape attitudes towards a complex, ongoing and fateful political reality, we came across Musolff's (2000) work on the British discourse on the European integration process throughout the 1990s. Working in the corpus linguistics tradition, Musolff (2000) studied a number of metaphors recurrent in this discourse, making available a selection of materials he used, marked with the metaphors.[8]

One caveat to directly using the database is the lack of clarity regarding the metaphor annotation procedure. In particular, the author does not report how many people participated, or any inter-annotator agreement figures. We therefore chose 4 out of Musolff's list of source domains, took all articles corresponding to them (128 documents), along with 23 articles from other source domains, and submitted them to a group of 8 undergraduate annotators, on top of Musolff's original markup that is treated as another annotator.

Annotators received the following instructions, reflecting our focus on the persuasive use of metaphor, as part of an argument:

> Generally speaking, a metaphor is a linguistic expression whereby something is compared to something else that it is clearly *literally* not, in order to make a point. Thus, in Tony Blair's famous "I haven't got a reverse gear", Tony Blair is compared to a car in order to stress his unwillingness/inability to retract his statements or actions. We would say in this case that a metaphor from a VEHICLE domain is used. In this study we will consider metaphors from 4 domains.

For the 4 chosen domains we provided the following descriptions, along with 2 examples for each:

---

[8]available from http://www.dur.ac.uk/andreas.musolff/Arcindex.htm

3

AUTHORITY Metaphors that have to do with discipline and authority, like school, religion, royalty, asylum, prison, etc.

LOVE Metaphors from love/romance and family.

BUILD Metaphors that have to do with building (the process) and houses and other buildings or constructions, their parts and uses.

VEHICLE Metaphors that have to do with land-borne vehicles, their parts, operation and maintenance.

People were instructed to mark every paragraph where a metaphor from a given domain occurs. They were also asked to provide a comment that briefly summarizes the ground for their decision, saying what is being compared to what.[9]

Table 2 shows the inter-annotator agreement figures.

Table 2: Inter-annotator agreement, measured on 2364 paragraphs (151 documents).[11]

Source Domain of Metaphor	κ
LOVE	0.66
VEHICLE	0.66
AUTHORITY	0.39
BUILD	0.43

LOVE and VEHICLE are close to acceptable reliability, with the other two types scoring low. In order to understand the nature of disagreements, we submitted the annotated materials plus some random annotations to 7 out of the original 8 people for validation, 4-8 weeks after they completed the annotations, asking them to accept or reject metaphor markups. We found that metaphors initially marked by at least 4 people (out of 9) were accepted as valid by people who did not initially mark them in 91% of the cases, on average across the metaphor types. These are thus uncontroversial cases, with the missing annotations likely due to attention slips rather than to genuine differences of opinion. Metaphors initially marked by 1-3 people were more controversial, with the average validation rate of 41% (Beigman Klebanov et al., 2008).

Evidently, some of the metaphors are clearer-cut than others, yet even the more difficult cases got non-negligible support at validation time from people who did not initially mark them. We therefore decided to regard the whole of the annotated data as valid for the purpose of the current research. Our focus is on finding metaphors (recall), and less on making sure all candidate metaphors are acceptable to all annotators; it suffices to know that even the minority opinion often finds support.

In the second stage of the research, we expanded the repertoire of the metaphor types to include additional source domains, mainly from Musolff's list. The dataset has so far been subjected to non-expert annotations by a group of the total of 15 undergraduate students. Metaphors from the source domains of VEHICLE, LOVE, BUILDING, AUTHORITY, WAR, SHOW, SCHOOL, RELIGION, MEDICINE were annotated by different subsets of the students.

The outcome of the second stage of the project is not sufficient for addressing the issue of discourse topics vs metaphors, however, as there are instances of metaphors in the text that do not fall into any of the source domains singled out by Musolff as recurrent ones in the discourse under consideration. We are now at an early stage of the third phrase we call OpenMeta, where annotators are asked to mark all metaphors they can detect, not confining themselves to a given list of source domains. Only annotators who participated in the previous, type-constrained, version of the task participate in OpenMeta project. So far, we have 44 documents annotated by 3 people for open-domain metaphors. This subset features as full a coverage of all metaphors used in the documents as we were able to obtain so far, and it is going to serve as test data for the topics vs metaphors hypothesis.

---

[9]In the topics vs metaphors experiment, we test the hypothesis on words rather than paragraphs. For metaphors from a pre-specified domain, such as VEHICLE or LOVE, it was usually clear which words in the paragraph belong to the domain and are used metaphorically. People's comments often explicitly used words from the paragraph, or made it otherwise clear through their description. For OpenMeta phase (please see below), where people were asked to mark metaphors from any source domain, they were also asked to single out the words in the paragraph that witness the metaphor, and these are the words used in the current experiment.

[11]These are results for binary classification for each metaphor type rather than a multiclass classification, since some articles have more than one type and some have none.

Our test set is thus biased towards recurrent metaphorical domains (those named by Musolff), and towards metaphors that are relatively salient to a naive reader, from recurrent or other source domains. Metaphors marked in the test data are those afforded a high degree of rhetorical presence in the discourse – either quantitatively, because they are repeated and elaborated, or qualitatively, because they are striking enough to arrest the naive reader's attention. According to the Presence Theory in rhetoric (Perelman and Olbrechts-Tyteca, 1969; Gross and Dearin, 2003; Atkinson et al., 2008), elements afforded high presence are key to the rhetorical design of the argument. These are not so much metaphors we live by without even noticing, such as those often studied in Conceptual Metaphor literature, like VALUE AS SIZE or TIME AS SPACE; these are metaphors that are clearly a matter of the author's conscious choice, closest in the current theorizing to Steen's (2008) notion of *deliberate* metaphors.

2.1 Pseudo sampling

The annotated data is not really a sample of the corpus. In fact, it is not known to us exactly how the documents were chosen; although all 44 metaphor annotated documents are from the newspapers and dates participating in the EUI corpus, only 20% are actually in the EUI corpus. How can we establish that there is a fit between the EUI collection and the annotated texts? We check how well discourse topics cover the documents, in the corpus and in the annotated material. Specifically, for a fixed k, is there a difference in the $\frac{|T_i^k|}{|D_i|}$ for annotated documents as opposed to the corpus at large? Using a random sample of 50 documents from EUI corpus, a 2-tailed t-test yielded $p < 0.05$, for all k, the trend being towards a better coverage of the EUI documents than of the metaphor annotated ones.

We hypothesized that this was due to the large discrepancy in the lengths of the texts: An average text in the EUI sample is 432 words long, whereas the metaphor annotated texts are 775 words long on average, with the shortest having 343 words. Shorter texts tend to be less elaborate and more "to the point", with a higher percentage of topical words.

To neutralize the effect of length on topical coverage, we chose from the EUI sample only documents that were at least 343 words long, resulting in 31 documents. Comparing those to the 44 metaphor annotated documents, we found $p > 0.37$ for every k, i.e. the annotated documents are indistinguishable in topical coverage from similar-length documents in the EUI corpus.

3 Experiment

3.1 Summary of notation

V All and only non-stop words containing only letters that appeared in at least 6 documents in the collection.

D$_i$ All words in document i that are in V.

T$_i^k$ All words in document i that are in V and are in the top k words for some topic active in document i according to LDA output.

M$_i$ All words in document i that are in V and are marked as metaphors in this document.

3.2 Hypothesis

We hypothesize that words in a given document that are high-ranking representatives of a common topic of discussion in a relevant corpus are less likely to be metaphorical than other words in the document. That is, such words would contain a smaller proportion of metaphors than their share in text. Using the definitions above: For an average document i and any k, $\frac{|T_i^k|}{|D_i|} > \frac{|M_i \cap T_i^k|}{|M_i|}$.

3.3 Results

As we hypothesized, metaphors are under-represented in topically used words. Thus, for k=25, about 15% of the indexed words in the document are deemed topical, containing about 3% of the metaphorically used indexed words in that document. For k=400, about 53% of the indexed words are topical, capturing only 22% of the metaphors.

4 Discussion

4.1 Metaphors from salient domains

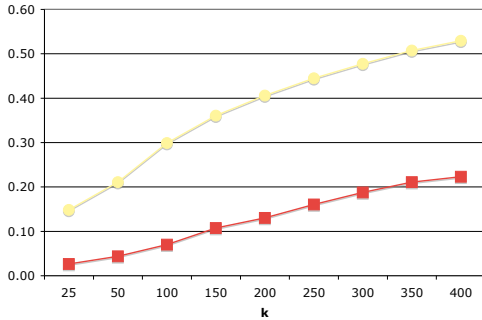A number of domains singled out by Musolff (2000) as being recurrent metaphors in the corpus, such

Figure 1: As hypothesized, $\frac{|T_i^k|}{|D_i|}$, shown in circles, is larger than $\frac{|M_i \cap T_i^k|}{|M_i|}$, shown in squares, for various k.

as VEHICLE or LOVE, are also things people care about politically, hence they also correspond to recurrent topics of discussion (see clusters titled *transport* and *childcare* in table 1). It has been shown experimentally that the subject's in-depth familiarity with the source domain is necessary for the metaphor to work as intended – see for example Gentner and Gentner (1983) work on using water flow metaphors for electricity. Our results suggest that participants in political discourse draw on domains not only familiar in general, but indeed highly salient in the specific discourse itself.

As a consequence, an extended metaphor from a discourse-topical domain can be easily mistaken by the topic detection software for a topical use of the relevant items. Consider, for example, an extract from a 19 December 1991 article in *Times*:

> Denis Healey, former Labour Chancellor of the Exchequer, urged the prime minister to stop playing Tory party politics with the negotiations over Europe and drew an image of Mr Major as a driver. He said: "I understand that if you are driving a car and sitting behind you is a lady with a handbag and a man with fangs, you may feel it wiser to drive in the slow lane. My own advice is that he should pull into a lay-by, turf the others out and then hand the wheel over to firmer and safer hands."

LDA considered {*drive driving*} to belong to the topic that deals with safety and road accidents, including in its 200 most likely words {*crash died accidents pedestrians traffic safety cars maps motorists*}, although additional metaphorically used items from the same semantic domain, such as *lane* and *wheel*, were not among the top 200 representatives of this topic.

It is an intriguing direction for future research to compare the topical and metaphorical uses of such domains, in order to determine which aspects loom large indeed, being both matters of literal concern and prolific generators of metaphors, and how these are manipulated for persuasive effects. The example above suggests that in the British EU-related discourse in 1990s safety of driving is both a topic-of-discussion ("Cyclists and pedestrians are more vulnerable on British roads than anywhere else in the European Union", proclaims *The Times* on 18 February 2000) and a metaphorical axis, stressing the importance of care and control, the hallmark of the Euro-sceptic stance towards the European integration process.

4.2 Topical metaphors

Putting aside topic detector's mistakes on extended metaphors from certain domains such as discussed in the previous section, what do metaphors in the topical vocabulary look like? The last topic shown in table 1 has to do with criticism towards EU bureaucracy, reflecting extensive discussions in the British media in the late 1990s of alleged corruption and mismanagement in the European Commission. Together with the words cited in the table, this topic lists *root* as one of its 300 most likely words.

This word shows up as a metaphor in 3 of our test documents. In two of them it is used precisely in the context projected by the topic:

> In limpid language, whose meaning no bureaucrat can twist, these four wise men and one wise woman delivered, to their great credit, a coruscating indictment not just of individual commissioners, but of the entire management and corporate culture of the European Commission. They have made an incontestable case, in Tony Blair's words, for "**root** and branch reform".

Here, *root* is used in the *root and branch* idiom suggesting a complete change, a reform, which comes as part of a bundle with severe criticism. Yet the figurative nature of this expression as a metaphor from PLANT domain is apparent to naive readers, making it an instance of imagery routinely going together with criticism in this corpus. A related metaphorical sense of *root* is attested in similar contexts in the corpus, further explaining its connection to the topic:

> Not unless they insist on credible systems to hold commissioners and bureaucrats to account. And not unless they appoint a new team with a brief not just to **root** out malpractices but to shut down entire programmes, such as tourism and humanitarian aid, which the Commission is incompetent to manage and which should never have been added to its ever-expanding empire.

> A bloodied European Commission looks likely to cling on to power today after an eleventh-hour threat to quit by its President, Jacques Santer, called the bluff of the European Parliament ... All week MEPs had been talking up the "nuclear option" of sacking the full Commission body over a burgeoning fraud and nepotism scandal that dates from 1995 ... Early 1997: Finnish Commissioner Erkki Liikanen announces plan to **root** out nepotism in Commission and improve financial controls.

In the third document with *root* metaphor, *root* is used in a different environment, and is not considered topical by LDA:

> For at the **root** of this conflict lies the German denial that unemployment has anything to do with cyclical fluctuations in the economy.

Our quantitative results show that cases such as *root* are more an exception than a rule. Yet, from the perspective of the argumentative use of metaphors, such cases are instructive of the way certain metaphors get "attached" to certain topics of discussion. In this case, the majority of mentions of *root* in this critical context come from Tony Blair's expression that was cited and referenced widely enough to acquire a statistical association with the discussion of the Commission's failings in the corpus. Indeed, the political significance of Blair's successful appropriation of the issue was not lost on the media:

> Tony Blair has swiftly positioned himself as the champion of "**root** and branch" reform. Not to be outdone, William Hague unveiled a "10-point plan" for reform of the Commission, no doubt drawing on his extensive McKinsey management expertise.

In future work, we plan to look closely at the topical metaphors, as they potentially represent outcomes of leadership battles fought in the media, and can thus have political consequences.

5 Conclusion

Using metaphor-annotated material that is sufficiently representative of the topical composition of a similar-length document in a large background corpus, we showed that words expressing a discourse-wide topic of discussion are less likely to be metaphorical than other words in a document.

This is, to our knowledge, the first quantitative demonstration of the connection between metaphoricity of a given word and its role in the relevant background discourse. It complements the traditionally "localistic" outlook on metaphors that is based on the observation that a metaphorically used item creates a local incongruity because there is a violation of a selectional restrictions between verbs and their arguments (Fass, 1991; Mason, 2004; Gedigian et al., 2006; Birke and Sarkar, 2006) or in the adjective-noun pairs (Krishnakumaran and Zhu, 2007). Global discourse-level information can potentially be used to focus metaphor detectors operating at the local level on items with higher metaphoric potential.

Reining and Lönneker-Rodman (2007) use minimal topical information to focus their search for metaphors. Working with a French-language

corpus discussing European politics, Reining and Lönneker-Rodman (2007) proposed harvesting salient collocates of the lemma *Europe*, that represents the main topic of discussion and is thus hypothesized to be the main target domain of metaphors in this corpus. Indeed, numerous instances of metaphors were collected using a 4-word window around the lemma in their corpus. Our work can be understood as developing a more nuanced approach to finding the likely target domains in the corpus – those words that represent a topic of discussion rather than the means to discuss a topic. Thus, it is not just Europe per se that is the target, but, more specifically, aspects such as monetary integration, employment, energy, immigration, transportation, and defense, among others. Our results suggest that to harvest deliberate metaphors more effectively, one is advised to consider words that do not represent a discourse topic.

References

Nathan Atkinson, David Kaufer, and Suguru Ishizaki. 2008. Presence and Global Presence in Genres of Self-Presentation: A Framework for Comparative Analysis. *Rhetoric Society Quarterly*, 38(3):1–27.

Beata Beigman Klebanov, Eyal Beigman, and Daniel Diermeier. 2008. Analyzing Disagreements. In *COLING 2008 Workshop on Human Judgments in Computational Linguistics*, pages 2–7, Manchester, UK.

Julia Birke and Anoop Sarkar. 2006. A clustering approach for nearly unsupervised recognition of nonliteral language. In *Proceedings of EACL*, pages 329–336.

David Blei, Andrew Ng, and Michael Jordan. 2003. Latent Dirichlet Allocation. *Journal of Machine Learning Resarch*, 3:993–1022.

Ari Chanen and Jon Patrick. 2007. Measuring correlation between linguists judgments and Latent Dirichlet Allocation topics. In *Proceedings of the Australasian Language Technology workshop*, pages 13–20, Melbourne, Australia.

Dan Fass. 1991. Met*: A method for discriminating metonymy and metaphor by computer. *Computational Linguistics*, 17(1):49–90.

Matt Gedigian, John Bryant, Srinivas Narayanan, and Branimir Ciric. 2006. Catching metaphors. In *Proceedings of NAACL Workshop on Scalable Natural Language Understanding*, pages 41–48.

Deidre Gentner and Donald Gentner. 1983. Flowing waters or teeming crowds: Mental models of electricity. In D. Gentner and A. Stevens, editors, *Mental models*. Hillsdale, NJ: Lawrence Erlbaum.

Alan Gross and Ray Dearin. 2003. *Chaim Perelman*. Albany: SUNY Press.

Eva Feder Kittay. 1987. *Metaphor: Its cognitive force and linguistic structure*. Oxford: Calderon Press.

Saisuresh Krishnakumaran and Xiaojin Zhu. 2007. Hunting elusive metaphors using lexical resources. In *Proceedings of the Workshop on Computational Approaches to Figurative Language*, pages 13–20, Rochester, New York.

Zachary J. Mason. 2004. CorMet: A computational, corpus-based conventional metaphor extraction system. *Computational Linguistics*, 30(1):23–44.

Andreas Musolff. 2000. *Mirror images of Europe: Metaphors in the public debate about Europe in Britain and Germany*. München: Iudicium. Annotated data is available at http://www.dur.ac.uk/andreas.musolff/Arcindex.htm.

Chaim Perelman and Lucie Olbrechts-Tyteca. 1969. *The New Rhetoric: A Treatise on Argumentation*. Wilkinson, J. and Weaver, P. (trans). Notre Dame, IN: University of Notre Dame Press.

Astrid Reining and Birte Lönneker-Rodman. 2007. Corpus-driven metaphor harvesting. In *Proceedings of the Workshop on Computational Approaches to Figurative Language*, pages 5–12, Rochester, New York.

Gerard Steen. 2008. The Paradox of Metaphor: Why We Need a Three-Dimensional Model of Metaphor. *Metaphor and Symbol*, 23(4):213–241.