

# Models of Metaphor in NLP

Ekaterina Shutova

Computer Laboratory  
University of Cambridge  
15 JJ Thomson Avenue  
Cambridge CB3 0FD, UK

Ekaterina.Shutova@cl.cam.ac.uk

## Abstract

Automatic processing of metaphor can be clearly divided into two subtasks: *metaphor recognition* (distinguishing between literal and metaphorical language in a text) and *metaphor interpretation* (identifying the intended literal meaning of a metaphorical expression). Both of them have been repeatedly addressed in NLP. This paper is the first comprehensive and systematic review of the existing computational models of metaphor, the issues of metaphor annotation in corpora and the available resources.

## 1 Introduction

Our production and comprehension of language is a multi-layered computational process. Humans carry out high-level semantic tasks effortlessly by subconsciously employing a vast inventory of complex linguistic devices, while simultaneously integrating their background knowledge, to reason about reality. An ideal model of language understanding would also be capable of performing such high-level semantic tasks.

However, a great deal of NLP research to date focuses on processing lower-level linguistic information, such as e.g. part-of-speech tagging, discovering syntactic structure of a sentence (parsing), coreference resolution, named entity recognition and many others. Another cohort of researchers set the goal of improving application-based statistical inference (e.g. for recognizing textual entailment or automatic summarization). In contrast, there have been fewer attempts to bring the state-of-the-art NLP technologies together to model the way humans use language to frame high-level reasoning processes, such as for example, creative thought.

The majority of computational approaches to

figurative language still exploit the ideas articulated three decades ago (Wilks, 1978; Lakoff and Johnson, 1980; Fass, 1991) and often rely on task-specific hand-coded knowledge. However, recent work on lexical semantics and lexical acquisition techniques opens many new avenues for creation of fully automated models for recognition and interpretation of figurative language. In this paper I will focus on the phenomenon of metaphor and describe the most prominent computational approaches to metaphor, as well the issues of resource creation and metaphor annotation.

Metaphors arise when one concept is viewed in terms of the properties of the other. In other words it is based on *similarity* between the concepts. Similarity is a kind of association implying the presence of characteristics in common. Here are some examples of metaphor.

- (1) Hillary *brushed aside* the accusations.
- (2) How can I *kill* a process? (Martin, 1988)
- (3) I *invested* myself fully in this relationship.
- (4) And then my heart with pleasure *fills*,  
And *dances* with the daffodils.<sup>1</sup>

In metaphorical expressions seemingly unrelated features of one concept are associated with another concept. In the example (2) the *computational process* is viewed as something *alive* and, therefore, its forced termination is associated with the act of killing.

Metaphorical expressions represent a great variety, ranging from conventional metaphors, which we reproduce and comprehend every day, e.g. those in (2) and (3), to poetic and largely novel ones, such as (4). The use of metaphor is ubiquitous in natural language text and it is a serious bottleneck in automatic text understanding.

<sup>1</sup>“I wandered lonely as a cloud”, William Wordsworth, 1804.

In order to estimate the frequency of the phenomenon, Shutova (2010) conducted a corpus study on a subset of the British National Corpus (BNC) (Burnard, 2007) representing various genres. They manually annotated metaphorical expressions in this data and found that 241 out of 761 sentences contained a metaphor. Due to such a high frequency of their use, a system capable of recognizing and interpreting metaphorical expressions in unrestricted text would become an invaluable component of any semantics-oriented NLP application.

Automatic processing of metaphor can be clearly divided into two subtasks: *metaphor recognition* (distinguishing between literal and metaphorical language in text) and *metaphor interpretation* (identifying the intended literal meaning of a metaphorical expression). Both of them have been repeatedly addressed in NLP.

## 2 Theoretical Background

Four different views on metaphor have been broadly discussed in linguistics and philosophy: the comparison view (Gentner, 1983), the interaction view (Black, 1962), (Hesse, 1966), the selectional restrictions violation view (Wilks, 1975; Wilks, 1978) and the conceptual metaphor view (Lakoff and Johnson, 1980)<sup>2</sup>. All of these approaches share the idea of an interconceptual mapping that underlies the production of metaphorical expressions. In other words, metaphor always involves two concepts or conceptual domains: the *target* (also called *topic* or *tenor* in the linguistics literature) and the *source* (or *vehicle*). Consider the examples in (5) and (6).

(5) He *shot down* all of my arguments. (Lakoff and Johnson, 1980)

(6) He *attacked* every weak point in my argument. (Lakoff and Johnson, 1980)

According to Lakoff and Johnson (1980), a mapping of a concept of *argument* to that of *war* is employed here. The *argument*, which is the target concept, is viewed in terms of a *battle* (or a *war*), the source concept. The existence of such a link allows us to talk about arguments using the war terminology, thus giving rise to a number of metaphors.

<sup>2</sup>A detailed overview and criticism of these four views can be found in (Tourangeau and Sternberg, 1982).

However, Lakoff and Johnson do not discuss how metaphors can be recognized in the linguistic data, which is the primary task in the automatic processing of metaphor. Although humans are highly capable of producing and comprehending metaphorical expressions, the task of distinguishing between literal and non-literal meanings and, therefore, identifying metaphor in text appears to be challenging. This is due to the variation in its use and external form, as well as a not clear-cut semantic distinction. Gibbs (1984) suggests that literal and figurative meanings are situated at the ends of a single continuum, along which metaphoricity and idiomaticity are spread. This makes demarcation of metaphorical and literal language fuzzy.

So far, the most influential account of metaphor recognition is that of Wilks (1978). According to Wilks, metaphors represent a violation of *selectional restrictions* in a given context. Selectional restrictions are the semantic constraints that a verb places onto its arguments. Consider the following example.

(7) My car *drinks* gasoline. (Wilks, 1978)

The verb *drink* normally takes an *animate* subject and a *liquid* object. Therefore, *drink* taking a *car* as a subject is an anomaly, which may in turn indicate the metaphorical use of *drink*.

## 3 Automatic Metaphor Recognition

One of the first attempts to identify and interpret metaphorical expressions in text automatically is the approach of Fass (1991). It originates in the work of Wilks (1978) and utilizes hand-coded knowledge. Fass (1991) developed a system called *met\**, capable of discriminating between literalness, metonymy, metaphor and anomaly. It does this in three stages. First, literalness is distinguished from non-literalness using selectional preference violation as an indicator. In the case that non-literalness is detected, the respective phrase is tested for being a metonymic relation using hand-coded patterns (such as CONTAINER-for-CONTENT). If the system fails to recognize metonymy, it proceeds to search the knowledge base for a *relevant analogy* in order to discriminate metaphorical relations from anomalous ones. E.g., the sentence in (7) would be represented in this framework as (*car,drink,gasoline*), which does not satisfy the preference (*animal,drink,liquid*), as *car*

is not a hyponym of *animal*. *met\** then searches its knowledge base for a triple containing a hypernym of both the actual argument and the desired argument and finds (*thing,use,energy\_source*), which represents the metaphorical interpretation.

However, Fass himself indicated a problem with the selectional preference violation approach applied to metaphor recognition. The approach detects any kind of non-literality or anomaly in language (metaphors, metonymies and others), and not only metaphors, i.e., it overgenerates. The methods *met\** uses to differentiate between those are mainly based on hand-coded knowledge, which implies a number of limitations.

Another problem with this approach arises from the high conventionality of metaphor in language. This means that some metaphorical senses are very common. As a result the system would extract selectional preference distributions skewed towards such conventional metaphorical senses of the verb or one of its arguments. Therefore, although some expressions may be fully metaphorical in nature, no selectional preference violation can be detected in their use. Another counterargument is bound to the fact that interpretation is always context dependent, e.g. the phrase *all men are animals* can be used metaphorically, however, without any violation of selectional restrictions.

Goatly (1997) addresses the phenomenon of metaphor by identifying a set of linguistic cues indicating it. He gives examples of lexical patterns indicating the presence of a metaphorical expression, such as *metaphorically speaking, utterly, completely, so to speak* and, surprisingly, *literally*. Such cues would probably not be enough for metaphor extraction on their own, but could contribute to a more complex system.

The work of Peters and Peters (2000) concentrates on detecting figurative language in lexical resources. They mine WordNet (Fellbaum, 1998) for the examples of systematic polysemy, which allows to capture metonymic and metaphorical relations. The authors search for nodes that are relatively high up in the WordNet hierarchy and that share a set of common word forms among their descendants. Peters and Peters found that such nodes often happen to be in metonymic (e.g. *publication – publisher*) or metaphorical (e.g. *supporting structure – theory*) relation.

The CorMet system discussed in (Mason, 2004) is the first attempt to discover source-target do-

main mappings automatically. This is done by “finding systematic variations in domain-specific selectional preferences, which are inferred from large, dynamically mined Internet corpora”. For example, Mason collects texts from the LAB domain and the FINANCE domain, in both of which *pour* would be a characteristic verb. In the LAB domain *pour* has a strong selectional preference for objects of type *liquid*, whereas in the FINANCE domain it selects for *money*. From this Mason’s system infers the domain mapping FINANCE – LAB and the concept mapping *money – liquid*. He compares the output of his system against the Master Metaphor List (Lakoff et al., 1991) containing hand-crafted metaphorical mappings between concepts. Mason reports an accuracy of 77%, although it should be noted that as any evaluation that is done by hand it contains an element of subjectivity.

Birke and Sarkar (2006) present a sentence clustering approach for non-literal language recognition implemented in the TroFi system (Trope Finder). This idea originates from a similarity-based word sense disambiguation method developed by Karov and Edelman (1998). The method employs a set of seed sentences, where the senses are annotated; computes similarity between the sentence containing the word to be disambiguated and all of the seed sentences and selects the sense corresponding to the annotation in the most similar seed sentences. Birke and Sarkar (2006) adapt this algorithm to perform a two-way classification: literal vs. non-literal, and they do not clearly define the kinds of tropes they aim to discover. They attain a performance of 53.8% in terms of f-score.

The method of Gedigan et al. (2006) discriminates between literal and metaphorical use. They trained a maximum entropy classifier for this purpose. They obtained their data by extracting the lexical items whose frames are related to MOTION and CURE from FrameNet (Fillmore et al., 2003). Then they searched the PropBank Wall Street Journal corpus (Kingsbury and Palmer, 2002) for sentences containing such lexical items and annotated them with respect to metaphoricity. They used PropBank annotation (arguments and their semantic types) as features to train the classifier and report an accuracy of 95.12%. This result is, however, only a little higher than the performance of the naive baseline assigning majority class to all instances (92.90%). These numbers

can be explained by the fact that 92.00% of the verbs of MOTION and CURE in the Wall Street Journal corpus are used metaphorically, thus making the dataset unbalanced with respect to the target categories and the task notably easier.

Both Birke and Sarkar (2006) and Gedigan et al. (2006) focus only on metaphors expressed by a verb. As opposed to that the approach of Krishnakumaran and Zhu (2007) deals with verbs, nouns and adjectives as parts of speech. They use hyponymy relation in WordNet and word bigram counts to predict metaphors at a sentence level. Given an IS-A metaphor (e.g. *The world is a stage*<sup>3</sup>) they verify if the two nouns involved are in hyponymy relation in WordNet, and if they are not then this sentence is tagged as containing a metaphor. Along with this they consider expressions containing a verb or an adjective used metaphorically (e.g. *He planted good ideas in their minds* or *He has a fertile imagination*). Hereby they calculate bigram probabilities of verb-noun and adjective-noun pairs (including the hyponyms/hypernyms of the noun in question). If the combination is not observed in the data with sufficient frequency, the system tags the sentence containing it as metaphorical. This idea is a modification of the selectional preference view of Wilks. However, by using bigram counts over verb-noun pairs Krishnakumaran and Zhu (2007) lose a great deal of information compared to a system extracting verb-object relations from parsed text. The authors evaluated their system on a set of example sentences compiled from the Master Metaphor List (Lakoff et al., 1991), whereby highly conventionalized metaphors (they call them *dead metaphors*) are taken to be negative examples. Thus they do not deal with literal examples as such: essentially, the distinction they are making is between the senses included in WordNet, even if they are conventional metaphors, and those not included in WordNet.

#### 4 Automatic Metaphor Interpretation

Almost simultaneously with the work of Fass (1991), Martin (1990) presents a Metaphor Interpretation, Denotation and Acquisition System (MIDAS). In this work Martin captures hierarchical organisation of conventional metaphors. The idea behind this is that the more specific conventional metaphors descend from the general ones.

<sup>3</sup>William Shakespeare

Given an example of a metaphorical expression, MIDAS searches its database for a corresponding metaphor that would explain the anomaly. If it does not find any, it abstracts from the example to more general concepts and repeats the search. If it finds a suitable general metaphor, it creates a mapping for its descendant, a more specific metaphor, based on this example. This is also how novel metaphors are acquired. MIDAS has been integrated with the Unix Consultant (UC), the system that answers users questions about Unix. The UC first tries to find a literal answer to the question. If it is not able to, it calls MIDAS which detects metaphorical expressions via selectional preference violation and searches its database for a metaphor explaining the anomaly in the question.

Another cohort of approaches relies on performing inferences about entities and events in the source and target domains for metaphor interpretation. These include the KARMA system (Narayanan, 1997; Narayanan, 1999; Feldman and Narayanan, 2004) and the ATT-Meta project (Barnden and Lee, 2002; Agerri et al., 2007). Within both systems the authors developed a metaphor-based reasoning framework in accordance with the theory of conceptual metaphor. The reasoning process relies on manually coded knowledge about the world and operates mainly in the source domain. The results are then projected onto the target domain using the conceptual mapping representation. The ATT-Meta project concerns metaphorical and metonymic description of mental states and reasoning about mental states using first order logic. Their system, however, does not take natural language sentences as input, but logical expressions that are representations of small discourse fragments. KARMA in turn deals with a broad range of abstract actions and events and takes parsed text as input.

Veale and Hao (2008) derive a “fluid knowledge representation for metaphor interpretation and generation”, called Talking Points. Talking Points are a set of characteristics of concepts belonging to source and target domains and related facts about the world which the authors acquire automatically from WordNet and from the web. Talking Points are then organized in *Slipnet*, a framework that allows for a number of insertions, deletions and substitutions in definitions of such characteristics in order to establish a connection between the target and the source

concepts. This work builds on the idea of *slippage* in knowledge representation for understanding analogies in abstract domains (Hofstadter and Mitchell, 1994; Hofstadter, 1995). Below is an example demonstrating how slippage operates to explain the metaphor *Make-up is a Western burqa*.

**Make-up =>**

- ≡ typically worn by women
- ≈ expected to be worn by women
- ≈ must be worn by women
- ≈ must be worn by Muslim women

**Burqa <=**

By doing insertions and substitutions the system arrives from the definition *typically worn by women* to that of *must be worn by Muslim women*, and thus establishes a link between the concepts of *make-up* and *burqa*. Veale and Hao (2008), however, did not evaluate to which extent their knowledge base of Talking Points and the associated reasoning framework are useful to interpret metaphorical expressions occurring in text.

Shutova (2010) defines metaphor interpretation as a paraphrasing task and presents a method for deriving literal paraphrases for metaphorical expressions from the BNC. For example, for the metaphors in “All of this *stirred* an unfathomable excitement in her” or “a carelessly *leaked* report” their system produces interpretations “All of this provoked an unfathomable excitement in her” and “a carelessly disclosed report” respectively. They first apply a probabilistic model to rank all possible paraphrases for the metaphorical expression given the context; and then use automatically induced selectional preferences to discriminate between figurative and literal paraphrases. The selectional preference distribution is defined in terms of selectional association measure introduced by Resnik (1993) over the noun classes automatically produced by Sun and Korhonen (2009). Shutova (2010) tested their system only on metaphors expressed by a verb and report a paraphrasing accuracy of 0.81.

## 5 Metaphor Resources

Metaphor is a knowledge-hungry phenomenon. Hence there is a need for either an extensive manually-created knowledge-base or a robust knowledge acquisition system for interpretation of metaphorical expressions. The latter being a hard task, a great deal of metaphor research resorted to

the first option. Although hand-coded knowledge proved useful for metaphor interpretation (Fass, 1991; Martin, 1990), it should be noted that the systems utilizing it have a very limited coverage.

One of the first attempts to create a multi-purpose knowledge base of source–target domain mappings is the Master Metaphor List (Lakoff et al., 1991). It includes a classification of metaphorical mappings (mainly those related to mind, feelings and emotions) with the corresponding examples of language use. This resource has been criticized for the lack of clear structuring principles of the mapping ontology (Lönneker-Rodman, 2008). The taxonomical levels are often confused, and the same classes are referred to by different class labels. This fact and the chosen data representation in the Master Metaphor List make it not suitable for computational use. However, both the idea of the list and its actual mappings ontology inspired the creation of other metaphor resources.

The most prominent of them are MetaBank (Martin, 1994) and the Mental Metaphor Databank<sup>4</sup> created in the framework of the ATT-meta project (Barnden and Lee, 2002; Agerri et al., 2007). The MetaBank is a knowledge-base of English metaphorical conventions, represented in the form of *metaphor maps* (Martin, 1988) containing detailed information about source-target concept mappings backed by empirical evidence. The ATT-meta project databank contains a large number of examples of metaphors of mind classified by source–target domain mappings taken from the Master Metaphor List.

Along with this it is worth mentioning metaphor resources in languages other than English. There has been a wealth of research on metaphor in Spanish, Chinese, Russian, German, French and Italian. The Hamburg Metaphor Database (Lönneker, 2004; Reining and Lönneker-Rodman, 2007) contains examples of metaphorical expressions in German and French, which are mapped to senses from EuroWordNet<sup>5</sup> and annotated with source–target domain mappings taken from the Master Metaphor List.

Alonge and Castelli (2003) discuss how metaphors can be represented in ItalWordNet for

<sup>4</sup><http://www.cs.bham.ac.uk/~jab/ATT-Meta/Databank/>

<sup>5</sup>EuroWordNet is a multilingual database with wordnets for several European languages (Dutch, Italian, Spanish, German, French, Czech and Estonian). The wordnets are structured in the same way as the Princeton WordNet for English. URL: <http://www.ilic.uva.nl/EuroWordNet/>

Italian and motivate this by linguistic evidence. Encoding metaphorical information in general-domain lexical resources for English, e.g. WordNet (Lönneker and Eilts, 2004), would undoubtedly provide a new platform for experiments and enable researchers to directly compare their results.

## 6 Metaphor Annotation in Corpora

To reflect two distinct aspects of the phenomenon, metaphor annotation can be split into two stages: identifying metaphorical senses in text (akin word sense disambiguation) and annotating source – target domain mappings underlying the production of metaphorical expressions. Traditional approaches to metaphor annotation include manual search for lexical items used metaphorically (Pragglejaz Group, 2007), for source and target domain vocabulary (Deignan, 2006; Koivisto-Alanko and Tisari, 2006; Martin, 2006) or for linguistic markers of metaphor (Goatly, 1997). Although there is a consensus in the research community that the phenomenon of metaphor is not restricted to similarity-based extensions of meanings of isolated words, but rather involves reconceptualization of a whole area of experience in terms of another, there still has been surprisingly little interest in annotation of cross-domain mappings. However, a corpus annotated for conceptual mappings could provide a new starting point for both linguistic and cognitive experiments.

### 6.1 Metaphor and Polysemy

The theorists of metaphor distinguish between two kinds of metaphorical language: *novel* (or *poetic*) metaphors, that surprise our imagination, and *conventionalized* metaphors, that become a part of an ordinary discourse. “Metaphors begin their lives as novel poetic creations with marked rhetorical effects, whose comprehension requires a special imaginative leap. As time goes by, they become a part of general usage, their comprehension becomes more automatic, and their rhetorical effect is dulled” (Nunberg, 1987). Following Orwell (1946) Nunberg calls such metaphors “dead” and claims that they are not psychologically distinct from literally-used terms.

This scheme demonstrates how metaphorical associations capture some generalisations governing polysemy: over time some of the aspects of the target domain are added to the meaning of a

term in a source domain, resulting in a (metaphorical) sense extension of this term. Copestake and Briscoe (1995) discuss sense extension mainly based on metonymic examples and model the phenomenon using lexical rules encoding metonymic patterns. Along with this they suggest that similar mechanisms can be used to account for metaphoric processes, and the conceptual mappings encoded in the sense extension rules would define the limits to the possible shifts in meaning.

However, it is often unclear if a metaphorical instance is a case of broadening of the sense in context due to general vagueness in language, or it manifests a formation of a new distinct metaphorical sense. Consider the following examples.

- (8) a. As soon as I *entered* the room I noticed the difference.
- b. How can I *enter* Emacs?
- (9) a. My tea is *cold*.
- b. He is such a *cold* person.

*Enter* in (8a) is defined as “to go or come into a place, building, room, etc.; to pass within the boundaries of a country, region, portion of space, medium, etc.”<sup>6</sup> In (8b) this sense stretches to describe dealing with *software*, whereby COMPUTER PROGRAMS are viewed as PHYSICAL SPACES. However, this extended sense of *enter* does not appear to be sufficiently distinct or conventional to be included into the dictionary, although this could happen over time.

The sentence (9a) exemplifies the basic sense of *cold* – “of a temperature sensibly lower than that of the living human body”, whereas *cold* in (9b) should be interpreted metaphorically as “void of ardour, warmth, or intensity of feeling; lacking enthusiasm, heartiness, or zeal; indifferent, apathetic”. These two senses are clearly linked via the metaphoric mapping between EMOTIONAL STATES and TEMPERATURES.

A number of metaphorical senses are included in WordNet, however without any accompanying semantic annotation.

## 6.2 Metaphor Identification

### 6.2.1 Pragglejaz Procedure

Pragglejaz Group (2007) proposes a metaphor identification procedure (MIP) within the frame-

<sup>6</sup>Sense definitions are taken from the Oxford English Dictionary.

work of the Metaphor in Discourse project (Steen, 2007). The procedure involves metaphor annotation at the word level as opposed to identifying metaphorical relations (between words) or source–target domain mappings (between concepts or domains). In order to discriminate between the verbs used metaphorically and literally the annotators are asked to follow the guidelines:

1. For each verb establish its meaning in context and try to imagine a more basic meaning of this verb on other contexts. Basic meanings normally are: (1) more concrete; (2) related to bodily action; (3) more precise (as opposed to vague); (4) historically older.
2. If you can establish the basic meaning that is distinct from the meaning of the verb in this context, the verb is likely to be used metaphorically.

Such annotation can be viewed as a form of word sense disambiguation with an emphasis on metaphoricity.

### 6.2.2 Source – Target Domain Vocabulary

Another popular method that has been used to extract metaphors is searching for sentences containing lexical items from the source domain, the target domain, or both (Stefanowitsch, 2006). This method requires exhaustive lists of source and target domain vocabulary.

Martin (2006) conducted a corpus study in order to confirm that metaphorical expressions occur in text in contexts containing such lexical items. He performed his analysis on the data from the Wall Street Journal (WSJ) corpus and focused on four conceptual metaphors that occur with considerable regularity in the corpus. These include NUMERICAL VALUE AS LOCATION, COMMERCIAL ACTIVITY AS CONTAINER, COMMERCIAL ACTIVITY AS PATH FOLLOWING and COMMERCIAL ACTIVITY AS WAR. Martin manually compiled the lists of terms characteristic for each domain by examining sampled metaphors of these types and then augmented them through the use of thesaurus. He then searched the WSJ for sentences containing vocabulary from these lists and checked whether they contain metaphors of the above types. The goal of this study was to evaluate predictive ability of contexts containing vocabulary from (1) source domain and (2) target

domain, as well as (3) estimating the likelihood of a metaphorical expression following another metaphorical expression described by the same mapping. He obtained the most positive results for metaphors of the type NUMERICAL-VALUE-AS-LOCATION ( $P(\textit{Metaphor}|\textit{Source}) = 0.069$ ,  $P(\textit{Metaphor}|\textit{Target}) = 0.677$ ,  $P(\textit{Metaphor}|\textit{Metaphor}) = 0.703$ ).

### 6.3 Annotating Source and Target Domains

Wallington et al. (2003) carried out a metaphor annotation experiment in the framework of the ATT-Meta project. They employed two teams of annotators. Team A was asked to annotate “interesting stretches”, whereby a phrase was considered interesting if (1) its significance in the document was non-physical, (2) it could have a physical significance in another context with a similar syntactic frame, (3) this physical significance was related to the abstract one. Team B had to annotate phrases according to their own intuitive definition of metaphor. Besides metaphorical expressions Wallington et al. (2003) attempted to annotate the involved source – target domain mappings. The annotators were given a set of mappings from the Master Metaphor List and were asked to assign the most suitable ones to the examples. However, the authors do not report the level of interannotator agreement nor the coverage of the mappings in the Master Metaphor List on their data.

Shutova and Teufel (2010) adopt a different approach to the annotation of source – target domain mappings. They do not rely on predefined mappings, but instead derive independent sets of most common source and target categories. They propose a two stage procedure, whereby the metaphorical expressions are first identified using MIP, and then the source domain (where the basic sense comes from) and the target domain (the given context) are selected from the lists of categories. Shutova and Teufel (2010) report interannotator agreement of 0.61 ( $\kappa$ ).

## 7 Conclusion and Future Directions

The eighties and nineties provided us with a wealth of ideas on the structure and mechanisms of the phenomenon of metaphor. The approaches formulated back then are still highly influential, although their use of hand-coded knowledge is becoming increasingly less convincing. The last decade witnessed a high technological leap in

natural language computation, whereby manually crafted rules gradually give way to more robust corpus-based statistical methods. This is also the case for metaphor research. The latest developments in the lexical acquisition technology will in the near future enable fully automated corpus-based processing of metaphor.

However, there is still a clear need in a unified metaphor annotation procedure and creation of a large publicly available metaphor corpus. Given such a resource the computational work on metaphor is likely to proceed along the following lines: (1) automatic acquisition of an extensive set of valid metaphorical associations from linguistic data via statistical pattern matching; (2) using the knowledge of these associations for metaphor recognition in the unseen unrestricted text and, finally, (3) interpretation of the identified metaphorical expressions by deriving the closest literal paraphrase (a representation that can be directly embedded in other NLP applications to enhance their performance).

Besides making our thoughts more vivid and filling our communication with richer imagery, metaphors also play an important structural role in our cognition. Thus, one of the long term goals of metaphor research in NLP and AI would be to build a computational intelligence model accounting for the way metaphors organize our conceptual system, in terms of which we think and act.

## Acknowledgments

I would like to thank Anna Korhonen and my reviewers for their most helpful feedback on this paper. The support of Cambridge Overseas Trust, who fully funds my studies, is gratefully acknowledged.

## References

R. Agerri, J.A. Barnden, M.G. Lee, and A.M. Wallington. 2007. Metaphor, inference and domain-independent mappings. In *Proceedings of RANLP-2007*, pages 17–23, Borovets, Bulgaria.

A. Alonge and M. Castelli. 2003. Encoding information on metaphoric expressions in WordNet-like resources. In *Proceedings of the ACL 2003 Workshop on Lexicon and Figurative Language*, pages 10–17.

J.A. Barnden and M.G. Lee. 2002. An artificial intelligence approach to metaphor understanding. *Theoria et Historia Scientiarum*, 6(1):399–412.

J. Birke and A. Sarkar. 2006. A clustering approach for the nearly unsupervised recognition of nonliteral language. In *In Proceedings of EACL-06*, pages 329–336.

M. Black. 1962. *Models and Metaphors*. Cornell University Press.

L. Burnard. 2007. *Reference Guide for the British National Corpus (XML Edition)*.

A. Copestake and T. Briscoe. 1995. Semi-productive polysemy and sense extension. *Journal of Semantics*, 12:15–67.

A. Deignan. 2006. The grammar of linguistic metaphors. In A. Stefanowitsch and S. T. Gries, editors, *Corpus-Based Approaches to Metaphor and Metonymy*, Berlin. Mouton de Gruyter.

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

J. Feldman and S. Narayanan. 2004. Embodied meaning in a neural theory of language. *Brain and Language*, 89(2):385–392.

C. Fellbaum, editor. 1998. *WordNet: An Electronic Lexical Database (ISBN: 0-262-06197-X)*. MIT Press, first edition.

C. J. Fillmore, C. R. Johnson, and M. R. L. Petruck. 2003. Background to FrameNet. *International Journal of Lexicography*, 16(3):235–250.

M. Gedigan, J. Bryant, S. Narayanan, and B. Ciric. 2006. Catching metaphors. In *In Proceedings of the 3rd Workshop on Scalable Natural Language Understanding*, pages 41–48, New York.

D. Gentner. 1983. Structure mapping: A theoretical framework for analogy. *Cognitive Science*, 7:155–170.

R. Gibbs. 1984. Literal meaning and psychological theory. *Cognitive Science*, 8:275–304.

A. Goatly. 1997. *The Language of Metaphors*. Routledge, London.

M. Hesse. 1966. *Models and Analogies in Science*. Notre Dame University Press.

D. Hofstadter and M. Mitchell. 1994. The Copycat Project: A model of mental fluidity and analogy-making. In K.J. Holyoak and J. A. Barnden, editors, *Advances in Connectionist and Neural Computation Theory*, Ablex, New Jersey.

D. Hofstadter. 1995. *Fluid Concepts and Creative Analogies: Computer Models of the Fundamental Mechanisms of Thought*. HarperCollins Publishers.

Y. Karov and S. Edelman. 1998. Similarity-based word sense disambiguation. *Computational Linguistics*, 24(1):41–59.



- P. Kingsbury and M. Palmer. 2002. From TreeBank to PropBank. In *Proceedings of LREC-2002*, Gran Canaria, Canary Islands, Spain.
- P. Koivisto-Alanko and H. Tissari. 2006. Sense and sensibility: Rational thought versus emotion in metaphorical language. In A. Stefanowitsch and S. T. Gries, editors, *Corpus-Based Approaches to Metaphor and Metonymy*, Berlin. Mouton de Gruyter.
- S. Krishnakumaran and X. Zhu. 2007. Hunting elusive metaphors using lexical resources. In *Proceedings of the Workshop on Computational Approaches to Figurative Language*, pages 13–20, Rochester, NY.
- G. Lakoff and M. Johnson. 1980. *Metaphors We Live By*. University of Chicago Press, Chicago.
- G. Lakoff, J. Espenson, and A. Schwartz. 1991. The master metaphor list. Technical report, University of California at Berkeley.
- B. Lönneker and C. Eilts. 2004. A Current Resource and Future Perspectives for Enriching Word-Nets with Metaphor Information. In *Proceedings of the Second International WordNet Conference—GWC 2004*, pages 157–162, Brno, Czech Republic.
- B. Lönneker-Rodman. 2008. The hamburg metaphor database project: issues in resource creation. *Language Resources and Evaluation*, 42(3):293–318.
- B. Lönneker. 2004. Lexical databases as resources for linguistic creativity: Focus on metaphor. In *Proceedings of the LREC 2004 Workshop on Language Resources for Linguistic Creativity*, pages 9–16, Lisbon, Portugal.
- J. H. Martin. 1988. Representing regularities in the metaphoric lexicon. In *Proceedings of the 12th conference on Computational linguistics*, pages 396–401.
- J. H. Martin. 1990. *A Computational Model of Metaphor Interpretation*. Academic Press Professional, Inc., San Diego, CA, USA.
- J. H. Martin. 1994. Metabank: A knowledge-base of metaphoric language conventions. *Computational Intelligence*, 10:134–149.
- J. H. Martin. 2006. A corpus-based analysis of context effects on metaphor comprehension. In A. Stefanowitsch and S. T. Gries, editors, *Corpus-Based Approaches to Metaphor and Metonymy*, Berlin. Mouton de Gruyter.
- Z. J. Mason. 2004. Cormet: a computational, corpus-based conventional metaphor extraction system. *Computational Linguistics*, 30(1):23–44.
- S. Narayanan. 1997. Knowledge-based action representations for metaphor and aspect (karma). Technical report, PhD thesis, University of California at Berkeley.
- S. Narayanan. 1999. Moving right along: A computational model of metaphoric reasoning about events. In *Proceedings of AAAI 99*, pages 121–128, Orlando, Florida.
- G. Nunberg. 1987. Poetic and prosaic metaphors. In *Proceedings of the 1987 workshop on Theoretical issues in natural language processing*, pages 198–201.
- G. Orwell. 1946. Politics and the english language. *Horizon*.
- W. Peters and I. Peters. 2000. Lexicalised systematic polysemy in wordnet. In *Proceedings of LREC 2000*, Athens.
- Pragglejaz Group. 2007. MIP: A method for identifying metaphorically used words in discourse. *Metaphor and Symbol*, 22:1–39.
- A. Reining and B. Lönneker-Rodman. 2007. Corpus-driven metaphor harvesting. In *Proceedings of the HLT/NAACL-07 Workshop on Computational Approaches to Figurative Language*, pages 5–12, Rochester, New York.
- P. Resnik. 1993. *Selection and Information: A Class-based Approach to Lexical Relationships*. Ph.D. thesis, Philadelphia, PA, USA.
- E. Shutova and S. Teufel. 2010. Metaphor corpus annotated for source - target domain mappings. In *Proceedings of LREC 2010*, Malta.
- E. Shutova. 2010. Automatic metaphor interpretation as a paraphrasing task. In *Proceedings of NAACL 2010*, Los Angeles, USA.
- G. J. Steen. 2007. Finding metaphor in discourse: Pragglejaz and beyond. *Cultura, Lenguaje y Representacion / Culture, Language and Representation (CLR)*, *Revista de Estudios Culturales de la Universitat Jaume I*, 5:9–26.
- A. Stefanowitsch. 2006. Corpus-based approaches to metaphor and metonymy. In A. Stefanowitsch and S. T. Gries, editors, *Corpus-Based Approaches to Metaphor and Metonymy*, Berlin. Mouton de Gruyter.
- L. Sun and A. Korhonen. 2009. Improving verb clustering with automatically acquired selectional preferences. In *Proceedings of EMNLP 2009*, pages 638–647, Singapore, August.
- R. Tourangeau and R. Sternberg. 1982. Understanding and appreciating metaphors. *Cognition*, 11:203–244.
- T. Veale and Y. Hao. 2008. A fluid knowledge representation for understanding and generating creative metaphors. In *Proceedings of COLING 2008*, pages 945–952, Manchester, UK.

- A. M. Wallington, J. A. Barnden, P. Buchlovsky, L. Fellows, and S. R. Glasbey. 2003. Metaphor annotation: A systematic study. Technical report, School of Computer Science, The University of Birmingham.
- Y. Wilks. 1975. A preferential pattern-seeking semantics for natural language inference. *Artificial Intelligence*, 6:53–74.
- Y. Wilks. 1978. Making preferences more active. *Artificial Intelligence*, 11(3):197–223.