

# Computational Approaches to Figurative Language

Birte Loenneker-Rodman and Srinu Narayanan

October 1, 2008

## 1 Introduction

The heading *figurative language* subsumes multiple phenomena that can be used to perform most linguistic functions including predication, modification, and reference. Figurative language can tap into conceptual and linguistic knowledge (as in the case of idioms, metaphor, and some metonymies) as well as evoke pragmatic factors in interpretation (as in indirect speech acts, humor, irony, or sarcasm). Indeed, the distinction between the terms *literal* and *figurative* is far from clear cut. While there is a continuum from literal to figurative, certain phenomena such as metaphor and metonymy are considered prototypical instances of figurative language.

To date, no comprehensive computational system addressing all of figurative language has been implemented, or even designed. Most computational work has focused on the conceptual and linguistic underpinnings of figurative language. Our discussion will focus on the two types of figurative language phenomena that have been studied computationally: metonymy (Section 2) and metaphor (Section 3). Each section introduces the phenomenon by way of examples and provides an overview of computational approaches. Section 4 concludes with a brief discussion of computational attempts to model other types of figurative language including humor and irony.

## 2 Metonymy

Metonymy is by far the most studied type of figurative language from a computational point of view. The discussion to follow situates the computational work on metonymy along several dimensions. Section 2.1 presents the phenomenon from the point of view of linguistics, introducing some of the problems computational approaches have to face. Corpus annotation gives some insights into subtypes and frequency of metonymy (Section 2.2); annotated data can also be further exploited by automatic systems, although some systems do not need such training data. Examples of computational methods and systems are discussed in Section 2.3.

### 2.1 Metonymy in Linguistics

In metonymy, one entity stands for another entity, and this is reflected in language. In the examples below, metonymical expressions are emphasized in variants (a.); variants (b.) are literal paraphrases spelling out what the metonymically used expression stands for.

- (1) a. I am reading *Shakespeare*.
- b. I am reading *one of Shakespeare's works*.

- (2) a. America doesn't want another *Pearl Harbor*.  
b. America doesn't want another *defeat in war*.
- (3) a. *Washington* is negotiating with *Moscow*.  
b. *The American government* is negotiating with *the Russian government*.
- (4) a. We need a better *glove* at third base.  
b. We need a better *baseball player* at third base.

According to commonly agreed upon views in Cognitive Linguistics (e.g. Kövecses, 2002, pp. 143–144), the metonymically used word refers to an entity that *provides mental access* to another entity (the one that is referred to in the literal paraphrase). The access-providing entity has been called *vehicle*; the one to which mental access is provided is the *target entity* (cf. Kövecses, 2002).

Entities participating in metonymy are members of the same conceptual domain, where they are *related* to one another. For example, both an author (such as *Shakespeare*) and written works belong to the ARTISTIC PRODUCTION domain. The coherence of a domain is brought about by the repeated co-occurrence of its component entities in the world, where they are experienced as being (closely) “together” (Kövecses, 2002, pp. 145–146); thus, the relation underlying metonymy has been called *contiguity*.

Although belonging to the same conceptual domain, metonymical vehicle and target are usually different types of entities; they are *ontologically different*. For example, in (1), above, a HUMAN stands for a PHYSICAL OBJECT. In (5), below, a PHYSICAL OBJECT stands for a PROCESS.

- (5) a. Mary finished *the cigarette*.  
b. Mary finished *smoking the cigarette*.

Metonymy is common in spoken and written communication. In fact, between some entities, the contiguity relation is so *salient* (Hobbs, 2001) or *noteworthy* (Nunberg, 1995) that it is more natural to use metonymy than a literal paraphrase. This has also been confirmed by corpus studies; we will come back to frequency issues in Section 2.2.

### 2.1.1 Regular metonymy

While there are obviously many relations between entities that are rarely or never exploited for metonymy, other relations between some types of entities give rise to larger clusters of metonymic expressions, or *regular metonymy*. Thanks to these relations, many members of an entity class can stand for members of a different class.

Examples can be found in (Lakoff and Johnson, 1980, pp. 38-39), including some of the following.

- (6) THE PLACE FOR THE EVENT
  - a. Let's not let *Thailand* become another *Vietnam*.
  - b. *Watergate* changed our politics.
- (7) OBJECT USED FOR USER
  - a. The *sax* has the flu today.
  - b. The *gun* he hired wanted fifty grand.
  - c. The *buses* are on strike.

d. [flight attendant, on a plane]: Ask *seat 19* whether he wants to swap. (Nissim and Markert, 2003, p. 56)

(8) THE TREE FOR THE WOOD (Nunberg, 1995)

a. The table is made of *oak*.

The regularity of some of these metonymic patterns has been noted by Apresjan (1973) and by Nunberg (1995), who discusses them under the label of *systematic polysemy*, noting that languages differ with respect to the set of regular metonymies available.

The grouping of metonymies into clusters presumably facilitates their understanding and makes language use more economic. However, a correct and detailed interpretation of metonymies always requires the activation of world knowledge. For example, the reader or listener must figure out which target entity the vehicle entity is related to, or among a set of related entities, select the one that is intended by the writer/speaker. To illustrate, in OBJECT USED FOR USER metonymies (Examples (7)), this interpretation involves accessing knowledge on the *function* of the user: *player* of the sax, *killer* (using a gun), *driver* (but not passenger) of the bus, *passenger currently occupying* a numbered seat. Probably because of the relative unpredictability of the function of the user, (Nissim and Markert, 2003, p. 57) call Example (7d) “unconventional”. In some cases, the meaning of metonymies is indeed underspecified, especially when they are presented out of context.

### 2.1.2 Challenges

A prominent discussion of certain complex cases of metonymy relevant to computing is presented by Nunberg (1995). Examples (9) and (10) illustrate the phenomenon.

(9) *Roth* is Jewish and widely read.

(10) *The newspaper* Mary works for was featured in a Madonna video.

The theoretical question invited by these examples is whether it is correct to consider the noun phrases as metonymical expressions, or whether one of the verb phrases (predicates) is not used literally. The difference can be exemplified by the respective paraphrases, as in (11b) versus (12b).

(11) a. *Roth* is Jewish and widely read.

b. *Roth* is Jewish and *the works by Roth* [are] widely read.

(12) a. Roth is Jewish and *widely read*.

b. Roth is Jewish and *the works by Roth are widely read*.

The second solution (12) is known as *predicate transfer*. If *widely read* can mean “the works by *X* are widely read”, then *widely read* can be predicated of a HUMAN. This interpretation is motivated by the idea that a noun phrase should not have two meanings concurrently: *Roth* should only refer to an entity of type HUMAN; the interpretation in (11b), where *Roth* means both *Roth* and *the works by Roth* should be avoided. Predicate transfer brings metonymy closer to syntactic phenomena (Nunberg, 1995; Copestake and Briscoe, 1995; Hobbs, 2001; Warren, 2002).

This discussion might seem theoretical and the examples constructed, had they not been confirmed by recent corpus studies. When annotating a corpus for metonymically used place names, Markert and Nissim (2003) found examples where two phrases within one sentence triggered two different meanings of the same noun, mentioned only once. For instance, (13) invokes both a literal reading and a metonymic reading of the noun *Nigeria*.

- (13) ... they arrived in *Nigeria*, hitherto a leading critic of the South African regime ... (Nissim and Markert, 2003, p. 57)

In (13), the initial verb phrase, headed by “arrived”, triggers a literal reading of *Nigeria*, whereas the final apposition, an elliptic predicate describing Nigeria as a “leading critic”, invites a metonymic reading following the pattern PLACE FOR PEOPLE. Markert and Nissim call this phenomenon a “mixed reading” of the noun. In their corpus study, it occurs especially often with coordinations and appositions.

Despite the challenge presented by the mixed reading phenomenon, most computational approaches to metonymy still consider the nouns to be used metonymically (and, if necessary, to be the bearer of two interpretations simultaneously); this solution is also known as **type coercion**. Others explicitly model the verb as having a metonymical meaning, following the **predicate transfer** view. Finally, computational approaches can also be constructed in such a way that they are compatible with either theoretical interpretation.

## 2.2 Computationally oriented corpus studies on metonymy

Corpus studies addressing qualitative and quantitative aspects of metonymy have mainly focused on selected types of regular metonymy, such as PHYSICAL OBJECT FOR PROCESS or PLACE FOR PEOPLE. These patterns are wide-spread among different languages. Some of them are typically instantiated by proper nouns, facilitating the retrieval of metonymy candidates in large corpora. Proper nouns are also particularly relevant to the practical tasks of automatic summarization and Information Extraction. To illustrate, the ACE corpus (LDC, 2005), created for purposes of Information Extraction and related tasks, has been manually annotated for some types of metonymy.

Markert and Nissim (2003) report on a corpus study concerning *country names* as vehicle entities. In a set of 1,000 occurrences, extracted from the British National Corpus (BNC, 100 million words) together with a three-sentence context, three metonymic patterns account for most of the metonymical instances:

1. PLACE FOR PEOPLE, as in (14) and (15), and further subdivided into several subclasses;

(14) *America* did try to ban alcohol

(15) a 29th-minute own goal from *San Marino* defender Claudio Canti

2. PLACE FOR EVENT as in Example (2a), above;

3. PLACE FOR PRODUCT: the place for a product manufactured there (e.g., *Bordeaux*).

Manually labeling the 1,000 occurrences according to eleven categories, including the three metonymic patterns given above, **literal**, and **mixed** (for “mixed readings”), yields 925 unambiguous classifications, with 737 or 79.7% of the occurrences being literal. Among the metonymic usages, the largest subgroup is PLACE FOR PEOPLE, attested by 161 examples. With three members, PLACE FOR EVENT metonymies constitute the smallest group in the corpus study.

Markert and Nissim (2006) report on the annotation of 984 instances of *organization names*. The largest portion is used literally (64.3%). Among five metonymic categories, defined by Markert and Nissim, ORGANISATION FOR MEMBERS (16) was the most frequent one (188 instances).

- (16) Last February *NASA* announced [...].

In ORGANISATION FOR MEMBERS metonymies, the concrete referents are often underspecified. Possible types of target entities include, for example, SPOKESPERSON, (certain classes of) EMPLOYEES, or ALL MEMBERS.

A different set of studies deals with metonymies of the type PHYSICAL OBJECT FOR PROCESS, addressing the issue whether their supposedly most likely interpretation(s) can be derived from corpus texts. To illustrate, (17b)–(17c) are likely interpretations of the metonymy in (17a).

- (17) a. John began *the book*.  
b. John began *reading the book*.  
c. John began *writing the book*.

In general, the question is: For the pattern *begin V NP*, given a head noun of the noun phrase, such as *book*, what are likely values of *V* in a corpus of English? Briscoe *et al.* (1990) use the Lancaster-Oslo/Bergen corpus (1 million words) to investigate this. However, they find the pattern *begin V NP* to be rare when the value of *V* corresponds to a highly plausible interpretation of *begin NP*. This finding has been confirmed by Lapata and Lascarides (2003) for the BNC. For example, summarizing the output of a partial parser, they find that *begin to pour coffee* is the most frequent instantiation of *begin V (DET) coffee*, and that no attestations of *begin to drink coffee* are contained in the corpus, although this might be a plausible interpretation as well. Another problem with this kind of pattern-based corpus analysis is that it ignores the wider context and thus does not attempt to disambiguate between the set of interpretations that can be derived from the corpus.

In the next section, several computational approaches to metonymy will be presented, including a paraphrasing approach to the PHYSICAL OBJECT FOR PROCESS metonymy (2.3.1) that addresses some of these problems.

## 2.3 Computational Approaches to Metonymy

Computational approaches to metonymy are usually not limited to a literal-nonliteral distinction. In other words, rather than just attempting to detect *whether* a given linguistic expression is used metonymically or not, these approaches also try to *interpret* metonymical expressions in some way, as opposed to most approaches to metaphor (Section 3) or idioms (Section 4). In what follows, we will present some examples, ranging from symbolic/logic systems (2.3.1) to paraphrasing approaches (2.3.2) to methods combining detection and interpretation (2.3.3).

### 2.3.1 Metonymy and inference

In symbolic applications using hand-crafted knowledge bases, such as expert systems and question-answering applications, metonymy is one of the phenomena hampering user interaction. If no processing or interpretation of a user query is undertaken, the encoded knowledge base must match the user input exactly, in order for the system to retrieve the adequate knowledge, make the necessary inferences, and finally return results. However, metonymical expressions in user-formulated queries can usually not be resolved directly against the contents of a knowledge base, given that (a.) axioms in knowledge bases use entity types as hard constraints, and (b.) in metonymy, one type of entity stands for another one. Interpreting metonymies and other types of figurative language is thus crucial for these systems.

Against this background, Hobbs (2001) discusses different varieties of metonymy. For a knowledge base of axioms and assertions, he designs additional axioms to deal with (noun) type coercion

and predicate transfer, facilitating the correct parsing of sentences such as (18) and (19). In (19b), “first” really modifies a LOSING event, not the tooth.

(18) John read Proust.

(19) a. She lost her first tooth.

b. She had her first loss of a tooth.

Fan and Porter (2004) address phenomena close to metonymy under the heading of *loose speak*, exemplified by user queries to a chemistry knowledge base. Since expressions produced by humans are less precise than the axioms in a knowledge base, the aim of their work is to find the right axiom given an “imperfect” (for example, metonymical) query.

### 2.3.2 Paraphrasing metonymy

Paraphrasing approaches to metonymy typically presuppose metonymic input. To each input instance, these systems assign a specific interpretation in the form of the most likely paraphrase. This approach is used by Lapata and Lascarides (2003) to process selected metonymies of the type PHYSICAL OBJECT FOR PROCESS. Each input instance takes the abstract form  $V - NP$  or  $v - o$ , with noun  $o$  being the head noun of the object  $NP$  of verb  $v$ , such as *begin cigarette* or *enjoy book*.

The task of the computational model is to automatically acquire possible interpretations of these metonymic verb-noun combinations from corpora, where an interpretation is given in the form of a verb referring to the process that is left implicit in the metonymic expression. For example, given the input *enjoy book*, the system generates *read*. This ultimately facilitates the paraphrasing of (20a) by (20b).

(20) a. John enjoyed *the book*.

b. John enjoyed *reading the book*.

The method is unsupervised insofar as it requires no lexical-semantic resources or annotations. To estimate probabilities for the implicit process, such as READING in (20), distributional information in the form of corpus co-occurrence frequencies is used. In particular, the method looks at verbs possibly referring to this process ( $p$ ), appearing as a complement to the submitted verb  $v$  (for example, *enjoy*), or  $p$  taking an object with, as its head, the submitted noun  $o$  (for example, *book*). The model combines the probabilities of

- a seeing a reference to the possibly implicit process ( $p$ ) in the corpus (the number of times a given verb referring to  $p$  is attested, divided by the sum of all verb attestations);
- b seeing the verb  $v$  with an explicit reference to process  $p$  (the number of times  $v$  takes  $p$  as its complement, as obtained from the output of a partial parser, divided by the number of times  $p$  is attested);
- c seeing the noun  $o$  as the object of  $p$  (the number of times  $o$  is the object of the process-denoting verb  $p$ , divided by the number of times  $p$  is attested).

The last parameter is an approximation of the actual paraphrase that does not take into account the contribution of the verb  $v$  to the paraphrase itself. This is so because it would be problematic to

model the entire paraphrase being searched: Corpus studies show that in many cases, the likelihood of uttering the metonymic construction is higher than that of uttering its full interpretation, as laid out in Section 2.2, above.

The intersection of the most frequent process verbs  $p$  related with the verb  $v$  and the most frequent process verbs  $p$  which take noun  $o$  as their object is the set of corpus-derived interpretations of the verb-noun metonymy. For example, for *enjoy film*, this set is constituted by the verbs *see*, *watch*, and *make*, using the BNC (Lapata and Lascarides, 2003). These can be ranked by probability.

An extended version of the model further takes into account the contribution of the sentential subject, thus preferring different paraphrases for (21) and (22).

(21) The composer began the symphony.

(22) The pianist began the symphony.

Lapata and Lascarides (2003) evaluate the performance of the model by measuring agreement with human paraphrase ratings. For selected metonymic inputs, paraphrase interpretations are derived with the model, and human judgments on these paraphrases are elicited. The human ratings correlate reliably with the model probabilities for the paraphrases, yielding a correlation of .64, against an upper bound (intersubject agreement) of .74.

This approach is compatible with either theoretical account of metonymy: coercion of the noun or predicate transfer (see Section 2.1.2, above).

### 2.3.3 Metonymy recognition as Word Sense Disambiguation

Word sense disambiguation approaches to metonymy integrate the *detection* of metonymical expressions with their *interpretation*, where to interpret the expression is tantamount to *classifying* it according to a pre-defined set of metonymical patterns. For example, Nissim and Markert (2003) distinguish between literally and metonymically used country names (detection), and further label metonymic instances with one of five pre-defined patterns (interpretation): 1. PLACE FOR EVENT, 2. PLACE FOR PEOPLE, 3. PLACE FOR PRODUCT, 4. other metonymy (rare for country names), and 5. mixed metonymy. This and all other approaches mentioned in this section thus consider the noun (the country name) as metonymic, as opposed to the predicate.

Nissim and Markert model the task as a word sense disambiguation (WSD) problem: each of the six labels (**l**iteral plus the five metonymy labels) represents one of the possible interpretations or “senses” of a country name. The method is a *supervised* machine learning approach, since it presupposes a manually annotated corpus of country name attestations. It is also *class based*, which makes it slightly different from traditional WSD methods. A traditional supervised WSD algorithm assigns a word sense label from a pre-defined set (for instance, *bank#1*, *bank#2*, etc.) to occurrences of a given word (*bank*) in an unseen test corpus, after having been trained on correctly sense-labeled occurrences of that word. The class-based metonymy recognition method assigns either the label **l**iteral or that of the relevant metonymic pattern to words of a given semantic class (for instance, the words *U.S.*, *Japan*, *Malta* from the country name class), after having seen the same *or different* words from the same class (for instance, *England*, *Scotland*, *Hungary*) in a manually labeled training corpus.

The class-based approach is supposed to take advantage of the regularities of systematic polysemy, even if exemplified by different lexical material. The model exploits several similarities that follow from regular polysemy:

1. *semantic class similarity* of the possibly metonymic noun: for example, *Japan* in the test data is similar to *England* in the training data,
2. *grammatical function similarity*: for example, a possibly metonymic noun appears in subject position both in the training and test data,
3. *head word (verb) similarity*: for example, in (23) and (24), similar events are denoted by the verbs *win* and *lose*, which are the semantic heads of metonymically used words (*Pakistan*, *Scotland*). A dependency grammar framework is used to identify semantic heads.

(23) *Pakistan* had **won** the World Cup.

(24) *Scotland* **loses** in the semi-final.

For WSD purposes in general, different machine learning techniques have been experimented with. For metonymy recognition, Nissim and Markert (2003) train a decision list classifier on a manually annotated corpus of 925 instances of country names (see also Section 2.2, above) and estimate probabilities via maximum likelihood with smoothing. Three different algorithms (Algorithm I to III) are experimented with and combined to achieve best results. Performance is evaluated by 10-fold cross-validation, where different subsets of the training data are withheld for testing and the results averaged.

**The basic feature – Algorithm I.** For this kind of machine learning approach, it is important to select relevant *properties* of the words to disambiguate or classify, called *training features*. These properties must be measurable in both the training and the test data. The basic version of the metonymy classifier uses exactly *one feature*, called **role-of-head**, composed of the following information:

1. the grammatical role (grammatical function) of the possibly metonymic word with respect to its head, to represent grammatical function similarity; Nissim and Markert (2003) use a set of seven functions, including *active subject (subj)*, *subject of a passive sentence*, and *modifier in a prenominal genitive*;
2. the lemmatized lexical head of the possibly metonymic word, to represent head word similarity.

For example, the value of the **role-of-head** feature for (23) with the possibly metonymic word *Pakistan* is **subj-of-win**. The feature is manually annotated on the training corpus. The *accuracy* of the classifier, defined as the number of correct decisions divided by the number of decisions made, is high (.902). However, due to data sparseness, the classifier has not always seen the exact value of the **role-of-head** feature in the training data, and thus makes decisions in only 63% of the cases. The remaining instances are submitted to a backoff procedure that always assigns the majority reading, **literal**. This results in a low *recall* for metonymies: Only 18.6% of the metonymies are identified.

Low coverage is the main problem of this algorithm. Therefore, two methods are used to generalize the **role-of-head** feature, *in those cases where a decision cannot be made* based on the full feature. This means that data already classified by the basic algorithm remains untouched by algorithms II or III, and generalization applies only to those examples that would otherwise be sent to the simple backoff procedure.



**Generalization – Algorithm II.** To generalize from an unseen test data value of `role-of-head` to a value seen in the training data, this method replaces the unseen lexical head by similar words, obtained from a thesaurus (Lin, 1998). For example, supposing the feature value `subj-of-lose` has not been seen during training, the basic algorithm cannot classify sentence (24), above. However, the training data might have contained sentence (23), above, which is similar and can help with the decision. The thesaurus suggests a generalization from *lose* to *win*, thus facilitating a decision for test sentence (24) based on information from training sentence (23).

For each submitted content word, the thesaurus returns a ranked list of similar words. The generalization algorithm first substitutes with the most similar word, and then works down the list. It stops as soon as a decision can be made; less similar words (up to 50) are thus less likely to influence the decision.

Thesaurus generalization raises the recall of metonymies to 41%. Accuracy before backoff drops to 0.877, but accuracy after backoff increases. Whereas the basic algorithm analyzes only approximately one third of all metonymies, sending the remaining ones to backoff or generalization, the thesaurus-based method deals with a much higher proportion of metonymies: It is applied to 147 instances, 46% of which are metonymies.

**Generalization – Algorithm III.** The second generalization method is to rely on the grammatical role as the only feature. A separate decision list classifier is trained for this parameter. Country names with the grammatical roles *subject* or *subject of a passive sentence* are predominantly metonymic, whereas all other roles are biased towards a literal reading. The algorithm thus assigns all target words in subject position to the non-literal group. It performs slightly better than thesaurus-based generalization.

**Combination.** Performance is best when the two generalization methods are combined such that grammatical role information is used only when the possibly metonymic word is a subject, and thesaurus information otherwise. This method yields a metonymy recall of 51%; accuracy is 0.894 before backoff and 0.87 after backoff.

**Discussion and extensions.** The data used for this experiment is very clean. First, all instances of names that were not country names were removed. Second, only those instances that are decidable for a human annotator were retained. Finally, the manual annotation of the training feature ensures high quality of the data. Deriving the feature value from the output of an automatic parser reduces performance by about 10%. Nissim and Markert (2003) do not break down their evaluation into different classes of metonymies; results are thus averaged over all metonymic patterns.

Other approaches to proper name metonymies include the work by Peirsman (2006), experimenting with an unsupervised approach based on an algorithm proposed by Schütze (1998) and reminding of Latent Semantic Analysis (Landauer and Dumais, 1997). Though failing to achieve the majority baseline for accuracy, this unsupervised WSD algorithm finds two clusters that significantly *correlate* with the manually assigned literal/non-literal labels. Finally, Markert and Nissim (2007) report on the SemEval-2007 competition on metonymy resolution, with five participants in the country name task. The best performing systems exploited syntactic features and made heavy use of feature generalisation, integrating knowledge from lexical databases. With respect to the individual classes of metonymies, only those covered by a larger number of examples in the training

data could be recognized with reasonable success; for country names, this is the PLACE FOR PEOPLE class. Interpreting metonymic expressions besides detecting them is thus still a challenging goal.

## 2.4 Summary

Linguistic theory has identified different subtypes of metonymy, clustered around the basic notion of relatedness or contiguity of two entities, giving rise to linguistic “short-cuts”, where one entity stands for another. From a theoretical point of view, the ensuing semantic incompatibilities have been treated as coercion of either the noun (altering its semantic type) or the predicate of the involved expression, with the aim of reconciling their semantics. Computational approaches reflect either view, and some have found ways to avoid this theoretical decision altogether, still coming up with a practical solution. Recently, a concentration on regular metonymy involving proper nouns can be observed. Although immediately relevant to some applications, where named entities play a central role, this concentration narrows down the phenomenon and might make it seem overly regular. Some types of metonymy, partly discussed here against the background of inference systems, are more creative and irregular, requiring larger amounts of world knowledge and relations for their resolution. In fact, even named entities can participate in more creative metonymic expressions, and when they do, current statistical approaches fail to interpret them.

## 3 Metaphor

In linguistics and in philosophy, there is rich continuing theoretical debate on the definition and use of metaphor. Computational approaches usually make reference to a conceptual model of metaphor that is based on empirical findings in *cognitive science*.

Several decades of cognitive science research suggest that there are powerful primary schemas underlying much of human language and thought (Feldman, 2006; Lakoff, 1987; Lakoff and Johnson, 1980; Langacker, 1987b; Johnson, 1987; Slobin, 1997; Talmy, 1988, 1999). These schemas arise from embodied interaction with the natural world in a socio-cultural setting and are extended via *conceptual metaphor* to structure the acquisition and use of complex concepts. Specifically, cross-cultural and cross-linguistic research has revealed that the structure of abstract actions (such as states, causes, purposes, means) is characterized cognitively in terms of *image schemas* which are *schematized recurring patterns* from the embodied domains of force, motion, and space.

Section 3.1 introduces the theory of metaphor that forms the basis of the existing computational models. Section 3.2 describes different computational implementations of the theory. A discussion of on-going research and an outlook are provided in Section 3.3.

### 3.1 Conceptual Metaphor

As opposed to metonymy, where the related entities are constituent elements of *the same* conceptual domain, a conceptual metaphor relates elements from two *different* domains. Typically, these domains are *experientially based concepts*, on the one hand, and *abstract concepts*, on the other. A conceptual metaphor is the *systematic* set of correspondences that exist between constituent elements of these two domains.

Conceptual metaphors typically employ a more abstract concept as *target domain* and a more concrete or physical concept as their *source domain*. For instance, metaphors of TIME rely on more

STATES ARE LOCATIONS (BOUNDED REGIONS IN SPACE)
CHANGES ARE MOVEMENTS (INTO OR OUT OF BOUNDED REGIONS)
CAUSES ARE FORCES
ACTIONS ARE SELF-PROPELLED MOVEMENTS
PURPOSES ARE DESTINATIONS
MEANS ARE PATHS (TO DESTINATIONS)
DIFFICULTIES ARE IMPEDIMENTS TO MOTION

Table 1: Metaphorical mappings that conceptualize a part of Event Structure

concrete concepts. This is illustrated by linguistic expressions, such as (25) and (26).

(25) the days [the more abstract or target concept] *ahead*

(26) *giving* my time

In (25), time is conceptualized as a path into physical space, evoked by the metaphorical use of *ahead*. In (26), time is a substance that can be handled as a resource in a transaction, or offered as a gift, a notion evoked by *give*. Different conceptual metaphors tend to be invoked when the speaker is trying to make a case for a certain point of view or course of action. For instance, one might associate “the days ahead” with leadership, whereas the phrase “giving my time” carries stronger connotations of bargaining.

A primary tenet of conceptual metaphor theory is that metaphors are matter of thought and not merely of language: hence, the term conceptual metaphor. The metaphor does not just consist of words or other linguistic expressions that come from the terminology of the more concrete conceptual domain, but conceptual metaphors underlie a system of related metaphorical expressions that appear on the linguistic surface. Similarly, the mappings of a conceptual metaphor are themselves motivated by image schemas, which are pre-linguistic schemas concerning space, time, moving, controlling, and other core elements of embodied human experience.

Many conceptual metaphors are cross-cultural and highly productive. An example is the Event Structure Metaphor (ESM), which has been found in all cultures studied to date (Lakoff and Johnson, 1980; Johnson, 1987; Langacker, 1987a; Lakoff, 1994). Through the ESM, our understanding of spatial motion (movement, energy, force patterns, temporal aspects) is projected onto abstract actions (such as psychological, social, political acts, or economic policies). The ESM is composed from a set of *Primary metaphors* (Grady, 1997), such as ACTIONS ARE SELF-PROPELLED MOVEMENTS, which compose in complex ways to build large metaphor systems. Figure 1 lists some examples of the primary metaphors that structure our conceptualization of events. Grady (1997) argues from developmental and linguistic evidence that the primary metaphor system forms the developmentally early mappings from which are composed more complex metaphors such as the ESM.

The compositional mapping (ESM) generalizes over an extremely wide range of *expressions* for one or more aspects of event structure. With respect to STATES and CHANGES, examples include (27) to (28).

(27) being *in* or *out* of a state (for example, trouble)

(28) *getting into* a state or *emerging from* a state (for example, depression)

Parts of the ESM metaphor interact in complex ways. Consider the submappings DIFFICULTIES ARE IMPEDIMENTS TO MOTION, ACTIONS ARE SELF-PROPELLED MOVEMENTS, and PURPOSES ARE DESTINATIONS. A difficulty is something that impedes motion to a destination, conceptualized as hindering purposeful actions. Metaphorical difficulties of this sort come in five types: blockages (29a); features of the terrain (29b); burdens (29c); counterforces (29d); and lack of an energy source (29e).

(29) DIFFICULTIES ARE IMPEDIMENTS TO MOTION

a. DIFFICULTIES ARE BLOCKAGES

- i. He *got over* his divorce.
- ii. He's trying to *get around* the regulations.
- iii. We *ran into a brick wall*.

b. DIFFICULTIES ARE FEATURES OF THE TERRAIN

- i. It's been *uphill* all the way.
- ii. We've been *bogged down*.
- iii. We've been *hacking* our way through a *jungle* of regulations.

c. DIFFICULTIES ARE BURDENS

- i. He's *weighed down* by lot of assignments.
- ii. He's been trying to *shoulder* all the responsibility.

d. DIFFICULTIES ARE COUNTERFORCES

- i. Quit *pushing* me *around*.
- ii. He's *holding* her *back*.

e. A DIFFICULTY IS A LACK OF AN ENERGY SOURCE

- i. I'm *out of gas*.
- ii. We're *running out of steam*.

Many abstract and contested concepts in politics, economics, and even mathematics may be metaphoric (Lakoff, 1994; Lakoff and Núñez, 2000). For example, Lakoff argues that metaphor is central to the core concept of FREEDOM, and that this abstract concept is actually grounded in bodily experience. Physical freedom is freedom to move – to go places, to reach for and get objects, and to perform actions. Physical freedom is defined in a frame in which there are potential impediments to freedom to move: blockages, being weighed down, being held back, being imprisoned, lack of energy or other resources, absence of a path providing access, being physically restrained from movement, and so on. Freedom of physical motion occurs when none of these potential impediments is present.

Various metaphors turn freedom of physical motion into freedom to achieve one's goals. The ESM, for instance, characterizes achieving a purpose as reaching a desired destination, or getting a desired object. Freedom to achieve one's purposes then becomes the absence of any metaphorical impediments to motion. Other ideas, like political freedom and freedom of the will, build on that concept. The concept of political freedom is characterized via a network of concepts that necessarily includes the ESM and the inferences that arise via that metaphor.

## 3.2 Computational Models of Metaphor

Computational approaches to metaphor have focused on one or more of the following issues: inference (Section 3.2.1), recognition in text and discourse (Section 3.2.2), and acquisition and representation (Section 3.2.3).

### 3.2.1 Metaphor and Inference

As illustrated by the examples provided above (Section 3.1), and further documented by corpus analyses (see Section 3.2.2), below), conventionalized metaphor is an everyday issue. Most Natural Language Processing (NLP) systems have to face it sooner or later. A successful handling of conventional metaphor is also the first step towards the processing of novel metaphor.

Obvious problems for NLP systems caused by metaphorical expressions consist in the incompatibility of metaphorically used nouns as arguments of verbs. In systems which constrain the type of arguments for every verb by semantic features like **human**, **living**, **concrete** or **abstract** (“selectional restrictions”), metaphors can cause inconsistencies that have to be solved. For example, if the grammatical subject of the English verb *go* was restricted to entities classified as **living** in a given system, the sentence (30) taken from Hobbs (1992) could not be parsed.

(30) The variable N *goes* from 1 to 100.

Obviously, there is an open-ended number of such sentences. To increase the ability of systems to deal with incompatibilities of this kind, caused by instantiations of conceptual metaphors, systems including metaphor processing mechanisms have been designed and implemented. These *knowledge-based* systems encode a representation of at least a part of the conventionalized mapping, leveraging “knowledge of systematic language conventions in an attempt to avoid resorting to more computationally expensive methods” (Martin, 1994). The systems generally perform much of the necessary knowledge-based *reasoning* in the source domain and transfer the results back to the target domain using the mapping representation. This procedure is applied, for example, in KARMA’s networks (Narayanan, 1999; Feldman and Narayanan, 2004) or in the rules of TACITUS (Hobbs, 1992) and ATT-Meta (Barnden and Lee, 2001).

Both KARMA (Narayanan, 1997a, 1999) and ATT-Meta (Barnden and Lee, 2001) are systems that can reason explicitly with metaphors, based on the theory of conceptual metaphors. These two systems have different motivations, and at least in their current state, offer a slightly different functionality. KARMA can be seen as a *story understanding system*, producing a temporally evolving state representation in the target domain, and an output that provides the Most Probable Explanation (MPE) of the input so far, combining metaphoric inference with background knowledge in the source and target domains. ATT-Meta can be seen as a *Question-Answering system* that verifies whether a fact (submitted as a user query) holds, given a possibly metaphorical representation of the current state of the world.

There are some fundamental differences in the computational design and particular abilities of the two systems. For example, at design and implementation level, KARMA uses x-schemas (“executing schemas” (Feldman and Narayanan, 2004, p. 387)) implemented as extended Stochastic Petri Nets for the source domain and Temporally Extended Bayes networks (aka Dynamic Bayes Nets) for the target domain. The ATT-Meta approach uses situation-based or episode-based first-order logic throughout (Barnden and Lee, 2001, p. 34).

In terms of the specific strengths of the systems, KARMA was developed in the context of a larger research program on the neural theory of language acquisition and use (Feldman, 2006) and modeled the role of imaginative simulation with sensory-motor representations for inference about abstract actions and events (Narayanan, 1997a). The system demonstrated the beginnings of the ability to detect interpreter’s bias in metaphorical utterances; for example, (31) implies that the speaker or writer in general opposes government control of economy, whereas (32) does not imply this.

(31) Government *loosened strangle-hold* on business.

(32) Government *deregulated* business.

KARMA detects this difference in speaker attitude because it can draw source domain inferences of “strangle-hold”, which has a detrimental effect on business (in the target domain); see (Narayanan, 1999).

The ATT-Meta system addresses deductive and simulative inference issues and requires only a minimal set of inter-domain mappings. This emphasises the importance of complex reasoning within domains, especially within the source domain; see (Barnden and Lee, 2001; Barnden *et al.*, 2002).

A commonality between KARMA and ATT-Meta is their general architecture. Both have three main components, or representational domains, which reflect the influence of central ideas from metaphor theory:

1. Knowledge about the source domain, including common-sense knowledge, factual, relational and active causal and procedural knowledge.
2. Knowledge about the target domain. This is usually less elaborate than source domain knowledge but contains facts, relations and causal knowledge about the target.
3. Mapping information (KARMA) or conversion rules (ATT-Meta) to transfer knowledge across domains. These mapping mechanisms can be of various types and may also include information about features or facts *not* to be varied or changed while producing the mapping, such as aspect (for example, ongoing vs. completed action).

Both systems rely on extensive domain knowledge which has to be manually coded by the designer or user. The computational engines of the systems (x-schemas and Bayes nets in KARMA or backchaining/backward reasoning of rules in ATT-Meta) then derive new facts or consequences within the domains—both within the source and target domains.

### 3.2.2 Metaphor recognition from Text

Linguistic realizations of conceptual metaphors are ubiquitous in everyday speech and text. In a text experiment by (Gedigian *et al.*, 2006), over 90% of the uses of motion terms (*fall, move, stumble, slide*) in the Wall Street Journal are abstract usages (about stock market activity, international economics, political acts). While this staggering figure may partly be explained by the high prior on articles about politics or the stock market in the selected sections, a more balanced corpus (the BNC) yields abstract uses in 61% of the instances where motion terms were used. This explains why semantically oriented language analyzers that are trained on carefully selected, gold standard

semantic resources such as FrameNet (Fillmore *et al.*, 2003) often perform poorly when applied to newspaper text: they cannot handle metaphoric uses.

Some recent experiments thus attempt to automatically identify metaphorical expressions in text. These approaches have been implemented in stand-alone systems. To our knowledge, non of them has been integrated into a general automatic analyzer. A further restriction that applies to most of the current work on metaphor recognition, including Mason (2004), Gedigian *et al.* (2006), and Birke and Sarkar (2006), is its focus on *verbs* as parts of speech. A further approach to metaphor recognition, addressing also nouns, has recently been presented by (Krishnakumaran and Zhu, 2007). In the remainder of this section, we will discuss metaphor recognition approaches based on two examples.

A clustering approach for separating literal from non-literal language use, has been implemented by Birke and Sarkar (2006) for verbs. The algorithm is a modification of the similarity-based WSD algorithm by Karov and Edelman (1998), where similarities are calculated between sentences containing the word to be disambiguated (target word), and collections of seed sentences for each word sense (feedback sets). In the case of metaphor recognition, there are only two feedback sets for each verb: literal or nonliteral. The original algorithm attracts a target sentence to the feedback set containing the (*one*) most similar sentence. In the modified version used to discriminate between literal and metaphorical uses, a target sentence is attracted to the feedback set to which it is most similar, *summing across* its similarities to each sentence per feedback set.

Birke and Sarkar do not make reference to a particular theory of metaphor or of figurative language in general. Consequently, their literal-nonliteral distinction is relatively vague: “*literal* is anything that falls within accepted selectional restrictions [...] or our knowledge of the world [...]. *Nonliteral* is then anything that is “not literal” [...]” (Birke and Sarkar, 2006, p. 330) This vagueness possibly decreases the conclusiveness of the results, given that a clear definition of the phenomenon is necessary to create consistently annotated feedback sets and test data.

The results of the clustering approach are reported in terms of f-score<sup>1</sup>, but precision and recall are not provided individually. The original WSD algorithm, attracting test instances to the set containing the most similar sentence, achieves an f-score of 36.9%. The f-score of the modified algorithm is 53.8%, averaged over 25 words in the test set. Active learning, where some examples are given back to a human annotator for decision during the classification, further increases f-score to 64.9%.

A different approach to metaphor identification, similarly discriminating between literal and metaphorical usages of verbs, has been implemented by Gedigian *et al.* (2006). They trained a maximum entropy classifier on examples chosen from concrete domains that are likely to yield metaphors. Appropriate lexical expressions (lexical units) are thus collected from concrete frames in FrameNet, especially frames related to MOTION and to caused motion, such as PLACING, but also CURE. Sentences from the PropBank (Kingsbury and Palmer, 2002) Wall Street Journal corpus containing these lexical units are extracted and annotated. As described earlier, more than 90% of the 4,186 occurrences of these verbs in the corpus data are metaphors.

The features used by the classifier include information on the semantic type of the arguments of the verbs. This reflects selectional preferences, introduced as selectional restrictions in Subsection 3.2.1, which are an important factor for determining whether a verb is being used metaphorically or not. Verb arguments and their semantic type are extracted from the PropBank corpus as follows:

---

<sup>1</sup> $2 \times \text{precision} \times \text{recall} / (\text{precision} + \text{recall})$ .

1. Argument information is extracted from existing PropBank annotations.
2. The head word of each argument is identified, using a method proposed by Collins (1999).
3. A semantic type is assigned to the head word. How this is done depends on the type of head word:
  - If the head is a pronoun, a pronoun type (human/non-human/ambiguous) is assigned.
  - If the head is a named entity, the semantic type of the argument is a tag assigned by an off-the-shelf named entity recognizer, thus generalizing over different classes of named entities.
  - Otherwise, the name of the head’s WordNet synset (Fellbaum, 1998) is used as the type of the argument, thus generalizing over synonyms.

A further feature is the bias of the target verb itself. This feature is useful because most verbs show a clear tendency towards either literal or metaphorical uses within the corpus. To determine the best feature combination for metaphor detection, the classifier is trained with different combinations and validated on a validation set. Results are best for the combination of the following features:

1. verb bias;
2. semantic type of argument 1 (ARG1), typically realized as the direct object in active English sentences such as (33);
 

(33) Traders *threw* **stocks** out of the windows.
3. (optionally,) semantic type of argument 3 (ARG3), the semantics and syntactic form of which is difficult to generalize due to verb-specific interpretations in PropBank.

The classifier trained with features 1 and 2 above, over all verbs in all frames, thereby generalizing over source domains. On a test set of 861 targets, achieves an accuracy of 95.12. This is above the baseline of 92.9 overall, achieved by selecting the majority class of the training set. The result also exceeds the alternative baseline accuracy of 94.89, achieved by selecting the majority class of each verb specifically.

Accuracy varies slightly across frames. It is equal or higher than the baselines in all frames, with the exception of the CURE frame. This might be due to verbs with strong biases in the training data; for example, *treat* had no metaphorical uses in the training data. Alternatively, generalizing across all source frames might result in models that do not represent the data from the CURE frame well enough, and a different feature set might be more appropriate for this frame.

### 3.2.3 Acquisition and Representation of Metaphors

As Martin (1994) points out, one of the problems for systems that deal with metaphor, is the acquisition of sufficient and suitable knowledge (Hobbs, 1992; Martin, 1994; Narayanan, 1997b; Barnden *et al.*, 2002). It would thus be useful to provide more knowledge about metaphor in lexical resources, which could be either directly used in Natural Language Processing (NLP) systems, or used as a basis for building rules and networks in systems designed especially for metaphor



handling. If well-studied linguistic knowledge supported by attestations in corpora was encoded in lexical resources, they could also be regarded as a common starting point for different systems, and the results of the systems would become more directly comparable.

Current general-domain lexical semantic resources (WordNet, PropBank, FrameNet) are of restricted usefulness for systems that aim at understanding or creating metaphorical expressions. One reason for this state of affairs is that metaphor captures generalizations across word senses and frames that are not represented in any of the popular linguistic resources. In English, there are specialized lists of metaphors, the most notable of which is the venerable Berkeley Master Metaphor list (Lakoff *et al.*, 1991) which is quite unsuitable for computational use; we will discuss some of its shortcomings below. For specialized use, there is a Mental Metaphor databank created by John Barnden at the University of Birmingham (<http://www.cs.bham.ac.uk/~jab/ATT-Meta/Databank/>) that deals with metaphors of the mind. Another ongoing effort is the *Metaphor in discourse* project (Steen, 2007), where subsets of the BNC are annotated for metaphors at word level, as opposed to the level of conceptual mappings. As far as we know, these databases and annotations have not been linked directly to any general purpose linguistic resource, such as WordNet or FrameNet.

There have also been efforts in other languages which could inform representation and annotation efforts. To our knowledge, the most advanced such effort is the Hamburg Metaphor Database (Lönneker and Eilts, 2004; Lönneker, 2004) which combines data from corpora, EuroWordNet and the Berkeley Master Metaphor List. For example, Reining and Lönneker-Rodman (2007) annotated more than 1,000 instances of lexical metaphors from the MOTION and BUILDING domains in a French newspaper corpus centered on the European Union, and integrated them into the Hamburg Metaphor Database.

The Master Metaphor List, while useful in the efforts described above, has fundamental flaws that severely restrict its wider applicability. The list was built almost two decades ago and subsequent research has made significant advances that directly bear on metaphor representation and annotation. One of the problems with the list is that the mapping ontology does not follow clear structuring principles (Lönneker-Rodman, 2008). Another central problem is that the list is non-compositional, in that there is no principled way to combine the existing mappings to create more complex ones. The main reason for this shortcoming is that projections from specific source to target domains are only partial (many aspects of the source are not mapped, many attributes of the target are not projected onto), making it very hard to generalize from existing mappings to find shared structure and to compose two maps into a more complex metaphor.

At the time of the construction of the Master Metaphor List, there was not enough known about either the developmental aspects of metaphor acquisition nor of the kind of basic metaphors that could provide a basis set for more complex compositions. It can thus be hoped that future resource creation and annotation efforts will take into account more recent research results from Cognitive Linguistics (see Section 3.1, above), such as the work on *primary metaphors* (Grady, 1997) that directly bears on the issue of compositionality. Also, the work by Gedigian *et al.* (2006), discussed in Section 3.2.2, suggests that linking to semantic frames such as those provided by the FrameNet project could significantly help in metaphor representation, by providing a connection via FrameNet frames to linguistic realizations of metaphor.

### 3.3 The road ahead

Existing systems have only scratched the surface of the information being communicated by metaphor. As an example of the subtlety and inferential richness of this information, consider the following types of metaphoric inferences that were revealed during an early evaluation of the KARMA system (Narayanan, 1997b). These mappings and inferences rely on the ESM to routinely communicate causal and relational information about abstract economic and political actions by projecting concepts from the physical domains of forces, spatial motion, and manipulation.

#### 3.3.1 Manner and metaphor

*Distances, speeds, force-values, sizes and energy-levels* are obviously important perceptual and motor control parameters, but with metaphor projections, they become important descriptive features of events in abstract domains, including impacting early parsing decisions of inferring semantic role assignments. Examples abound in ordinary discourse (economies crawl, goals remain far or near, we take giant steps, recoveries are anemic etc.).

#### 3.3.2 Aspectual inferences

Narayanan (1997c) outlined a model of *aspect* (the internal temporal structure of events) which is able to detect and model subtle interactions between grammatical devices (such as morphological modifiers like *be + V-ing* (progressive aspect) versus *has V-ed* (perfect aspect)) and the inherent aspect of events (such as the inherent iterativity of *tap* or *rub*, or the punctuality of *cough* or *hit*). In examining the KARMA metaphor database, Narayanan (1997b) found aspectual distinctions to be *invariantly projected* across domains. Furthermore, the high-frequency of aspectual references in describing events makes it important to model the relevant semantic distinctions.

#### 3.3.3 Goals, resources

It is well known (Wilensky, 1983; Schank and Abelson, 1977) that narratives are generally about *goals* (their accomplishment, abandonment, etc.) and *resources*. KARMA pilot experiments showed that metaphors such as the ESM may in fact be *compactly coding* for these features. Narratives are able to exploit aspects of spatial motion, forces, and energy expenditure to assert information about changing goals and resources. For instance, the amount of *energy* usually maps to *resource* levels as in *slog, anemic, sluggish or bruised and bloodied, or stagger to their feet*. Similarly *slippery slopes, slipperiest stones, slide into recessions*, get projected as the possible thwarting of goals due to unanticipated circumstances. In general, stories in the abstract domain are often about the complex notion of controllability, monitoring problems, and policy adjustments. Again, monitoring, changing directions, rates, etc. are obviously common in sensory-motor activity, and so using these features and appropriate projections allows the speaker to communicate monitoring and control problems in abstract plans.

#### 3.3.4 Novel expressions and blends

As Lakoff (1994); Gibbs (1994); Fauconnier and Turner (2002) and other researchers point out, a variety of novel expressions in ordinary discourse as well as in poetry make use of highly conventionalized mappings such as the ones described here. The expressions “slippery slope”, “crossroads”,

or “rushing headlong on the freeway of love” are all immediately interpretable even if one has no previous exposure to these expressions in the abstract domains of their usage. Indeed, these expressions are interpretable due to the event structure metaphor that maps motion to action.

It appears that even in the cases where there are blends from multiple source domains, as long as they are interpretable and coherent in the *target*, these expressions are interpretable by humans. As far as we know, there has been no model that can scalably deal with such blends (Fauconnier and Turner, 2002).

### 3.3.5 Agent attitudes and affects

Agent attitudes often encode anticipatory conditions, motivation and determination of agents involved. Some of this is implemented in the KARMA system and annotating and studying performance on a larger dataset could be productive for computational approaches. For instance, the expression *taking bold steps* encodes determination in the face of anticipated obstacles/counterforces ahead.

### 3.3.6 Communicative intent and metaphor

One of the important aspects of communication involves evaluative judgments of situations to communicate speaker intentions and attitudes. While the KARMA system showed some ability to handle this phenomenon, we hypothesize that the cross-linguistic prevalence of the use of embodied notions of force and motion to communicate aspects of situations and events is linked to the ease with which evaluative aspects can be communicated in experiential terms. This is likely to be an extremely productive area for computational research in natural language understanding.

## 3.4 Summary

Metaphor is an important aspect of everyday language which provides a great deal of information about both the content and pragmatic aspects of communication. There is now a body of theory within cognitive science and a set of empirical data that can support a systematic analysis of the information content of metaphors. To perform such an analysis would require a scalable version of existing metaphor systems that is applied to a significantly enhanced inference task. To accomplish this would also require adapting and/or the building of a new set of semantic resources to identify, represent, and annotate metaphors. The pieces are there and the time is ripe for these enhancements.

## 4 Other Figurative Language

Metaphor and metonymy have received considerable attention in philosophy, linguistics, media science and similar fields. These two phenomena are clearly central types of figurative language, and have been shown to be frequent in everyday language as well as in more specialized texts. Against this background, it is surprising how relatively little attention has been paid to them from a computational point of view, unless we hypothesize that they have been considered as “too difficult”, and constantly postponed to later research. The individual approaches and systems we have discussed in this chapter are selected examples, but we can safely state that the pool we drew them from does not offer a wealth of additional material.

Many further phenomena can be treated under the heading of figurative language. In particular, idioms such as (34) are another important subcategory.

- (34) a. It's been a while since we last *shot the breeze*.  
b. It's been a while since we last *had a relaxed conversation*.

Idioms extend over a stretch of words. These can be treated as one linguistic entity, a relatively fixed expression with a meaning of its own that cannot be compositionally derived from its constituents. Again, several corpus studies have confirmed that idioms are wide-spread (e.g. Fellbaum *et al.*, 2006; Villavicencio *et al.*, 2004). As their *form* is more constrained, detection of idioms might be easier than that of other non-literal language phenomena. Consequently, idiom extraction (Degand and Bestgen, 2003; Fazly and Stevenson, 2006) exploits lexical and syntactic fixedness of idioms. However, some questions remain open. In particular, the results of corpus investigations (Fellbaum *et al.*, 2006, pg. 350) show that lexical and syntactic variants of idioms are far from rare. For example, a noun participating in an idiom can have lexical variants that *conserve* idiom meaning, and those nouns are likely to be just among the most similar ones in a semantic resource. It is thus not clear whether the usage of a thesaurus to extract similar words when creating variants of a possible idiom (Fazly and Stevenson, 2006) is helpful at all in idiom detection. What is more important, the crucial step from idiom recognition to idiom interpretation has not yet been attempted.

As opposed to the previously discussed phenomena, non-literalness can also arise from the context. In those cases, the figurative meaning of an expression cannot be ascribed to the usage of single words or fixed phrases any more. Rather, we are dealing with *pragmatic* phenomena, including indirect speech acts (35a), irony (36a), and certain types of humor (37).

- (35) a. Do you have a watch?  
b. Please tell me what time it is.
- (36) a. I just love spending time waiting in line.  
b. I hate spending time waiting in line.
- (37) Why did the elephant sit on the marshmallow? – Because he didn't want to fall into the hot chocolate.

The joke in (37) does not have a literal paraphrase because its communicative effect relies on the fact that the situation it describes is absurd, i.e., against the laws of nature, logic or commonsense. Absurdness is one of the factors that have been exploited in computational approaches to humor. Such approaches almost unanimously come along as *humor generation* (as opposed to detection). Humor generation has long been restricted to puns, largely influenced by Binsted's (1996) seminal work on JAPE, a *Joke Analysis and Production Engine*, generating question-answer punning riddles. Recently, two new lines of computational humor generation have appeared: word-play reinterpretation of acronyms (Stock and Strapparava, 2005) and jokes based on the ambiguity of pronouns (Njiholt, 2006; Tinholt, 2007). Sample outputs are presented in (38) to (40).

- (38) What do you call a gruesome investor? *A grisly bear*. (Binsted, 1996, p. 96)
- (39) FBI – Federal Bureau of Investigation - *Fantastic Bureau of Intimidation* (Stock and Strapparava, 2005, p. 115)

- (40) The members of the band watched their fans as they went crazy. *The members of the band went crazy? Or their fans?* (Tinholt, 2007, p. 66)

As with other approaches to figurative language, the strategy of humor generation systems is to reduce the complexity of the task. First, each of them focuses on a particular type of humor. Second, the problem is formulated in terms of certain *patterns* at different levels:

1. a more or less flexible syntactic pattern;
2. an inventory of lexical entities and semantico-conceptual relations between them;
3. in the case of pun generation and acronym deformation, phonologic aspects.

In summary, figurative devices are ubiquitous in language and require conceptual, linguistic, and pragmatic knowledge for interpretation. World knowledge and common sense, and the representation of beliefs, attitudes, and emotional states play an equally important role. Figurative language currently presents a serious barrier to building scalable Natural Language Understanding (NLU) systems. To date, computational models of figurative language have been fairly sparse, and have tended to focus on a single type of language, such as metonymy, idioms, or metaphor. However, taken together, these models have demonstrated the capability of handling a wide variety of phenomena and the time is ripe for a serious computational effort that harnesses the essential insights of these efforts to capture the interactions between multiple literal and figurative devices in an integrated framework.

## References

- Apresjan, J. D. (1973). Regular polysemy. *Linguistics*, **142**, 5–32.
- Barnden, J., Glasbey, S., Lee, M., and Wallington, A. (2002). Reasoning in metaphor understanding: The ATT-Meta approach and system. In *Proceedings of the 19th International Conference on Computational Linguistics (COLING-2002)*, pages 1188–1192, San Francisco, CA.
- Barnden, J. A. and Lee, M. G. (2001). Understanding open-ended usages of familiar conceptual metaphors: an approach and artificial intelligence system. CSRP 01-05, School of Computer Science, University of Birmingham.
- Binsted, K. (1996). *Machine humour: An implemented model of puns*. Ph.D. thesis, University of Edinburgh.
- Birke, J. and Sarkar, A. (2006). A clustering approach for the nearly unsupervised recognition of nonliteral language. In *Proceedings of the 11th Conference of the European Chapter of the Association for Computational Linguistics*, pages 329–336, Trento, Italy.
- Briscoe, E., Copestake, A., and Boguraev, B. (1990). Enjoy the paper: lexical semantics via lexicology. In *Proceedings of the 13th International Conference on Computational Linguistics (COLING-90)*, pages 42–47, Helsinki, Finland.
- Collins, M. (1999). *Head-Driven Statistical Models of Natural Language Parsing*. Ph.D. thesis, University of Pennsylvania.

- Copestake, A. and Briscoe, E. J. (1995). Semi-productive polysemy and sense extension. *Journal of Semantics*, **1**(12), 15–67.
- Degand, L. and Bestgen, Y. (2003). Towards automatic retrieval of idioms in French newspaper corpora. *Literary and Linguistic Computing*, **18**(3), 249–259.
- Fan, J. and Porter, B. (2004). Interpreting loosely encoded questions. In *Proceedings of the Nineteenth National Conference on Artificial Intelligence (AAAI 2004)*, pages 399–405, San Jose, CA.
- Fauconnier, G. and Turner, M. (2002). *The way we think: Conceptual blending and the mind's hidden complexities*. Basic Books, New York.
- Fazly, A. and Stevenson, S. (2006). Automatically constructing a lexicon of verb phrase idiomatic combinations. In *Proceedings of the 11th Conference of the European Chapter of the Association for Computational Linguistics (EACL-2006)*, pages 337–344, Trento, Italy.
- Feldman, J. (2006). *From Molecule to Metaphor*. The MIT Press, Cambridge, MA.
- Feldman, J. and Narayanan, S. (2004). Embodied meaning in a neural theory of language. *Brain and Language*, **89**, 385–392.
- Fellbaum, C., editor (1998). *WordNet: An Electronic Lexical Database*. The MIT Press, Cambridge, MA.
- Fellbaum, C., Geyken, A., Herold, A., Koerner, F., and Neumann, G. (2006). Corpus-based studies of German idioms and light verbs. *International Journal of Lexicography*, **19**(4), 349–360.
- Fillmore, C. J., Johnson, C. R., and Petruck, M. R. L. (2003). Background to FrameNet. *International Journal of Lexicography*, **16**(3), 235–250.
- Gedigian, M., Bryant, J., Narayanan, S., and Ciric, B. (2006). Catching metaphors. In *Proceedings of the 3rd Workshop on Scalable Natural Language Understanding*, pages 41–48, New York City.
- Gibbs, R. (1994). *The Poetics of Mind: Figurative Thought, Language, and Understanding*. Cambridge University Press, Cambridge, UK; New York.
- Grady, J. (1997). *Foundations of meaning: primary metaphors and primary scenes*. Ph.D. thesis, University of California, Berkeley.
- Hobbs, J. R. (1992). Metaphor and abduction. In A. Ortony, J. Salck, and O. Stock, editors, *Communication from an Artificial Intelligence Perspective: Theoretical and Applied Issues*, pages 35–58. Springer, Berlin.
- Hobbs, J. R. (2001). Syntax and metonymy. In P. Bouillon and F. Busa, editors, *The Language of Word Meaning*, pages 290–311. Cambridge University Press, Cambridge, United Kingdom.
- Johnson, M. (1987). *The Body in the Mind: The Bodily Basis of Meaning, Imagination and Reason*. The University of Chicago Press, Chicago, IL.
- Karov, Y. and Edelman, S. (1998). Similarity-based word sense disambiguation. *Computational Linguistics*, **24**(1), 41–59.

- Kingsbury, P. and Palmer, M. (2002). From Treebank to PropBank. In *Proceedings of the 3rd International Conference on Language Resources and Evaluation (LREC-2002)*, Las Palmas, Gran Canaria, Canary Islands, Spain.
- Kövecses, Z. (2002). *Metaphor: A Practical Introduction*. Oxford University Press, New York.
- Krishnakumaran, S. and Zhu, X. (2007). Hunting elusive metaphors using lexical resources. In *Proceedings of the Workshop on Computational Approaches to Figurative Language*, pages 13–20, Rochester, New York.
- Lakoff, G. (1987). *Women, Fire, and Dangerous Things. What Categories Reveal about the Mind*. The University of Chicago Press, Chicago, IL.
- Lakoff, G. (1994). What is metaphor. In J. A. Barnden and K. J. Holyoak, editors, *Advances in Connectionist and Neural Computation Theory: Analogy, Metaphor and Reminding*, volume 3, pages 203–258. Ablex, Norwood, NJ.
- Lakoff, G. and Johnson, M. (1980). *Metaphors we live by*. The University of Chicago Press, Chicago, IL.
- Lakoff, G. and Núñez, R. E. (2000). *Where Mathematics Comes From: How the Embodied Mind Brings Mathematics into Being*. Basic Books, New York.
- Lakoff, G., Espenson, J., and Schwartz, A. (1991). Master metaphor list. Second draft copy. Technical report, Cognitive Linguistics Group, University of California Berkeley. <http://cogsci.berkeley.edu>.
- Landauer, T. K. and Dumais, S. T. (1997). A solution to plato’s problem: The latent semantic analysis theory of acquisition, induction and representation of knowledge. *Psychological Review*, **104**, 211–240.
- Langacker, R. (1987a). *Foundations of Cognitive Grammar I: Theoretical Prerequisites*. Stanford University Press.
- Langacker, R. W. (1987b). *Foundations of Cognitive Grammar, Vol. 1*. Stanford University Press.
- Lapata, M. and Lascarides, A. (2003). A probabilistic account of logical metonymy. *Computational Linguistics*, **29**(2), 261–315.
- LDC (2005). ACE (Automatic Content Extraction) English Annotation Guidelines for Entities.
- Lin, D. (1998). An information-theoretic definition of similarity. In *Proceeding of the 15th International Conference on Machine Learning*, pages 296–304, San Francisco, CA. Morgan Kaufmann.
- Lönneker, B. (2004). Lexical databases as resources for linguistic creativity: Focus on metaphor. In *Proceedings of the LREC 2004 Satellite Workshop on Language Resources and Evaluation: Language Resources for Linguistic Creativity*, pages 9–16, Lisbon, Portugal. ELRA.
- Lönneker, B. and Eilts, C. (2004). A current resource and future perspectives for enriching WordNets with metaphor information. In *Proceedings of the Second International WordNet Conference – GWC 2004*, pages 157–162, Brno, Czech Republic.

- Lönneker-Rodman, B. (2008). The Hamburg Metaphor Database Project. Issues in Resource Creation. *Language Resources and Evaluation*. In press.
- Markert, K. and Nissim, M. (2003). Corpus-based metonymy analysis. *Metaphor and Symbol*, **18**(3), 175–188.
- Markert, K. and Nissim, M. (2006). Metonymic proper names: A corpus-based account. In A. Stefanowitsch and S. T. Gries, editors, *Corpus-based approaches to metaphor and metonymy*, pages 152–174. Mouton de Gruyter, Berlin and New York.
- Markert, K. and Nissim, M. (2007). Semeval-2007 task 08: Metonymy resolution at semeval-2007. In *Proceedings of the 4th International Workshop on Semantic Evaluations (SemEval-2007)*, page 3641, Prague. ACL.
- Martin, J. H. (1994). MetaBank: A knowledge-base of metaphoric language conventions. *Computational Intelligence*, **10**(2), 134–149.
- Mason, Z. J. (2004). CorMet: A computational, corpus-based conventional metaphor extraction system. *Computational Linguistics*, **30**(1), 23–44.
- Narayanan, S. (1997a). *Knowledge-Based Action Representations for Metaphor and Aspect*. Ph.D. thesis, University of California at Berkeley.
- Narayanan, S. (1997b). *Knowledge-based Action Representations for Metaphor and Aspect (KARMA)*. Ph.D. thesis, Computer Science Division, University of California at Berkeley.
- Narayanan, S. (1997c). Talking the talk is like walking the walk: A computational model of verbal aspect. In *Proc. 19th Cognitive Science Society Conference*.
- Narayanan, S. (1999). Moving right along: A computational model of metaphoric reasoning about events. In *Proceedings of the National Conference on Artificial Intelligence (AAAI '99)*, pages 121–129, Orlando, Florida. AAAI Press.
- Nissim, M. and Markert, K. (2003). Syntactic features and word similarity for supervised metonymy resolution. In *Proceedings of the 41st Annual Meeting of the Association for Computational Linguistics*, pages 56–63.
- Njiholt, A. (2006). Embodied conversational agents: “a little humor too”. *IEEE Intelligent Systems*, **21**(2), 62–64.
- Nunberg, G. (1995). Transfers of meaning. *Journal of Semantics*, **1**(12), 109–132.
- Peirsman, Y. (2006). What’s in a name? The automatic recognition of metonymical location names. In *Proceedings of the EACL-2006 Workshop on Making Sense of Sense: Bringing Psycholinguistics and Computational Linguistics Together*, pages 25–32, Trento, Italy. ACL.
- Reining, A. and Lönneker-Rodman, B. (2007). Corpus-driven metaphor harvesting. In *Proceedings of the HLT/NAACL-07 Workshop on Computational Approaches to Figurative Language*, pages 5–12, Rochester, NY.



- Schank, R. C. and Abelson, R. P. (1977). *Scripts, Plans, Goals and Understanding. An Inquiry into Human Knowledge Structures*. Lawrence Erlbaum Associates, Hillsdale, NJ.
- Schütze, H. (1998). Automatic word sense discrimination. *Computational Linguistics*, **24**(1), 97–123.
- Slobin, D. I. (1997). The origins of grammaticizable notions: Beyond the individual mind. In D. I. Slobin, editor, *Expanding the Contexts*, volume 5 of *The Crosslinguistic Study of Language Acquisition*, chapter 5. Lawrence Erlbaum Associates, Mahwah, New Jersey; London.
- Steen, G. J. (2007). Finding metaphor in discourse: Pragglejazz and beyond. *Cultura, Lenguaje y Representación / Culture, Language and Representation (CLR), Revista de Estudios Culturales de la Universitat Jaume I*, **5**, 9–26.
- Stock, O. and Strapparava, C. (2005). HAHACronym: A computational humor system. In *Proceedings of the ACL Interactive Poster and Demonstration Sessions*, pages 113–116, Ann Arbor, MI.
- Talmy, L. (1988). Force dynamics in language and cognition. *Cognitive Science*, **12**, 49–100.
- Talmy, L. (1999). Spatial schematization in language. Presented at Spatial Cognition Conference, U.C. Berkeley.
- Tinholt, J. W. (2007). *Computational Humour. Utilizing cross-reference ambiguity for conversational jokes*. Master’s thesis, University of Twente, Faculty of Electrical Engineering, Mathematics and Computer Science.
- Villavicencio, A., Copestake, A., Waldron, B., and Lambeau, F. (2004). Lexical encoding of MWEs. In *Proceedings of the Second ACL Workshop on Multiword Expressions: Integrating Processing*, pages 80–87, Barcelona, Spain.
- Warren, B. (2002). An alternative account of the interpretation of referential metonymy and metaphor. In R. Dirven and R. Pörings, editors, *Metaphor and Metonymy in Comparison and Contrast*, pages 113–130. Mouton de Gruyter, Berlin.
- Wilensky, R. (1983). *Planning and Understanding*. Addison-Wesley, Cambridge, MA.