

The Cat and the Brocaded Bag: Using Metaphor Analysis to Computationally Process Creatively Modified Idioms

Sylvia Weber Russell¹, Ingrid Fischer² and Ricarda Dormeyer³

¹ Department of Computer Science, University of New Hampshire
Durham, NH 03824, USA

² ALTANA Chair of Bioinformatics and Information Mining, University of Konstanz, Box
M712, 78457 Konstanz, Germany

³ imbus AG, Kleinseebacher Str. 9, 91096 Möhrendorf, Germany

Abstract. Theories and computational models of natural language understanding that handle idioms generally circumvent the question of novel modifications to idioms. Yet such variations are prevalent in the media. This paper addresses the perhaps most challenging type of idiom variation, i.e., variation of decomposable idioms through nontrivial metaphoric modifications in the source domain, i.e., the domain of the words of the idiom in their literal senses. An existing metaphor representation system is used as a basis for interpreting such idioms.

1 Introduction

“...from p. 269 on, she *lets the ideological cat out of the brocaded bag of prose*.”⁴ This expression is characterized by rich metaphoric modification, especially through extension of ‘brocaded’ from its source domain as an attribute of the (physical) ‘bag’ to its target domain as an attribute of the (mental) ‘prose.’ Many such metaphorically based variations can be found in ordinary prosaic texts:

- (1) starts to feel *the upwind in the sails* of his customers
- (2) *blows fresh wind into the slack sails* of the church
- (3) The cat hopped completely out of the bag
- (4) The education director did not yet *let the cat out of the bag*,
but [*did let out*] *at least the head*
- (5) He *lets the cat out of the bag*, which is supposed to eat the bad rats
- (6) He *lets the cat at least peek out of the bag*
- (7) The *cat unfortunately escaped out of the bag*
- (8) However, the public only got to *see the tail [of the cat]*
- (9) He *let the cat partly out of the Christmas grab bag*
- (10) They are always *rubbing salt and pepper into their open wounds*

⁴ The examples in this section are (translations) taken from the COSMAS corpus at the University of Mannheim, Germany www.idsmanheim.de/kt/cosmas.html.

These metaphorically based, decomposable idiom variations⁵ are less easily analyzed than some other variations of the same class, such as: variations of verbs involving an unmodified inference or effect (*The cat is out of the bag* [1–4]); target domain modifications of nouns, which need no analysis as word sense extensions (...*federally-funded wing* [5]); quantification and other modifications seen as applicable to both source and target domains (*he stirred up three hornets' nests*); insertion of a temporal verb (*he often stirred up a hornet's nest*); and purely syntactic changes (*a hornet's nest was stirred up*). (See [6] for a theoretical and computational treatment of syntactic flexibility.) These do not provide the richness of the above examples.

“Creative” is of course not a well-defined term in the interdisciplinary literature. More narrowly characterized, the above examples can be considered to be instances of novel, semantically transparent (retroactively or not), idiom-internal, semantic source-domain variations of decomposable, analyzable, metaphorically motivated idioms (with or without “word play”). How is a computer program to understand such creative variations?

This paper builds on the implemented method of Dormeyer, Fischer and Russell [4] in proposing a systematic way of computationally processing this arguably most interesting type of variation, as represented by our introductory example. We use an approach that we find necessary for any kind of novel metaphor, namely metaphor analysis based on the semantics of the metaphorically used word(s). We make no claims about human processing of idioms here, which in any case does not yet have a strong consensus [7, 8]. However, interdisciplinary research has shown that many idioms are compositional and analyzable, with metaphoric links between source and target domains playing a role in their human interpretation [9–11]. The metaphor-analytic approach we use is also consistent with psycholinguistic evidence that the literal meanings of at least some of the words of a decomposable idiom and of its variations play a role in its interpretation [12, 13, 6].

The following section summarizes the metaphor analysis method used. In the succeeding section, the computational interpretation process, with a focus on the relevant representations, is outlined for several examples of source domain modifications. This is followed by an analysis of our introductory example, based on our existing set of metaphor representation components.

2 Metaphor Interpretation

A linguistic metaphor is a linguistic expression in which one or more concepts underlying a piece of text is extended from one semantic context (source domain), to another (target domain), on the basis of some perceived or imposed similarity between the source and target expressions. The task of a text understanding system that encounters a metaphor consists of detecting the metaphor and finding an interpretation in terms of the target domain.

The approach used here for interpretation of metaphoric variations of idioms is that of **MAP** [14, 15]. This approach is consistent with our assumption that a metaphoric

⁵ The designation “decomposable.” here does not exclude the idea that some parts of the idiom have meaning only for the idiom as a whole.

usage of a word (at least at some point in its evolution) is based on semantic knowledge of the word in its so-called literal or source domain sense. Only the idiom-relevant ideas of the metaphor analysis are presented here. For a particular literal sense of a verb or of an adjective of a certain class, the working assumption is that certain components of its underlying concept are always extended to (i.e., present in) a metaphoric interpretation. (This does not necessarily imply that humans continue to make these extensions from the literal meaning for partially or completely “frozen” usages. However, our computational process can and does treat frozen usages as extensions “from scratch.”) Representing this knowledge in the lexicon in terms of extensible components is therefore critical for the computer interpretation of metaphoric text.

The determination of what can be metaphorically extended from the action or state represented by a verb is guided by the hypothesis that a purpose of metaphor is to call attention to an effect on someone. Recognition of this effect can be represented as an inference, implication or result proceeding from the verbal concept. For example, for the verb ‘to sow,’ the effect is that its object is in another location:

(11) ...*a fertile intellectual soil* in which Freud sowed his ideas. [16]

In keeping with this consideration, the extensible part of a verb definition includes the following parts relevant to our discussion:

1. An abstract structure that separates causation (if relevant) from an effected (resultant) state.
2. Embedded (evaluative or emotional) effects on particular entities, whether part of the definition of an attribute, such as NEGATIVE for pain, or a subjective connotation, such as POSITIVE for ‘breeze’ and NEGATIVE for ‘draft.’
3. A small set of “conceptual⁶ features” applicable to both literal and metaphorical usages.

These “abstract” components are those which apply to both the source and target domains. We also need to indicate the nature of the target domain and, if a target domain paraphrase is desired, how to interpret these components in the target domain. To implement this observation, a set of four (nonextensible) high-level conceptual domains: PHYSICAL, MENTAL, SENSORY, and CONTROL, is recognized, each with a further sub-level of categorization. The objects of these domains are nouns such as *bag*, *story*, *sight*, and *privilege*, respectively. Verbs and adjectives in these domains are illustrated by *to sweep*, *small*; *to think*, *intelligent*; *to watch*, *beautiful*; and *to buy*, *wealthy*.

Given such representations in the lexicon, “understanding” a phrase that contains a metaphorically used verb involves retrieving the representation of the verb in terms of the above components, together with the target domain, which is determined by the object noun of the phrase or simple sentence in which it occurs. In example (11), the metaphorically used verb *sow* occurs in conjunction with the literally used direct object ‘ideas’; the metaphoric use is recognized through the difference in domains for the verb and direct object, namely, PHYSICAL vs. MENTAL–INTELLECTUAL. The difference in conceptual domain indicated by ‘intellectual’ triggers a metaphorical interpretation in

⁶ in a sense similar to that of Schank [17]

the domain of the object. To support this kind of processing, each noun must be defined in the lexicon in terms of its conceptual domain, and each verb with its domain in its literal sense. As ‘sow’ is literally ‘to cause things to be in another location,’ a *minimal* literal paraphrase of (11) resulting from transfer of the abstract structure of ‘sow’ to the target domain is ‘Freud caused his ideas to be in the minds of others.’

While previous knowledge of metaphors may be built into the system (Martin [18]) the described approach need not rely on such knowledge; rather, like humans trying to understand a metaphor they have not previously encountered, it relies only on knowledge of definitions which can be retrieved from the lexicon. Thus there are no lexicalized metaphors in this system; conventional and novel metaphors are treated the same way.

3 Idiom Interpretation

Recognition of idioms has received much interdisciplinary attention [7]. The present discussion focuses on the problem of *interpretation* of metaphoric expressions already hypothesized as idioms.

In applying our method of metaphor processing to metaphorically modified idioms, we have somewhat of a head start. Given the basic idiom, for example, we can identify key words of the idiom. (Pulman [2] discusses the extent to which the presence of certain lexical items points to a particular idiom.) For *he let the cat out of the bag*, ‘cat’ always represents a ‘secret’ (which would not of course be discernible independently of the idiom). Furthermore, the target domain is always a mental one, regardless of the nature of the secret. In the lexicon, each idiomatic entry can therefore be provided with both its literal source domain and its idiomatic target domain, as well as the mapping between them. Moreover, we know what the general interpretation of the basic form of an idiom is. Therefore we can enter any inference which generates this interpretation into the extensible representation of the idiom. As our basic idiom is decomposable, parts of it may be modified to vary the “picture” called up by the basic idiom. The program can then base its interpretation of variations of idioms on metaphoric extension from this inference. We need not deal with the multitude of inferences which might proceed from this sentence if intended literally. For example, the relevant inference of *let the cat out of the bag*, which generates the interpretation that a secret is known, is that if the cat is out of the bag, people can see it. This allows us to implement isomorphic correspondences between states underlying the literal verb representations with inference states that can then be expressed metaphorically (illustrated in Fig. 1).

In our previous work [4], feature structure representations for *let the cat out of the bag* and *rub salt into someone’s [weeping] wounds* were used in an implementation of idiom modeling. In this paper the representational framework of that approach is used to show how interpretation of various richer source domain modifications of *let the cat out of the bag* can be processed in a non-*ad hoc* fashion. The more “creative” usages will be those that need interpretation of novel metaphor.

Because of the perceived analogy between literal and metaphoric interpretations of idioms, the meaning representations of idioms and their components are the most critical part of the interpretation process, whatever notation is used, and are the focus

Predication

```

STRUCTURE
[ CAUSE (EVENT STATEb STATEe)
%Idiom event analyzed as causing change of state %
EVENT ((AGENT +) (VOL +| -))
% if (animate) AGENT, then VOLUNTARY or not %
OR (AGENT -)
TIME: STATEb < EVENT < STATEe
% Result: %
[ DOMAIN: PHYSICAL
STATEb: AT:in (OBJ:cat1 LOC:bag1)
STATEe: NOT AT (OBJ:cat1 LOC:bag1) ] ]

```

Inference

```

STRUCTURE
[ % Idiomatically interpreted result %
SOURCE DOMAIN: SENSORY-SIGHT
TARGET DOMAIN: MENTAL-INTELLECTUAL
STATEb: AT (OBJ LOC:AGENT) & NOT (STATEe)
STATEe: AT (OBJ LOC:others)
% In the SENSORY-SIGHT domain, OBJ is a view
and LOC is the sight faculty of the indicated human;
in the MENTAL-INTELLECTUAL domain,
OBJ is intellectual (what is thought or known)
and LOC is the intellectual faculty (mind) % ] ]

```

Fig. 1. The representation of *let the cat out of the bag*.

of our interpretation approach here. The representations are in Russell's abstract representation language [14, 15], adapted for readability. In these representations, verbs are represented spatially in terms of OBJECTS, LOCATIONS and, if relevant, AGENTS. The DOMAIN indication applies only to the resulting STATE STRUCTURE rather than to the CAUSE component or to the predication as a whole. Uppercase (primitive) words are representation elements; lowercase words are instantiations. STATEb is "beginning state" and STATEe is "end state." Temporal relationships between components are indicated (TIME) through a notation associated with Reichenbach [19], where < represents "temporally before" and "EVENT" represents the time of the represented event. The character % indicates comments. Quantification is ignored for the present purpose. In the following subsections examples are given.

3.1 Common Variations

An entry for an idiom in the lexicon has two parts: its **representation** (an "abstract" version of elements characterizing that idiom) and its **interpretation**, which involves substitutions. Before proceeding to idiom modifications for which metaphor analysis is

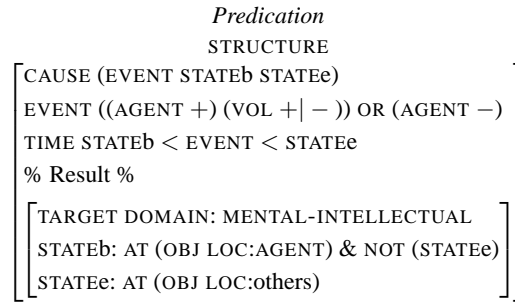


Fig. 2. Interpretation of *let the cat out of the bag*.

indicated, we look at the entry for an unmodified idiom and consider some of the “less creative” variations based on this entry. Our running example is *let the cat out of the bag*. Its representation is given in Fig. 1.

As with the inference postulates of Pulman [2], it is the inferred STATES of the basic idiom representation which anchor the interpretation of modifications. For example, the ‘cat’ being ‘out of the bag’ is linked to a fact being public (known to others), and the ‘cat’ being ‘in the bag’ is linked to a fact being secret (not known to others). Fig. 1 can be used for both the basic idiom and for the modification consisting of extraction of the “result” part of an idiom based on an agentive verb, i.e., *the cat is out of the bag* as derived from *let the cat out of the bag* [4]. Result negation (*the cat is [still] in the bag*) is also relatively simple. Both of these modifications are in contrast to modifications that deviate from the representation STRUCTURE, which are handled flexibly, i.e., according to their semantic extensions. Thus if there is a match between the parse of the hypothesized idiom into its logical form (abstract representation) and either the entire *Predication* STRUCTURE or its result portion, then the (abstract) interpretation consists of the *Predication* STRUCTURE, with the *Inference* STRUCTURE replacing the result part of the *Predication* STRUCTURE without the source domain (see Fig. 2). The corresponding paraphrase of *let the cat out of the bag* (the “effect” part of the translation is derived from the lexicon search for a representation corresponding to STATE_e) is: **‘do something such that the mental (intellectual) object that was not known to others becomes known to others.’**⁷ *[the cat] be out of the bag* is paraphrased as: **‘be effected that the intellectual object that was not known to others becomes known to others.’**

For *the cat is [still] in the bag*, the resulting STATES take the following form:

(12) STATE_b: AT:in (OBJ:cat1 LOC:bag1)
STATE_e: STATE_b

⁷ For readability, this paraphrase and the representation on which it is based do not include the *intention* that the proposition be kept secret, or the possibility that the secret may be shared. Of relevance to the possibility of “complete” interpretations, however, is Gibbs’ (1993) reference to experiments demonstrating that idioms cannot be expected to be equivalent to their paraphrases, at least not as simple predications.

STATE_e here matches STATE_b of the basic idiom *Predication* and therefore takes the interpretation of STATE_b in Fig.2. This STATE interpretation is then inserted into the STATE configuration of Fig.3. The entire interpretation of *to put the cat back into the bag* is shown in Fig.3 on the right. The paraphrase is then: **‘do something such that the intellectual object that became known to others became not known to others, as it was originally.’**

The preceding modifications were “logical” in that they had to do with the presence, negation or reversal of states. More subtle modifications add components to the verb but leave the logical result the same. In order to interpret novel modifications of this kind in a way that makes sense, extensible components must be defined in a way that is general for both the source and the target domains, as described in the above section on metaphor. For example, in the case of *the cat hopped out of the bag*, the abstract definition of ‘hop’ would include (beside the structure representing the change of location) components representing NON-CONTINUOUS and a high SPEED. (A detailed literal description of ‘hop’ in the lexicon - including, e.g., ‘to propel oneself upward and land again,’ while useful for other purposes, would not necessarily lead to a meaningful interpretation.) Using such abstract components as a guide through a discrimination net, the program arrives at either a target domain word with the same components, if there is one, or simply uses the target domain translation of the components themselves - in this case, *the cat came out of the bag suddenly*, i.e., ‘the proposition that was not known to others suddenly could become known to others.’

3.3 Adjectival Source Domain Variations - The Brocaded Bag

The idiom variation, *she lets the ideological cat out of the brocaded bag of prose*, rests on an analogical form, where the target domain concepts (‘ideology’ and ‘prose’) are explicitly given. In accordance with the general rule for metaphoric analogical expressions, ‘X of Y \rightarrow X=Y’ (cf. ‘water of life’), ‘bag’ is mapped to the (MENTAL-INTELLECTUAL) ‘prose’ of the AGENT, therefore becoming the LOC of the ‘cat/ideology.’ Interpreting this variation involves an analysis of a clearly novel metaphor. The revelation of a secret, as usual, is inferred from the basic structure and key words of the input phrase. From the basic idiom representation, the ‘ideological *cat*’ is revealed to the public. The question is, what does ‘brocaded’ mean in the target domain?

The past participial adjective ‘brocaded’ represents a conceptual *attribute* intended to apply to the (target domain) MENTAL OBJECT, i.e., ‘prose,’ that a human relates to in the above STATES. We need to show the extension of the literal sense of ‘brocaded’ to its metaphoric sense as a descriptor of the ‘prose.’ Adjectives representing conceptual attributes, i.e., simple properties of nominals, are defined abstractly in the lexicon in terms of extensible conceptual features. For the PHYSICAL-domain ‘brocaded,’ one feature is ‘(COMPLEX +).’ Another corresponds to the word ‘beautiful,’ which has the feature value ‘(EVALUATION POSITIVE).’ One might argue that these representation components, which also apply to ‘fancy’ and other adjectives, are overly general. However, Aarts and Calbert [20] found experimental support for the thesis that only evaluation and intensity are general candidates for extension from adjectives; Osgood [21] proposed evaluation, potency (intensity) and activity (cf. a COMPLEX pattern, sometimes referred to as a busy pattern). Moreover, our feature values distinguish such adjectives

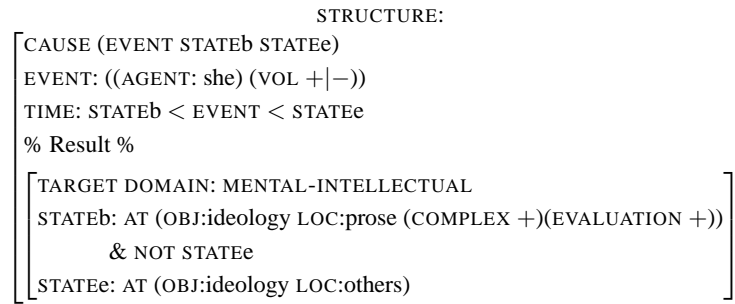


Fig. 4. Interpretation of *She let the ideological cat out of her brocaded bag of prose.*

from those that, for example, are ‘simple’ (COMPLEX –), such as ‘red,’ and ‘messy’ (EVALUATION –), such as ‘blotchy.’

Interpretation, then, consists of applying these feature values of ‘brocaded’ to the nominal that the adjective modifies (‘prose’), in the MENTAL–INTELLECTUAL domain. If these components are integrated into the basic idiom representation for *let the cat out of the bag*, we have for *She let the ideological cat out of her brocaded bag of prose* the interpretation in Fig. 4, giving the paraphrase: ‘**She did something such that the ideology in her beautiful, complex prose becomes known to others.**’ It might be noted that this idiom variation incorporates more semantic complexity than the usual metaphoric extension from *let the cat out of the bag*. Usually, a specific utterance or action causes the revelation; here, the “revelation” is a potential judgment or recognition made in the course of reading. This recognition is “as if” a secret were being revealed, perhaps requiring another layer of analysis for a full interpretation. Neither is the layer of irony perceptible in this expression addressed. However, the metaphoric elements are basically accounted for.

4 Other Computational Research on Metaphor

The metaphoric basis of our method invites comparison with other computational work on metaphor. In his “abduction” approach to metaphor, Hobbs [22] proves the *coherence* of a “Congress” schema and a “baseball” schema that provide the background of a “novel” metaphoric *veto-hit* usage. This usage is certainly creative, but Hobbs’ approach does not address the interpretation of a metaphor that is novel *in the sense that* no schema or conventional metaphor is available.

Other approaches, e.g., that of Martin [18] as cited earlier, rely on references to known, Lakovian conventional metaphors, such as *IDEAS ARE FOOD*. These approaches, which exploit, as we do, the invariance of certain basic components such as change of state, are productive for a large range of metaphoric expressions. However, as in the case of Hobbs’ theory, they would not succeed for the less prevalent but frequently encountered metaphors for which an underlying conventional metaphor or schema cannot be identified.

In this respect the work of Narayanan [23] and the more recent work of Barnden et al. [24] is of interest. Narayanan’s implemented use of various crosslinguistic mappings

and Barnden et al.'s addition of "view neutral mapping adjuncts" to their metaphor reasoning system are theoretically similar to the extensible abstract structures and features used in the MAP system to define verbs (and predications underlying nouns) in physical and nonphysical domains. The application of these mappings to idioms is discussed by Fischer and Chang [25] based on variations of German idioms. Narayanan's existing mappings do not appear to be applicable to attributes such as *brocaded*, while Barnden et al.'s give some attention to the necessary evaluative connotations; in any case, our three approaches appear to at least partially converge. The latter researchers [26] also provide techniques that deal with the "interaction" theory of metaphor and target-to-source influence, as they use contextual information from the target domain to highlight or confirm properties hypothesized as extended from the source domain. Presumably these methods could help to *confirm* extensions made from, e.g., *brocaded* above, within or outside the context of idiom variations.

Our next task in handling source domain adjectives as idiom variations will focus on a pilot implementation with a small but diverse vocabulary, with close attention to generality. For now, our idiom-variation paraphrase system is just "peeking out of the bag."

References

1. Erbach, G., Krenn, B.: Idioms and support verb constructions. In Nerbonne, J., Netter, K., Pollard, C., eds.: *German in Head-Driven Phrase Structure Grammar*. Number 46 in Lecture Notes. CSLI Publications, Stanford (1994) 365–396
2. Pulman, S.: The recognition and interpretation of idioms. [7]
3. Cacciari, C.: The place of idioms in a literal and metaphorical world. [7] 27–56
4. Dormeyer, R., Fischer, I., Russell, S.W.: A lexicon for metaphors and idioms. In Schnorbusch, D., Langer, S., eds.: *Semantik im Lexikon*. Tübinger Beiträge zur Linguistik. Gunter Narr Verlag, Tübingen (2005) 205–224
5. Ernst, T.: Grist for the linguistic mill: Idioms and "extra" adjectives. *Journal of Linguistic Research* (1981)
6. Stock, O., Slack, J., Ortony, A.: Building castles in the air: Some computational and theoretical issues in idiom comprehension. [7]
7. Cacciari, C., Tabossi, P., eds.: *Idioms: Processing, Structure, and Interpretation*. Lawrence Erlbaum, Hillsdale (1993)
8. Everaerts, M., van der Linden, J.E., Schenk, A., Schreuder, R., eds.: *Idioms: Structural and Psychological Perspectives*. Lawrence Erlbaum, Hillsdale (1995)
9. Nayak, N., Gibbs, R.W.: Conceptual knowledge in the interpretation of idioms. *Journal of Experimental Psychology: General* **119** (1990) 315–330
10. W.Gibbs, R.: Idiomaticity and human cognition. [8] 97–116
11. Geeraerts, D.: Specialization and reinterpretation in idioms. [8] 5776
12. Nunberg, G., Sag, I.A., Wasow, T.: Idioms. *Language* **70** (1994) 491–538
13. Cacciari, C., Glucksberg, S.: Understanding idiomatic expressions: The contribution of word meanings. In Simpson, G.B., ed.: *Understanding Word and Sentence*. Elsevier, Amsterdam (1990) 217–240
14. Russell, S.W.: Information and experience in metaphor: A perspective from computer analysis. *Metaphor and Symbolic Activity* **1** (1986)
15. Russell, S.W.: Metaphoric coherence: Distinguishing verbal metaphor from 'anomaly'. *Computational Intelligence* **8** (1992)

16. Samson, M.: German-American Dialogues and the Modern Movement before the "Design Migration," 1910-1933. PhD thesis, Harvard (1988)
17. Schank, R.: Conceptual Information Processing. North Holland, Amsterdam (1975)
18. Martin, J.: A Computational Model of Metaphor Interpretation. New York edn. Academic Press (1990)
19. Reichenbach, H.: Elements of Symbolic Logic. Macmillan, London (1947)
20. Aarts, J., Calbert, J.: Metaphor and Non-metaphor: The semantics of adjectivenoun combinations. Max Niemeyer Verlag, Tübingen (1979)
21. Osgood, C.: The cognitive dynamics of synesthesia and metaphor. In Honeck, R., Hoffman, R., eds.: Cognition and figurative language. Lawrence Erlbaum (1980) 203–238
22. Hobbs, J.: Metaphor and abduction. In Ortony, A., Slack, J., Stock, O., eds.: Communication from an Artificial Intelligence Perspective: Theoretical and Applied Issues. Springer, Berlin (1992) 35–58
23. Narayanan, S.: Moving right along: A computational model of metaphoric reasoning about events. In: AAAI. (1999)
24. Barnden, J., Glasbey, S., Lee, M., Wallington, A.M.: Domain-transcending mappings in a system for metaphorical reasoning. In: EACL. (2003) 57–61
25. Chang, N., Fischer, I.: Understanding idioms. In Zhlke, W., Schukat-Talamazzini, E., eds.: Proc. 5. Konferenz zur Verarbeitung natürlicher Sprache - ITG-Fachtagung "SPRACHKOMMUNIKATION", Berlin, VDE-Verlag (2000) 33–38
26. Barnden, J., Glasbey, S., Lee, M., Wallington, A.M.: Varieties and directions of inter-domain influence in metaphor. *Metaphor and Symbol* **19** (2004) 1–30

