


Grounding Figurative Language Use in Incompatible Ontological Categorizations

Katja Markert

Language Technology Group
HCRC, Edinburgh University
Edinburgh EH8 9LW, U.K.
markert@cogsci.ed.ac.uk

Udo Hahn

 Text Understanding Lab
Freiburg University
D-79085 Freiburg, Germany
<http://www.coling.uni-freiburg.de>

Abstract

We propose a formal criterion for delineating literal from figurative speech (metonymies, metaphors, etc.). It is centered around the notion of categorization conflicts that follow from the context of the utterance. In addition, we consider the problem of granularity, which is posed by the dependence of our approach on the underlying ontology, and compare our distinction with alternative reference-based explanations.

Introduction

Figurative language use comes in different varieties (e.g., as metonymy in example (2) and as metaphor in example (3) below), and is typically contrasted with literal language use (e.g., example (1)) on the basis of some notion of deviance.

- (1) “*The man left without paying.*”
- (2) “*The ham sandwich left without paying.*”
- (3) “*The Internet is a gold mine.*”

Cognitive linguists have been struggling for decades to draw a proper distinction between literal and figurative utterances. Their interest derives from the question how a basic, lexical meaning representation must be conceived from which figurative (and possibly literal) readings can be derived. This viewpoint implies to assume a computational process and, hence, requires to be explicit about the representational foundations from which to proceed.

Currently, two approaches prevail, which spell out this distinction. The first one, e.g., Lakoff & Johnson (1980), simply regards *deviation from literal reference* as a sufficient condition for figurativeness. No formal criteria for the nature of such a deviation are given so that the discrimination of literal and figurative meaning rests on subjective ascription.

The second approach (Fass, 1991; Pustejovsky, 1991; Stalard, 1993) introduces such a formal criterion. Each time *selectional restrictions* are violated, e.g., through type conflicts, an instance of figurative speech is encountered. Special reasoning patterns are then activated, like type coercion for metonymies (Pustejovsky, 1991) or analogy-based structure mapping for metaphors (Carbonell, 1982; Gentner et al., 1989), in order to cope with the triggering instance such that a reasonable interpretation can be derived, one that no longer violates the underlying constraints. The proponents of this approach present a lot of supporting evidence for their methodological claims (cf. example (4)) but obviously fail to

cover a wide range of residual phenomena (example (5) lacks any violation of selectional restrictions though being figurative, at least, assuming the “writings-by-Chaucer” reading):

- (4) “*I read Chaucer.*”
- (5) “*I like Chaucer.*”

In this paper, we aim at providing a formal framework from which a proper distinction between literal and figurative language use can be made. Rather than formalizing the notion of deviation with recurrence to selectional restrictions, we will base our distinction on conceptual criteria that incorporate the influence from the context of an utterance. These criteria allow us further to focus on the dependence of literal and figurative speech on individual ontologies. Considering granularity issues of ontologies we may even overcome the influence of subjectivity by taking additional formal criteria into account.

Lexical Meaning

We will base our considerations on the notion of *context-independent lexical meaning* of lexemes, from which the notions of literal and figurative meaning in context will be derived. Lexical meaning will be a function from lexemes to categories (concepts) of an ontology.

So, let \mathcal{L} be the set of *lexemes* of a given natural language and let $\mathcal{L}' \subset \mathcal{L}$ be the subset of lexemes containing nouns, main verbs and adjectives only (e.g., *man* or *policeman* are elements of \mathcal{L}'). We also assume an ontology composed of a set of concept types $\mathcal{F} := \{\text{MAN, POLICEMAN, HAM-SANDWICH, ...}\}$, a set of instances $\mathcal{I} := \{\text{man-1, policeman-2, ...}\}$ related to concept types, and a set of relations $\mathcal{R} := \{\text{has-part, part-of, agent, ...}\}$, which link concept types or instances. We take a set theoretical semantics for granted as is commonly assumed in description logics (Woods & Schmolze, 1992). The *lexical meaning* β_{lex} can then be defined as a relation $\beta_{lex} \subset \mathcal{L}' \times \{\mathcal{F} \cup \mathcal{I}\}$. While we refrain from considering the linkage between lexemes and ontological entities in depth (cf., e.g., Cruse (1986) or Jackendoff (1990)), we require the relation β_{lex} to fulfill the following properties:

1. If *lexeme* $\in \mathcal{L}'$ is a proper name, then a unique lexeme .i $\in \mathcal{F} \cup \mathcal{I}$ with $(\text{lexeme}, \text{lexeme} .i) \in \beta_{lex}$ exists such that $\text{lexeme} .i \in \mathcal{I}$. Thus, every proper name is linked to a single instance in the domain knowledge base.

2. If $lexeme \in \mathcal{L}'$ is not a proper name, then a concept $lexeme.CON \in \mathcal{F}$ must exist so that $(lexeme, lexeme.CON) \in \beta_{lex}$. Also, no instance $lexeme.i \in \mathcal{I}$ exists such that $(lexeme, lexeme.i) \in \beta_{lex}$.
3. For reasons of simplicity, we will now restrict β_{lex} appropriately. If $lexeme \in \mathcal{L}'$ is not a proper name, then we require for all $i \in \beta_{lex}(lexeme)$ that i can be referred to by $lexeme$ in a *context-independent* way. Hence, we assume that reference to any i via $lexeme$ is always possible. (We cannot, e.g., relate the lexeme *fool* to MAN as not every man can be referenced by *fool* independent of the context.) The condition of context-independence may, however, still hold for several concepts that stand in a subsumption relation to each other. So, when we regard the lexeme *man*, this condition holds for both the concepts MAN and POLICEMAN, as all $i \in POLICEMAN$ and all $i \in MAN$ can be referenced by *man*. We then regard the most general concept to which this unconditioned reference relation applies (here, MAN) as the lexical meaning, and, in general, consider *lexical meaning* as a function¹ $\beta_{lex} : \mathcal{L}' \mapsto \mathcal{F} \cup \mathcal{I}$. By convention, we denote $\beta_{lex}(lexeme)$ by $lexeme.CON$.

Lexical meaning is thus considered as a context-independent function from lexemes to categories (concepts) of an ontology. As there is no agreement on canonical ontologies, this mapping introduces subjective conceptualizations.

Finally, we extend our definition to words w of a discourse so that their corresponding lexeme be $w.lex \in \mathcal{L}'$. We simply assume $\beta_{lex}(w) := \beta_{lex}(w.lex)$. We distinguish the range of that mapping by $w.i$ for proper names and $w.CON$ in all other cases. Hence, the lexical meaning of the word “*man*” in example (1) is given by the concept MAN.²

Literal vs. Figurative Meaning

While in the previous section we have been dealing with the isolated lexical meaning of a word only, in this section we will incorporate the *context of an utterance* in which a word appears. Hence (cf. Fig. 1), we introduce the word w' with respect to which word w is syntactically related — w' is either head or modifier of w . Such a dependency relation (either a direct one or a well-defined series of dependency relations) at the linguistic level induces a corresponding conceptual relation $r \in \mathcal{R}$ at the ontological level (Romacker et al., 1999). The conceptual relation r links the conceptual correlates, $w.sf$ and $w'.sf$, of w and w' , respectively. Accordingly, we may now say that w *StandsFor* a corresponding domain entity $w.sf$; alternatively, $w.sf$ is called the (intended) meaning of w . The comparison of $w.sf$ with $w.CON$ or $w.i$ lies at the heart of the decision criterion we

¹In order to make β_{lex} a function we assume in the case of *polysemy* one of several meaning alternatives to be the primary one from which the others can be derived. In the case of *homonymy*, we assume the existence of different lexemes which can be mapped directly to mutually exclusive concepts.

²The lexical meaning of a word w must be distinguished from the concrete referent of w in the given discourse.

propose for judging whether a reading is literal or figurative. So, in the well-known example (2), “*ham sandwich*” ($= w$) *StandsFor* “the man who ordered the ham sandwich” ($= w.sf$), which is distinct from its lexical meaning, HAM-SANDWICH ($= w.CON$).




Figure 1: Framework for Contextual Interpretation

We may now consider some examples to distinguish several cases how $w.sf$ can be related to $w.CON$ or $w.i$. This will also lead us to clarify the notion of distinctiveness between the items involved. Let $w.sf$ be an instance from \mathcal{I} , and let $w.CON_{sf}$ be the least general concept such that $w.sf \in w.CON_{sf}$.³ This assumption will be shortcut as $w.sf$ *inst-of* $w.CON_{sf}$.

In the simplest case, $w.sf$ and $w.CON / w.i$ are related by an *inst-of* relation. Then $w.CON_{sf} = w.CON$ holds. In the utterance

- (6) “A *man* left without paying.”

we have $w = \text{“man”}$ and $w' = \text{“left”}$. Furthermore, $w.sf = \text{man-1 inst-of MAN} = w.CON = w.CON_{sf}$. So, in the example (6), lexical meaning and actual meaning coincide.

If we consider all relations other than equality as deviant, we characterize a class of phenomena that is certainly larger than the one containing figurative speech only. Example

- (7) “A *policeman* left without paying. *The man* lost his job.”

illustrates an anaphoric, non-figurative relation between “*man*” ($= w$) and “*policeman*”. A subsumption relation holds between $w.CON_{sf}$ ($= POLICEMAN$) and $w.CON$ ($= MAN$), which means that $w.CON$ is either more general than or equal to $w.CON_{sf}$. In particular, we have ($\text{policeman-1} =$) $w.sf \in w.CON$, but not $w.sf$ *inst-of* $w.CON$, in general (as in example (6)).

Loosening ties a bit more, we may abandon the subsumption relation between $w.CON_{sf}$ and $w.CON$ as in example

- (8) “A *policeman* left without paying. *The fool* lost his job.”

We have ($\text{policeman-1} =$) $w.sf \in w.CON$ ($= FOOL$), but the specialization relation between $w.CON_{sf}$ ($= POLICEMAN$) and $w.CON$ ($= FOOL$) no longer holds. Instead, we are set back to $w.sf \in w.CON_{sf} \cap w.CON$ and, therefore, $w.CON_{sf} \cap w.CON \neq \emptyset$. We say that $w.CON_{sf}$ and $w.CON$

³The least general concept $w.CON_{sf}$ with $w.sf \in w.CON_{sf}$ is the intersection of all concepts $C \in \mathcal{F}$ with $w.sf \in C$.

are *compatible*, as no categorization conflict arises. This also holds for all previously discussed examples. As a consequence, the notion of *categorization conflict* turns out to become crucial for our distinction between literalness and figurativeness — the latter being based on an underlying categorization conflict, whereas the former is not. We summarize these observations in the following definition:

Definition 1 (Literalness via Syntactic Constraints)

A word w in an utterance U is used according to its *literal meaning*, if for every instance $w.sf \in \mathcal{I}$ which w StandsFor, one of the following two conditions hold:

$$w.sf = w.i \quad \text{if } w \text{ is a proper name} \quad (1)$$

$$w.sf \in w.CON \quad \text{else} \quad (2)$$

Especially, $w.CON_{sf} \cap w.CON \neq \emptyset$ holds for non-proper nouns.

We here restrict the notion of figurative speech to those relationships between $w.sf$ and the lexical meaning of w in terms of $w.CON$, which are not inclusive ones. A literal use of the word w for an instance $w.sf$ *inst-of* $w.CON_{sf}$ is only possible, if $w.CON_{sf} \cap w.CON \neq \emptyset$. If, however, a categorization conflict occurs, i.e., $w.CON_{sf} \cap w.CON = \emptyset$, then we call the use of w *figurative* (as illustrated by “ham sandwich” in example (2) or by “gold mine” in example (3)). We would like to stress the following implications:

1. We can determine exactly the place where *subjectivity* comes in when a distinction between literalness and figurativeness is made — it is mirrored by subjectivity in categorization. “fool” in example (8) can only be considered as literal, if the concepts FOOL and POLICEMAN are considered as being compatible (in the set theoretic sense introduced above). If one does not share this conceptualization, this usage of “fool” must be considered as figurative (or even absurd). Thus, we capture the subjectivity of figurativeness formally in the ontological premises, not via intuitive considerations.
2. Definition 1 does not depend on the *violation of selectional restrictions*. The example (5) (“I like Chaucer.”) allows for the same analysis as example (4) (“I read Chaucer.”), because the intended patient of *like* are, in both cases, Writings-by-Chaucer (= $w.sf$), although this is not indicated by selectional restrictions at all. In both cases $w.CON_{sf} \cap w.CON = \emptyset$, i.e., figurativeness holds.

Granularity

The (non-)inclusion criterion we have set up for the distinction between literal and figurative speech in Definition 1 introduces a particularly strong tie to the underlying ontology. One of the problems this might cause lies in *granularity* phenomena of domain knowledge bases and their impact on literal/figurative distinctions. Given different levels of granularity, it may well happen that a word w StandsFor an instance $w.sf$ *inst-of* $w.CON_{sf}$ with $w.CON_{sf} \cap w.CON = \emptyset$, though,

intuitively, one would rate the usage of w as a literal one. Assume we have a knowledge base KB_1 in which CPU happens to be PART-OF MOTHERBOARD, while MOTHERBOARD itself turns out to be PART-OF COMPUTER. If we analyze the example

(9) “The CPU of the computer ...”

accordingly, we end up with the determination of a figurative usage for ($w =$) “computer”, since $MOTHERBOARD \cap COMPUTER = \emptyset$ (cf. Fig. 2).




Figure 2: Example (9) Assuming KB_1

If we assume, however, a representation in a knowledge base KB_2 such that CPU is an *immediate* PART-OF COMPUTER, then we derive a literal usage for w (cf. Fig. 3).




Figure 3: Example (9) Assuming KB_2

In order to lower the dependence on knowledge base granularity we may derive a weaker condition of literalness from Definition 1. Assume $w.sf$ and $w'.sf$ to be related by a conceptual relation r (technically, $w'.sf r w.sf$). Thus, for literal usage of w the following conditions hold:

$$w'.sf r w.i \quad \text{if } w \text{ is a proper name} \quad (3)$$

$$\exists i \in w.CON: w'.sf r i \quad \text{else} \quad (4)$$

(3) immediately follows from (1) in Definition 1, since $w'.sf r w.sf (= w.i)$ holds. (4) can be deduced from (2) by defining $i := w.sf$.

Since these conditions provide weaker conditions of literal language use than the ones we have agreed upon in Definition 1, all literal usages determined by the strong condition still remain literal (in particular, example (9) is considered a literal usage of “computer” given KB_2). Considering the granularity effects for example (9) with respect to

KB_1 , we may determine the literal usage of “computer” by the following consideration. Since CPU is PART-OF MOTHERBOARD, and MOTHERBOARD is PART-OF COMPUTER, we conclude with the transitivity of the PART-OF relation⁴ that CPU is PART-OF COMPUTER. Hence, criterion (4) is fulfilled (assuming $i = \text{computer-2}$, $w'.sf = \text{cpu-1}$, $w.sf = \text{motherboard-2}$, and $r = \text{PART-OF}$). Unlike the examples discussed previously, we do not have $w.sf \in w.CON$ (criterion (2) from Definition 1). So, by moving from the strict criteria in Definition 1 to the weaker ones stated by criteria (3) and (4) we are able to balance granularity phenomena of knowledge bases, to a certain extent at least.

Figurativeness and Reference

One might argue that the problem just discussed, the dependence of the distinction between literal and figurative usage on knowledge base structures, follows from the dependence of *StandsFor* on syntactic context. Accordingly, some researchers, e.g. Lakoff & Johnson (1980), have proposed to build the definition of figurative speech upon the notion of *reference*. The assumption being made is that w uniquely refers to a knowledge base item $w.ref \text{ inst-of } w.CON_{ref}$ and that figurativeness results from the deviation of this reference from literal meaning. Although their notion of deviance is not formalized, referentially-based literalness can now be defined straightforwardly in our approach by proceeding exactly along the lines of Definition 1:

Definition 2 (Literalness in the Referential Approach)

A word w is called *literal in the referential approach*, if:

$$w.ref = w.i \quad \text{if } w \text{ is a proper name} \quad (5)$$

$$w.ref \in w.CON \quad \text{else} \quad (6)$$

Without doubt, we here circumvent the granularity problem, since no change in reference occurs for example (9), no matter whether KB_1 or KB_2 is assumed.⁵ But the reference approach runs into severe problems when one considers, e.g., classical examples of metonymies such as

(10) “I like to read Chaucer. He was a great writer.”

We have $w = \text{“Chaucer”}$ as a typical example for a writer-for-writings metonymy (Lakoff & Johnson, 1980).⁶ The assumption to link literal/figurative usage to reference relations is flawed by the fact that $w = \text{“Chaucer”}$ does not refer to the “writings by Chaucer”, because in this case the referentially determined anaphor “He” could not be resolved. In particular, we have $\text{Chaucer.ref} = \text{Chaucer}$, therefore $w.ref = w.i$. Hence, “Chaucer” must be considered,

⁴We are aware of empirical observations about the transitivity of PART-WHOLE relations made by Chaffin (1992) and Winston et al. (1987), and take their constrained notion of transitivity for granted.

⁵Note that this definition is, nevertheless, still dependent on the knowledge base and on the lexical meaning of w .

⁶This example is ambiguous in several ways as, e.g., also the “style of Chaucer” could be another reasonable metonymic reading, thus giving rise to a writer-for-style metonymy. The following arguments hold for those alternative metonymic readings as well.

counterintuitively, as a literal use according to criterion (5) (similar problems have been discussed at length by Stallard (1993)). Given our context-dependent definitions ((1) or (3)), we get $w.sf = \text{Writings-by-Chaucer}$ so that $w.sf \neq \text{Chaucer}$. Thus, $w = \text{“Chaucer”}$ is analyzed figuratively in our approach.

Summarizing our discussion, we combine criteria (1) to (6) by the following conditions:

1. A word w is used in its *literal* meaning in all definitions, if $w.sf \in w.CON$ (analogously, $w.sf = w.i$) for all $w.sf$, and $w.ref \in w.CON$ hold (combining Definition 2 and 1 with respect to literal usage).
2. If $w.sf \notin w.CON$ (analogously, $w.sf \neq w.i$) for some $w.sf$ (with respect to a relation r and another word w'), but $w.ref \in w.CON$ ($w.ref = w.i$), two cases must be distinguished:
 - In cases of granularity effects criteria (4) or (3) hold. By this, an $i \in w.CON$ exists with $w'.sf r i$ (analogously, $w'.sf r w.i$). We can include this in our definition of literal usage as its analysis is only due to implications a particular ontology design brings to bear.
 - In cases of figurative speech like the one in example (10) the criteria (4) / (3) do not hold. We include these cases in our definition of figurative usage as phrased below.
3. A word w is used in its *figurative* meaning according to the syntactic and the referential definition, if $w.ref \notin w.CON$ holds and there exists a $w.sf \notin w.CON$. This is the case, e.g., in example (2).

Having only considered the figurative usage of a word w so far, we end up by defining U as a *figurative utterance*, if it contains at least one word w that is used in a figurative way.

Related Work

We consider as the main contribution of this paper the introduction of a *formal notion of deviance* that is both general and simple. To the best of our knowledge, no comparable work has been done so far on this issue. Although there exist formal characterizations of metaphors (Indurkha, 1988; Gentner et al., 1989) these studies rather account for *structural properties* of metaphors (e.g., constraints on domain mappings, aptness conditions), than they deal with the *distinction* between literal and figurative speech.

Usually, however, utterly vague characterizations of what constitutes figurative language prevail such as the famous and often cited description by Lakoff & Johnson (1980) who characterize a metonymy by the use of “one entity to refer to another that is related to it”. There is no restriction on the kind of relatedness between the objects. For example, relatedness might include class inclusion, similarity or part-whole relations. But only the latter are included in metonymy, in general, and the examples Lakoff and Johnson put forward suggest that it is this conventional kind of metonymy they

are talking about. The relation of class inclusion even leads to literal meaning as we have shown. Similar criticism applies to Tourangeau & Sternberg (1982), Fauconnier (1984), Kittay (1987), Turner (1988), Nunberg (1995), and many others. The same shadowy definitions of figurative language are then often adopted by computational linguists (Fass, 1991; Martin, 1992). This leads to the fact that it is mostly not clear at all, which phenomena are treated by these approaches.

In addition, a tendency can be observed in more formal approaches — pressed by the need to look for computationally feasible definitions of metaphor or metonymy — to consider figurative language as a violation of selectional restrictions (Carbonell, 1982; Fass, 1991; Hobbs et al., 1993; Pustejovsky, 1991) or communicative norms (Grice, 1975; Searle, 1979). Such an approach *equates* an often used triggering condition, *viz.* constraint violation, with the phenomenon of figurative language (or, subsets, like metonymies). Hence, it confuses the possible, but not necessary effects of a phenomenon with the phenomenon to be explained.

Despite the lack of formal rigor in previous work, it is worth to investigate how our formal criterion is compatible with other views on figurative speech from cognitive science. The tendency to see figurative speech rooted in conceptual categories, as we do, is becoming consensus in cognitive linguistics. The main trend is, e.g., to treat metaphors as a means of categorization by way of similarity (Gibbs, 1992) and to retrace figurative speech to cognitive procedures involving categorization and (subjective) experience (Lakoff & Johnson, 1980; Fauconnier, 1984; Lakoff, 1987). So, Lakoff and Johnson see metaphors rooted in our way of conceptualization via mappings. Kittay (1987) and Turner (1988) regard some kind of conceptual incompatibilities as the basis of metaphorization. Nevertheless, they do neither explicate their theory of categorization and incompatibility, nor do they recognize that these incompatibilities are relevant for other kinds of figurative speech, as well as for metaphors in the strict sense of the word. The dependence of lexical, literal and figurative meaning on ontologies is, therefore, realized, but no explicit formal treatment is given of particular problems this implies.

This is where we see the second major contribution of the paper. Once a formal distinction between literal and figurative meaning is given, it allows us to characterize *subjectivity*, so far an entirely informal notion, by reference to the particular ontology underlying the natural language understanding process. We aim at adapting different ontologies such that by way of abstracting away different *granularities* of representation structures (e.g., by generalizing more fine-grained representations to a coarser grain size, as in criterion (4)) disagreement might turn into consensus (e.g., considering example (9)). Contrary to that, the majority of researchers in our field of study attribute the difference in opinion to the existence of different, incompatible ontologies, and leave it with that explanation without further attempt at smoothing (Lakoff & Johnson, 1980; Turner, 1988).

An exception to this rule is the work by Veale & Keane (1994). While the authors still adopt an entirely informal *definition* of metaphors (close to the one from Turner (1988)), with all its drawbacks, Veale & Keane incorporate a concise, formal explication of how different viewpoints of different speakers influence metaphor interpretation. In contrast to our work, they offer the possibility to accept or reject beliefs in their knowledge base, depending on whether the speaker believes the propositions to be true or not. This then accounts for connotations which might arise in the metaphor interpretation. Instead, we focus on the problem of granularity, offering the possibility to derive coarse-grained views of the ontology from fine-grained views, thus not (de)activating certain propositions but viewing the same propositions in different grain sizes. This is not meant to explain metaphorical effects, but to reconcile different notions of literalness.

The third major proposal we make relates to the contextual embedding of figurative speech. The criterion we formulate is entirely based on syntactic relations that guide conceptual interpretation. In particular, and unlike most algorithmic accounts (Norvig, 1989; Fass, 1991; Pustejovsky, 1991; Hobbs et al., 1993), it does not rely at all upon the violation of selectional restrictions (for a notable exception, cf. Martin (1992)), since this criterion accounts only for a subset of the phenomena naturally recognized as figurative language. In addition, the syntax-based proposal we make avoids to consider reference changes as an indicator of figurativeness as is commonly assumed (e.g. by Lakoff & Johnson (1980)). Instead, our proposal is inspired by Fauconnier's (1984) "connector" function. Though Fauconnier aims at an embedding of figurative language into syntax, there exists no formalization of this notion in relation to an established grammar framework nor is the notion of conceptual incompatibility (and other purely conceptual issues of figurative language) formalized. A more formal criticism of the reference changes proposal was made by Stallard (1993) who, nevertheless, then only dealt with figurative language violating sortal constraints. Another notion of context is again used by Veale & Keane (1994), who do not use syntactic or selectional restriction properties for explaining effects of metaphorical speech, but rely on the speaker's belief system.

Our approach is fully compatible with viewing figurative language as *regular and not violating* linguistic norms. Whereas literal language is grounded in inclusion relations to lexical meaning, figurative language is grounded in other relations to lexical meaning. These can, nonetheless, be systematic and conventionalized relations like part-whole relations. There is no need to claim that inclusion relations are prior or are to be preferred to other (conventional) relations. This is in accordance with the conventional metaphor view first stipulated by Lakoff and Johnson. It is also in accordance with psycholinguistic research showing that figurative speech is in most cases as easily understood as literal speech. This is especially the case when the instance of figurative speech is conventional, i.e., grounded in systematic and pervasive onto-

logical relationships (Blasko & Connine, 1993). The essence of this is that pervasive and structured relations or relations made salient by the context (Inhoff et al., 1984) may be as easily available to comprehension as inclusion relations.

Conclusion

In this paper, we have drawn a distinction between literal and figurative speech which is based on *formal* criteria. These are grounded in the solid framework of description logics, in particular, by relying on its set theoretic semantics. A crucial condition of whether language use is considered literal or figurative is introduced by the particular *ontology* referred to. While earlier formal approaches appeal to semantic types, sortal constraints, etc., this is not fully convincing, since the entire structure and *granularity* of the theory of the domain being talked about contributes to the understanding process, whether literally or figuratively based. In particular, we captured the notion of subjectivity in ontological premises and explained how granularity problems may be overcome.

A *recognition procedure* for figurative language is reported in Hahn & Markert (1997). Contrary to almost all competing approaches, we do not rely on a special triggering mechanism to start figurative interpretation when the literal one has failed. Rather we compute both interpretations in parallel, i.e., without preference for literal interpretations. The distinction between both forms of interpretation (and the need for a corresponding criterion) comes in, finally, when the text understander is required to disambiguate between competing readings. Among the preference criteria we apply are the distinction whether a reading is literal (preferred) or figurative.

The model we have presented does currently not account for neologisms, as those have no *a priori* lexical meaning, and many tricky cases of quantification and the use of proper names. In addition, we considered only rather simple figurative descriptions (words or compounds), not touching the issue of compositionality of figurative speech. From a more technical perspective, we have also not scrutinized the different kinds of relations that are still required to hold between $w.CON_{s,f}$ and $w.CON$, if $w.CON_{s,f} \cap w.CON = \emptyset$. So, a necessary condition for figurative speech has been established that needs to be supplemented by sufficient ones. We also have no criteria available right now that lead us to distinguish between various types of figurative speech (e.g., metaphors vs. irony). Finally, we stop short of distinguishing between innovative figurative speech (like in the *ham sandwich* example) and conventionalized figurative speech (*systematic polysemy* (Pustejovsky, 1991; Nunberg, 1995)).

Acknowledgements. K. Markert was a member of the Graduate Program *Human and Machine Intelligence* at Freiburg University, funded by DFG.

References

Blasko, D. G. & C. M. Connine (1993). Effects of familiarity and aptness on metaphor processing. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 19(2):295–308.
 Carbonell, J. G. (1982). Metaphor: an inescapable phenomenon in natural-language comprehension. In W. G. Lehnert & M. H.

Ringle (Eds.), *Strategies for Natural Language Processing*, pp. 415–434. Hillsdale, N.J.: Lawrence Erlbaum.
 Chaffin, R. (1992). The concept of a semantic relation. In A. Lehrer & E. F. Kittay (Eds.), *Frames, Fields and Contrasts*, pp. 253–288. Hillsdale, N.J.: Lawrence Erlbaum.
 Cruse, D. (1986). *Lexical Semantics*. Cambridge: Cambridge University Press.
 Fass, D. C. (1991). *met**: A method for discriminating metonymy and metaphor by computer. *Computational Linguistics*, 17(1):49–90.
 Fauconnier, G. (1984). *Espace Mentaux*. Paris: Editions de Minuit.
 Gentner, D., B. Falkenhainer & J. Skorstad (1989). Viewing metaphors as analogy: the good, the bad, and the ugly. In Y. Wilks (Ed.), *Theoretical Issues in Natural Language Processing*, pp. 171–177. Hillsdale, N.J.: Lawrence Erlbaum.
 Gibbs, R. (1992). Categorisation and metaphor understanding. *Psychological Review*, 99(3):572–577.
 Grice, H. (1975). Logic and conversation. In P. Cole & J. Morgan (Eds.), *Syntax and semantics*, Vol. 3, pp. 41–58. New York: Academic Press.
 Hahn, U. & K. Markert (1997). In support of the equal rights movement for literal and figurative language: A parallel search and preferential choice model. In *Proceedings of the 19th Annual Conference of the Cognitive Science Society*, pp. 609–614.
 Hobbs, J. R., M. E. Stickel, D. E. Appelt & P. Martin (1993). Interpretation as abduction. *Artificial Intelligence*, 63(1-2):69–142.
 Indurkha, B. (1988). Constrained semantic transference: a formal theory of metaphors. In A. Prieditis (Ed.), *Analogica*, pp. 129–157. Los Altos, CA: Morgan Kaufmann.
 Inhoff, A., S. Lima & P. Carroll (1984). Contextual effects on metaphor comprehension in reading. *Memory and Cognition*, 12:558–567.
 Jackendoff, R. (1990). *Semantic Structures*. Cambridge, MA: MIT Press.
 Kittay, E. (1987). *Metaphor: Its Cognitive Force and Linguistic Structure*. Oxford: Clarendon Press.
 Lakoff, G. (1987). *Women, Fire, and Dangerous Things. What Categories Reveal about the Mind*. Chicago, IL: University of Chicago Press.
 Lakoff, G. & M. Johnson (1980). *Metaphors We Live By*. Chicago, IL: University of Chicago Press.
 Martin, J. (1992). Computer understanding of conventional metaphoric language. *Cognitive Science*, 16:233–270.
 Norvig, P. (1989). Marker passing as a weak method for inferencing. *Cognitive Science*, 13(4):569–620.
 Nunberg, G. (1995). Transfers of meaning. *Journal of Semantics*, 12:109–132.
 Pustejovsky, J. (1991). The generative lexicon. *Computational Linguistics*, 17(4):409–441.
 Romacker, M., K. Markert & U. Hahn (1999). Lean semantic interpretation. In *IJCAI'99 – Proceedings of the 16th International Joint Conference on Artificial Intelligence*.
 Searle, J. (1979). Metaphor. In A. Ortony (Ed.), *Metaphor and Thought*, pp. 93–123. Cambridge: Cambridge University Press.
 Stallard, D. (1993). Two kinds of metonymy. In *Proceedings of the 31st Annual Meeting of the ACL*, pp. 87–94.
 Tourangeau, R. & R. Sternberg (1982). Understanding and appreciating metaphors. *Cognition*, 11:203–244.
 Turner, M. (1988). Categories and analogies. In D. Helman (Ed.), *Analogical Reasoning*, pp. 3–24. Dordrecht: D. Kluwer.
 Veale, T. & M. T. Keane (1994). Belief modelling, intentionality and perlocution in metaphor comprehension. In *Proceedings of the 16th Annual Conference of the Cognitive Science Society*, pp. 910–915.
 Winston, M., R. Chaffin & D. Herrmann (1987). A taxonomy of part-whole-relations. *Cognitive Science*, 11:417–444.
 Woods, W. A. & J. G. Schmolze (1992). The KL-ONE family. *Computers & Mathematics with Applications*, 23(2-5):133–177.