Tilburg University

Tilburg University

Idioms, non-literal language and knowledge representation

van der Linden, H.J.B.M.

Publication date:
1992

Document Version
Publisher's PDF, also known as Version of record

Citation for published version (APA):
van der Linden, H. J. B. M. (1992). *Idioms, non-literal language and knowledge representation*. (ITK Research Memo). Institute for Language Technology and Artifical Intelligence, Tilburg University.

ITK MEMO

Idioms, non-literal language and knowledge representation *

Erik-Jan van der Linden

No. 12   1992/12

Abstract

There are two pertinent themes in the study of idioms in the area of Natural Language Processing. Firstly, idioms should be defined and located in the space of non-literal expressions. This will be the first aim of this paper. Secondly, a processing model should be developed. In this paper, the application of knowledge representation techniques in three different models for the representation and processing of idioms are discussed. The first, a *symbolic procedural* model extends the *two-level model* which was originally developed in computational morphology. The second is a simple *localist connectionist* model. The third, a symbolic hierarchical model, represents idioms as part of a lexicon conceived as an *inheritance hierarchy*. A comparison between the models is made in which the focus lies on the resolution of the ambiguity of idioms, the relation between the literal and non-literal interpretation and the syntactic flexibility of idiomatic expressions.

Key words: Natural Language Processing; Non-Literal Language; Idioms; Metaphors; Connectionism; Inheritance.

2

1 Introduction

Two issues are of importance in any computational theory of idioms. Firstly, a definition of idioms should be provided (section 2). Definitions of idioms in the linguistic literature are not adequate, as will be argued here, since they define what idioms are *not*: a positive definition should be supplied, that defines idiomaticity as a property. Furthermore, idioms should be located in the space of non-literal expressions in order to understand why these expressions are non-literal (3-4). Secondly, models for the representation and processing of idioms should be designed. In section (5), three different models for the representation and processing of idioms will be presented, which use different knowledge representation techniques. The first extends the two-level model which was originally developed in computational morphology. The second is a simple localist connectionist model. [2] The third represents idioms in a lexicon that is modelled as an inheritance hierarchy. [3] The focus in comparing the three models will be on the resolution of the ambiguity between the idiomatic and non-idiomatic interpretation of an idiom, the relation between the literal and non-literal interpretation and the syntactic flexibility of idiomatic expressions. (For a more elaborate discussion of other aspects like syntactic-semantic processing and prosodic properties of idioms see van der Linden (in prep.)).

2 Idiomatic expressions and non-literal language

2.1 Idioms and compositionality

In the present section two attempts to account for idioms on the assumption of compositionality will be discussed and rejected. It will be concluded that the meaning of idioms cannot be subject to compositionality. This is important for a proper classification of idioms as non-literal expressions.

In the first subsection compositionality will be introduced. Next, the two attempts are discussed. Then, a definition of idioms will be provided.

2.2 Compositionality

The description of the relation between the form of the expressions of a language and their meanings is a central goal of linguistic theory. The *compositionality principle* (henceforth CP) is one of the principles that describe this relation. In its most general form it goes as follows:

---

[2] The first two models are reported on in van der Linden and Kraaij (1990).

[3] A variant of this model can be found in van der Linden (1992)

> "The meaning of an expression is a function of the meanings of its parts and of the way in which they are syntactically combined" (Partee 1984:281)

CP accounts for the ability of the language user to understand the meanings of sentences not encountered before. It

> "(...) is required to explain how a finitely representable language can contain infinitely many nonsynonymous expressions" (Fodor and Pylyshyn 1988:43)

Opposed to CP is the strong version of the so-called principle of contextual interpretation, which can be stated as follows.

> "Nur in Zusammenhange eines Satzes bedeuten die Wörter etwas" (Frege 1884:73)

This strong version of the contextuality principle would imply that all meanings of sentences are "primitive in a sense" (Hoeksema 1984:35). However, a system in which every concept could be expressed by any sound "(..) would amount to no communication system at all (...)" (Makkai 1978:405). Some (aspects of) word meanings should be invariant across contexts. Here, compositionality is considered a default from the linguistic point of view for the interpretation of syntactically complex expressions. Hoeksema mentions idioms and indexical expressions as 'exceptions' to CP (see Partee (1984) for solutions to other problematic phenomena for CP).

2.3 Compositionality and the meaning of idioms

Although intuitively the meaning of an idiom is not a function of the meaning of its constituent parts, attempts have been made to account for the meaning of idioms under the principle of compositionality. A trivial argument against this are cases where parts of idioms do not have a meaning outside the idiom. Examples are *queer the pitch* and *spic and span*. The meaning of these idioms cannot be a function of the meanings of the constituents because the parts have no meaning (Wasow et al. 1983).[4] Secondly, some idioms have an idiosyncratic syntactic structure. Since semantic principles are formulated to combine the meanings of syntactically well-formed expressions, they don't apply in these cases (Wasow et al. 1983). Examples are *by and large* or *trip the light fantastic*.[5]

---

[4]Because of the existence of these idioms Boatner et al. (1975) are wrong in describing idioms as "the assigning of a new meaning to a group of words which already have their meaning" (Boatner et al. 1975:iv).

[5]As Wasow et al. remark this is not the whole story for idioms like *long time no see* which appear to be semantically composed.

For idioms with non-idiosyncratic syntactic structures, the parts of which can be assigned meaning outside the idiom, it follows from the definition of CP that if CP applies this can only be accomplished if parts of the meaning of the idiom can be assigned to parts of the idiom. Two possibilities exist.

One part carries the whole idiomatic meaning It could be possible that the meaning of the idiom is a property of one of the parts, and that the other part has *no* meaning (Ruhl, cited in Wood 1986; Partee 1984). In the case of *kick the bucket* the meaning die is assigned to *kick* and no meaning to the other part. The fact that *kick* means *to die* in slang seems to contribute to the plausibility of this claim. It raises the question, however, why one cannot say *Pat rested the bucket* to mean *Pat rested* (Wasow et al. 1983). Also, the origin of *kick the bucket* has little to do with the meaning of *kick* in slang. This approach thus fails.

Both parts carry part of the idiomatic meaning If the parts that constitute the expression can be assigned part of the idiomatic meaning, compositional combination of these meanings results in an idiomatic meaning for the whole expression (Gazdar et al. 1985). For some expressions the relation between form and meaning is not arbitrary: a relation exists between parts of the idiom and parts of the meaning of the idiom. It follows that parts of idioms could be semantic units (see also Makkai (1978)). Evidence could be sentences where parts of expressions are modified (3), quantified (1) or parts are omitted in elliptical constructions (2) (Wasow et al. 1983).[6]

(1)   He pulled a string or two.

(2)   My goose is cooked, but yours isn't.

(3)   He left no legal stone unturned.

The fact that these idioms are regular from a syntactic point of view and that the words constituting them have a meaning outside the idiom, has led Gazdar et al. (1985) to include a treatment of these expressions under compositionality in GPSG. (A similar line of reasoning can be found in the work of Gibbs and his co-workers (Gibbs 1980; Gibbs and Gonzales 1985; Gibbs 1986; Gibbs and Nayak 1989; Gibbs, Nayak and Cutting 1989; Gibbs, Nayak, Bolton and Keppel 1989. These publications will be referred to in this paper as (Gibbs various)). To for instance the verb *spill* two meanings are assigned: spill', the non-idiomatic sense and spill'', the idiomatic sense meaning *divulge. Beans* also has two

---

[6] The acceptability of the sentences under consideration is discussed by Schenk (1992).

5

senses, where one means approximately *information*. Spill" then is a partial function, that can take only one argument: the(beans)". The compositional semantic principles in GPSG combine these two, resulting in the idiomatic meaning of the phrase. There are some problems with and arguments against this compositional analysis. Firstly, what Gazdar et al. do not mention, is that it has to be prevented that other functors combine with the(beans)". This could be accomplished in two ways: for *all* functors in the lexicon it could be stated which arguments they do not combine with. This would, however, imply the stipulation of a large number of negative facts about *all* functors that are related to items in the lexicon: for instance spill' and take' etc. do *not* take the(beans)" as an argument. This kind of linguistic description, namely massively stating 'negative' facts, is not common in linguistics. Another way out is to describe the(beans)" as a 'partial argument'. However, whereas functors can be partial this is not the case for arguments.[7]

This second compositional account fails as well. The fact that parts of idioms seem to carry meaning can yet be accounted for in another way (section (3)).

The conclusion of the present section should be that the meaning of idioms cannot be accounted for on the assumption of compositionality. Note that contextuality cannot give an account either: the meanings of idioms do not differ from the meanings of other lexical elements with respect to their invariance across contexts. Apparently, a different principle is needed.

2.4 Definition

In the present section a principle will be introduced, *idiomaticity*, which describes the discrepancy between form and meaning of idioms. With this principle, it becomes possible to present a definition of idioms.

According to Gazdar et al. (1985:327) 'Traditional wisdom dictates that an idiom is by definition a constituent or series of constituents where interpretation is not a compositional function of the interpretation of its parts.'. Comparable definitions can be found in Hocket (1958); Fraser (1970); Katz (1973); Heringer (1976); Chomsky (1980); Wood (1986): an idiom is 'wholly non-compositional in meaning'[8]; Di Sciullo and Williams (1987): listemes do "not have the form or interpretation specified by the recursive definitions of the objects of the language"; Abeillé and Schabes (1988); Schenk (1992):

---

[7] Vergnaud (1985) hypothesises that nouns that occur in idioms can only be inserted in their canonical context. This is a general rule and not a property of the idiomatic noun, and therefore such a notion is not equal to that of partial argument.

[8] Wood (1986) also gives a nice overview of the literature on idioms up to 1980.

6

"expressions consisting of more than one word, for which a literal interpretation does not give the correct meaning"; Erbach (1991).[9] Three aspects of these definitions need consideration. (a) should idioms always be multi-lexemic expressions? (b) do these definitions demarcate idioms from other expressions? (c) should idioms be defined as a class of expressions, or should idiomaticity be defined as a property of expressions?

Idioms are multi-lexemic expressions Defining an idiom as *any* grammatical form the meaning of which is not deducible from its structure (Hockett 1958, cited in Wood 1986, my emphasis), entails that single morphemes are the simplest case of idioms (Fraser 1970:22)). It would imply that every morpheme is granted the status of idiom. The important difference between morphemes and idioms under the definition of Hockett is that for morphemes there exists no structure which enables deduction of meaning, whereas in the case of idioms, such a structure does exist, but cannot be used for deduction. So, although this may seem trivial it has to be stated explicitly that *this paper will limit the notion idiomaticity to complex expressions that are made up of more than one lexeme.*

Demarcation from other classes Most definitions in the literature do not provide properties that distinguish idioms from other types of expressions. They describe what idioms are *not*, compositionality does not apply, but do not indicate which principles *do* apply (contextuality, meaning postulates, etc.). A positive definition of idioms which says what the meaning of an idiom *is*, is preferable because it makes stronger claims.[10]

Idioms or idiomaticity Idiomatic expressions do not form a homogeneous class. It is not the case that the meaning of some expressions is completely compositional, and of others completely non-compositional. Expressions that are no idioms *proper* may be partly idiomatic. A first example are collocations, which are idiomatic with respect to generation but not with respect to analysis (Fillmore et al. 1988). If a language user merely knows the meanings of the words *school* and *whales*, he will be able to arrive at the interpretation of a group of fishes when encountering the expression *school of whales* without knowledge of the collocation. Generation of such an expression without this knowledge, however, is not possible. It is likely that the language user will generate an expression like *group of*

---

[9] Also various papers in Everaert and van der Linden (1989) and Everaert et al. 1992

[10] Compare Wasow et al. 1983, who assert that "the idiomatic meaning is assigned to the whole phrase" (Wasow et al. 1983:110). See also Fillmore et al. 1988:501. See also Wilensky and Arens (1980): "... these constructs are phrasal in that the language user must know the meaning of the construct as a whole to use it correctly" (Wilensky and Arens 1980:117)

whales. Note that this asymmetry does not apply to other expressions. A language user that does not know the meaning of the word *bank* will neither be able to analyse the word, nor to generate it; a language user that knows the word, can do both. A second example is a construction like *it is raining cats and dogs* in which *it is raining* can be assigned a compositional interpretation, although the expression as a whole is idiomatic (as will be argued below). Compositionality thus seems to apply to some aspects of meaning in a construction, whereas other principles apply to other aspects. Therefore it seems fruitful to define a notion of *idiomaticity* as a property and to apply this notion to parts of the meanings of expressions, rather than to claim that a certain class of elements should be described as idioms with an all-or-none property of non-compositionality that distinguishes them from all other expressions (Wood 1986; Schenk 1992). Like compositionality and contextuality, idiomaticity is a property that may apply to parts of expressions (Wood 1986; Napoli 1988:331).

Taking into account these three points, the definitions of *idiomaticity* and *idiomatic expression* may run as follows.

Definition 1 (Idiomaticity) *Idiomaticity is a property of aspects of the meaning of complex (multi-lexemic) expressions. Idiomaticity implies that these aspects are exclusively a part of the meaning of the expression as a whole.*

Definition 2 (Idiomatic expression) *An* idiomatic expression *is an expression some aspect(s) of the meaning of which is (are) subject to idiomaticity.*

With these definitions, it becomes possible to define idioms.

Definition 3 (Idioms) *Idioms are expressions all aspects of the meaning of which are subject to idiomaticity.*

From this definition it follows that expressions in which one of the parts has its non-idiomatic meaning will not be considered idioms but idiomatic expressions. Dutch examples[11] are *op de kleintjes letten* (*of the little_ones take_care, to be careful with ones money*) (Everaert 1989) and *het regent pijpestelen* (*it rains pipe_stems, it's raining cats and dogs*) which are not idioms, but idiomatic expressions in which *letten op* (*to be careful with*), and *het regent* (*it's raining*) retain their non-idiomatic interpretation.

---

[11]Examples of Dutch idioms will be presented with a word-by-word translation (if this exists) and a well-formed translation in English (if this exists) in which, if possible, a comparable English idiom is used. This does not imply, however, that the analysis of the Dutch idiom applies to that of the English idiom.

Pijpestelen is in itself not an idiom since its meaning is not a property of *pijpestelen* itself, but of the expression *het regent pijpestelen* as a whole. This same line of reasoning applies to such expressions as *aanstalten maken* (*get ready*) and *in aantocht zijn* (*be on the way*): *aanstalten* and *aantocht* can only occur in these expressions, so the expressions must be idiomatic since this aspect of the meaning is an exclussive property of the expression as a whole. They are, however, not idioms, since the verbs in these expressions retain there non-idiomatic meaning.

With the definition of *idiomaticity, idiomatic expression* and *idiom* a more precise classification of expressions in which meaning is a property of the whole expression can be given. The definition of idioms in for instance Wood (1986) only defines idioms and cannot relate these to other expressions.

3 Metaphorical properties of idioms

Idiomaticity does not imply *arbitrariness* of meaning. In the present section metaphorical properties of idioms, which are important in this respect, will be discussed. Two notions, *motivation* and *isomorphism* will be introduced. Some attention will also be paid to the relation between metaphorical properties and compositionality, since metaphorical properties have mistakenly been taken as an argument in favour of the compositionality of the meaning of idioms.

3.1 Motivation and isomorphism

Metaphors are general principles that link some domain to some target.[12] An example might be ANGER IS THE HEAT OF A FLUID IN A CONTAINER. Metaphors like this may underlie several metaphorical expressions (4, taken from Lakoff (1987:380-381)). Metaphors may underlie basic (4c) and complex expressions (4a;b) .

(4) a. You make my *blood boil.*
    b. He's just *letting off steam.*
    c. He *exploded.*

Most idioms are frozen metaphorical expressions. For some idioms, like *kick the bucket,* the underlying metaphor is no longer visible. For other idioms, the metaphor is visible for language users and determines the appropriateness of the idiom in certain contexts and constructions. In cognitive linguistics (Lakoff 1987) and psycholinguistics (Gibbs various), the latter claim has been provided with a number of arguments. Firstly, there is a certain agreement between speakers about the metaphors and images that are underlying idioms. If no conventional images or metaphors would underly idioms,

---

[12] See Lakoff and Johnson (1980); for computational models of metaphorical language see for instance Martin (1989).

such agreement would not be expected. Secondly, the syntactic behaviour of idioms can partly be explained in terms of metaphorical properties.[13] In section (6) it is argued that at least where this concerns passivisation, there is a strong correlation between metaphorical properties and occurrence of the idiom in the passive: with these properties absent, passivisation is not possible. Napoli (1988) presents an *extension*-test to see whether idioms can be analysed. An example of an extension is for instance *The cat got out of the bag and wrecked havoc.* "Extendability can call for building up a story around the idiom which is plausible at the metaphorical level." (Napoli 1988:330). Although one could argue that Napoli's extensions are wordplay, and therefore do not provide evidence on which can build a linguistic theory, the agreement speakers have about extendability shows that underlying metaphors are important. Thirdly, Nayak and Gibbs (1990) showed that contexts which have a certain underlying metaphor will affect the appropriateness of idioms, in that idioms with the same underlying metaphor are more acceptable in this context.

There is a variety of terms in the literature to refer to the meraphorical properties of idioms. Here, the notions *motivation* and *isomorphism* will be used (Geeraerts 1992; the interpretation of these notions here differs somewhat from Geeraerts' interpretation). In the next section, motivation and isomorphism will be introduced, and compared to other notions proposed in the literature. Then, the relation between metaphorical properties of idioms and compositionality will be discussed.

Motivation   The conventional image underlying an idiom, or part of it, may result in the possibility of establishing a relation, a *motivating* link, between the idiomatic interpretation of the idiom, and the non-idiomatic interpretation of the idiom (Lakoff 1987). For instance *blow the fuse* offers an image for loss of temper; *spill the beans* offers an image for making secret information public; *saw logs*, meaning *to be sound asleep*, can also be interpreted on the basis of a conventional metaphor. The relationship between the two is *motivated* just in case there are independently existing elements of the conceptual system that link the idiomatic and non-idiomatic meaning (Lakoff 1987:451-452). This link may be metaphorical or may be a conventional image. Note that this does not imply that meaning or form of the idiom are *predictable*. Motivating links make sense of idiomatic expressions and therefore make them easier to understand, learn, remember, and use than random pairings.

---

[13] Although it is argued in van der Linden (in prep.) that some doubt can be cast upon psycholinguistic experiments by Gibbs (various) that support this claim.

10

Isomorphism Not only may a relation exist between the non-idiomatic interpretation as a whole and the idiomatic interpretation as a whole, it may also be the case that parts of the idiomatic and non-idiomatic interpretation maintain relations: an *isomorphism* (Geeraerts 1992) may exist between the parts of the idiomatic and the non-idiomatic interpretation. In for instance *blow the fuse* it is possible to find a part-to-part-correlation. A *fuse* refers to strained patience and *blow*, "colloquially, makes sense (thus also *blow one's top/lid/cool/gasket*)" (Wood 1986:36). *The beans* in *spill the beans* may refer to the information that is supposed to be kept secret. *Spill* refers to making that information public (Lakoff 1987:451). Isomorphism is thus a relation between the parts of the meaning of the idiom, and the parts of the idiom.

Wordly motivation For some idioms it may be the case that the meaning of parts of the idiom equals (an extension of) the meaning of the part outside the idiom. The part of the idiom has a similar referent inside and outside the idiom. A *wordly motivation* exists for parts of the idiom. Geeraerts mentions *parels voor de zwijnen gooien* (*pearls for the swines throw; cast pearls before swine*). Here *parel* can be interpreted as *something with a special value* independently of the expression: it is even listed as such in the lexicon. *Zwijn* can be interpreted a an extension of its lexical meaning *unworthy person*. However, the meanings of these constituents in the idiom are a property of the expression as a whole, and therefore, the expression is idiomatic.

3.1.1 Examples

Motivated and isomorphic

(5) a. het paard achter de wagen spannen.
   the horse behind the cart to_set.
   set the cart before the horse.

   b. de koe bij de horens grijpen.
   the cow by the horns to_take.
   to take the bull by the horns.

Non-Motivated and isomorphic

(6) a. de lakens uitdelen.
   the sheets to_hand_out
   to play first fiddle.

   b. een hak zetten.
   a cut to_set.
   to play a nasty trick.

Motivated and non-isomorphic

(7) a. geen lange draad meer spinnen.
no long thread more to_spin.
to die soon.
b. de geest geven.
the ghost to_give.
to give up the ghost.

Non-Motivated and non-isomorphic

(8) a. de kat uit de boom kijken.
the cat out_of the tree to_look.
to wait to see which way the wind blows.
b. de kat de bel aanbinden.
the cat to_bell.
to bell the cat.

3.1.2 Terminology

In order to avoid terminological confusion and to indicate that the account presented here subsumes those in the literature, figure (1) contains an overview of terminology. For every term in the literature it is indicated how the term coincides with one term or a combination of terms used in this paper. If the term is indifferent with respect to a certain factor, this is denoted as *indif.*.

Table 1: Terminology

term	author(s)	motivated	isomorphic	wordly mot.
analyzable	Gazdar et al. (1985)	*indif.*	*yes*	*indif.*
	Napoli (1988)	*indif.*	*yes*	*indif.*
unanalyzable	Gazdar et al. (1985)	*indif.*	*no*	*indif.*
	Napoli (1988)	*indif.*	*yes*	*indif.*
normally decomposable	Gibbs and Nayak (1989)	*indif.*	*yes*	*yes*
	Nunberg (1978)			
abnormally decomposable	Gibbs and Nayak (1989)	*yes/no*	*yes*	*indif.*
	Nunberg (1978)			
non-decomposable	Gibbs and Nayak (1989)	*yes/no*	*no*	*indif.*
	Nunberg (1978)			
opaque	Gibbs and Nayak (1989)	*no*	*indif.*	*indif.*
transparent	Gibbs and Nayak (1989)	*yes*	*indif.*	*indif.*
imageable	Lakoff (1987)	*yes*	*indif.*	*indif.*
metaphorical referents	Lakoff (1987)	*indif.*	*yes*	*indif.*
explanation	Zernik (1987)	*yes*	*indif.*	*yes*

- **analyzable**: "analyzable into lexical subparts" (Napoli 1988:329)

- **unanalyzable**: "syntactically complex lexical items with a single undecomposable semantic interpretation" (Gazdar et al. 1985:244)

- **(normally) decomposable**: "each of the components refers in some way to the components of their idiomatic referents" (Gibbs and Nayak 1989:105); "(...) an idiomatic transitive VP is DECOMPOSABLE just in case it is used to refer to a state or activity such that it would normally

12

be believed that that activity could be identified as an open relation Rxb, such that the object NP of the idiom refers to b, and the verb to R" (Nunberg 1978:124)

- **abnormally decomposable**: "the object NP (...) does not *itself* refer to some component of the idiomatic referent, but only to some metaphorical relation between the component and the referent (...)" (Gibbs and Nayak 1989:106)

- **non-decomposable**: "idioms whose individual components did not make a contribution to the overall figurative meaning" (Gibbs and Nayak 1989:108)

- **imageable**: "(...) idioms that have associated conventional images" (Lakoff 1987:447)

- **explanation** an association between a pattern and a concept (Zernik 1987:106)

3.1.3 Motivation, isomorphism and compositionality

Just because of the fact that parts of idioms may have metaphorical referents, it has been claimed that the meaning of these idiomatic expressions is controlled by the principle of compositionality (Gazdar et al. 1985; Gibbs and Nayak 1989). In section (2) this approach has been argued against. For instance, although Gibbs and Nayak (1989) are right when they claim that parts of *decomposable idioms* have identifiable meaning, this does *not* imply that the property of having this meaning is a property of the lexeme outside the idiom, and that the meaning of the idiom is subject to compositionality. Looking at dictionaries one observes the same: idioms are listed in the entry of one (or more) of the content-words in the idiom as a *unit*. The dictionary user does not find 'idiomatic meaning' of every content word leaving him to find out the meaning of the whole himself. Besides, the relations between the meaning of parts of the idiom and parts of the meaning of the idiom do not involve linguistically significant generalisations (Nagy 1978:296) and should therefore be stipulated with the representation of the idiom within the lexicon. Thus although this might seem paradoxical at first sight, it is possible to enable distribution of meaning while adhering to the definition of idiomaticity. The crucial point is that *the distribution should be a property of the idiom as a whole*: it is a property of the idiom as a whole whether the parts can be assigned metaphorical referents (and whether the idiomatic expression as a whole can be motivated). Geeraerts (1992) argues that the principle of compositionality could be extended in order to capture this. Compositionality could be extended with a static interpretation besides the usual dynamic interpretation. If some relation exists between the meaning of the constituent parts of

13

an expression and the meaning of the expression as a whole, then a static notion of compositionality applies. In this paper the principle will not be given a dynamic extension. This would make it look like compositionality remains the only design principle for natural language, whereas, again, the existence of other principles should be warranted.

A similar line of reasoning holds for attempts to regard *idiomatic meanings* as *literal meanings* (Dascal's (1987) 'moderate literalism'). Note that Dascal's notion of literal language would necessitate stretching the usual conception of literal language: *to literally kick the bucket* means *to hit some designated pail with the foot*, and not *to die*.

4 Comparison to other classes of expressions

4.1 Non-literal language in general

Given the definition of idiomaticity presented here, idioms differ from other kinds of non-literal language such as "indirect" speech acts, implicature, metonymy, irony, simile and sarcasm. In the case of idioms non-literal meaning is a property of the expression as a whole which is represented within lexical entries, whereas in the case of other non-literal expressions meaning is derived on the basis of other information sources (like metaphorical principles (Lakoff and Johnson 1980; Martin 1989), Grice's maxims, etc.).

4.2 Idiomatic and metaphorical expressions compared

Motivated idioms are conventionalised metaphorical phrases that are still to some extent transparent, i.e. for which the underlying metaphor is recognisable.

The line between complex metaphorical expressions and idioms is rather thin; the main difference is that within a complex metaphorical expression the meaning of the whole expression is a function of the metaphorical parts. One can thus also observe that in a complex metaphorical expression all parts have their own metaphorical meaning, whereas for an idiom this possibly metaphorical meaning cannot exist outside the idiom.

Compare for instance *his heart sank* or *his heart jumped* with *she broke his heart*. The metaphor OUR HEART IS THE PLACE WHERE OUR EMOTIONAL LIFE RESTS underlies these three expressions. The verbs have a metaphorical interpretation as well. However, only one of the three is idiomatic and fixed, namely *she broke his heart*. It is claimed here that this is due to the semantic specialisation of the expression. This becomes clear from the fact that *break someone's heart* is more limited in its distribution, it can either refer to love (*the beautiful woman broke his heart*) or to endearment (*the*

14

hairy-nosed wombat broke his heart). All possible metaphors underlying *heart* and *jumped* apply for *his heart jumped: Mary scared Jack. His heart jumped; Mary told Jack she loved him. His heart jumped; Mary made Jack angry. His heart jumped. Fixedness* is thus not just an unexplainable property of expressions but correlates with semantic specialisation. *His heart jumped is thus a complex metaphorical expression; broke his heart* is an idiomatic expression.

If the expression *take the bull by the horns* is classified as a metaphor (Schenk 1992) *take* would mean *deal with; the bull* means *a problematic matter;* and *by the horns* means *at the most important part of the matter*. Notice, however, that none of the subexpressions can occur outside the expression carrying this meaning (9, 10) (# indicates that no idiomatic interpretation is possible).

(9)   # The bull bothered me.

(10)   # He decided that he would take the bull by the lead.

Metaphorical reference is thus a property of the whole expression, and not of the individual parts: it is distributed by the expression to the parts. Therefore these expressions are considered idiomatic as well: their meaning is a property of the expression as a whole.

4.3   Collocations

A category that is related to idioms is that of *collocations*. Collocations consist of a head-argument combination (*een moord begaan, commit murder*), or a head-adjunct combination (*een school vissen, a school of fish*). In defining collocations the move of distinguishing the *property* idiomaticity from *classes* of expressions turns out to be useful. Idiomaticity applies to *encoding* for collocations, but not to *decoding* (Fillmore et al. 1988). This means that with respect to analysis, decoding, a collocation can be interpreted compositionally on the basis of the literal meaning or a metaphorical extension of its parts. With respect to generation, encoding, however, a speaker who does not know the expression does not know what head-argument combination to use when he wants to express something about the argument. For example, from (the metaphorical extension of) the meaning of *begaan* (*commit*) and the literal meaning of *moord* (*murder*) it is possible to compositionally form the meaning of *een moord begaan* (*commit murder*). However, without knowledge of this expression, the chance to use the verb *begaan* (*commit*) in a generation process is in principle equal to the chance of the use of other verbs

15

Table 2: Dimensions of expressions

```
literal
          simple
                    conv.
                              words sleep
          complex
                    novel
                              compositionality John sleeps
non-literal
          simple
                    conv.
                              metaphor (I) exploded
                    novel
                              metaphor (He) detonated (me)
          complex
                    conv.
                              routine formulae Good God!
                              idiomaticity kick the bucket
                    novel
                              metaphor John is a rat
                              sarcasm You're smart
                              speech acts It's cold here
                              ...............
```

(execute a murder, carry out a murder, do a murder, make a murder).

4.4   Conclusion

To summarise the current and the previous section, the space in which idioms should be located can be sketched with three dimensions: literal vs. non-literal expressions, simple vs. complex expressions. In table (2) a third dimension is added: novel vs. conventional expressions. Considering idioms as simple expressions does not do justice to their internal structure. Although idioms have metaphorical properties, metaphorical aspects of idioms are conventional properties of the expression as a whole. One can for instance not say *John wanted the beans* to mean *John wanted the information*. Idioms are thus complex, non-literal expressions, with a conventional meaning. The table also indicates wy idiomaticity cannot be defined in terms of non-compositionality: a *positive* definition is required to demarcate idiomatic expression from other expressions.

5   Representation and processing of idioms

Now that it is clear what expressions are to be considered idiomatic and what their properties are, it becomes possible to present models of for the representation and processing of idioms. Three models will be presented here: a simple symbolic algorithmic model, a localist, connectionist model and a model in which the lexicon is viewed as an inheritance hierarchy.

The issue concerning the representation and processing of idioms that will be concentrated upon in the models to come, will be the resolution of the ambiguity of idioms. The next section concerns syntactic flexibility. The general approach to NLP here, is that the NL processor operates efficiently if it adopts an incremental mode of interpretation, and interprets input as immediate as possible (Thibadeau et al. 1982). Ambiguities are resolved on the basis of a best-first strategy. The question, then, is which possibility is the best one, and on the basis of what knowledge choices should be made.

5.1   Conventionality

A choice between the literal and non-literal reading of an idiom can be made using various kinds of linguistic information, but the claim here is that the mere fact that one of the analyses is idiomatic suffices. Besides, this choice does not have to be stipulated explicitly. Rather it follows naturally from the architecture of the lexicon and the retrieval process, provided an appropriate model of the lexicon is used.

Phrases consisting of idioms *can* in most cases be interpreted non-idiomatically as well. Very rarely, however, an idiomatic phrase *should* in fact be interpreted non-idiomatically (Koller 1977:13; Chafe 1968:123; Gross 1984:278; Swinney 1981:208). Psycholinguistic research indicates that there is clear preference for the idiomatic reading (Gibbs 1980; Schweigert and Moates 1988). We will refer to the fact that phrases should be interpreted according to the lexical, non-literal meaning, as the 'conventionality' principle. If this principle could be modeled in an appropriate way, this would provide a heuristic that would render the interpretation process more efficient since other than lexical knowledge is not nessecary for the resolution of ambiguities. So, the resolution of the ambiguity occurs as soon as the idiom has been encountered in the input.

When can and does an incremental processor *start* looking for idioms? Psycholinguistic research indicates that idioms are not activated when the 'first' (content) word is encountered (Swinney and Cutler 1979). There is, from the computational point of view, no need to start 'looking' for idioms,

17

Figure 1: Continuation class model: lexicon structure and algorithm

```
(a)                              (b)  DO read a letter
                                        IF word has been found THEN
k-i-c-k*---b-u-c-k-e-t*                    IF this word forms a lexical item
          h-a-b-i-t*                          with previous word(s)
          \e-e-l-s*                        THEN make its information
                                              available to syn/sem process
                                        ELSE make word information
                                              available to syn/sem process
                                      UNTIL no more letters in input.
```

when only the first word has been found since this would only result in increase of the processing load at higher levels. In Stock's (1989) approach to ambiguity resolution the idiomatic and the non-idiomatic analysis are processed in parallel. An external scheduling function gives priority to one of these analyses. Also, the disambiguation process already starts when the 'first' word has been encountered. As we have stated, this increases the load on higher processes.

5.2 An extension of the notion *continuation class*

The first model presented here extends the notion *continuation class* from two-level morphology.

Lexical representation Lexical entries in two-level morphology are represented in a trie structure, which enables incremental lookup of strings. A lexical entry consists of a lexical representation, linguistic information, and a so-called continuation class, which is a list of sublexicons "the members of which may follow" (Koskenniemi 1983, p. 29) the lexical entry. In the continuation class of an adjective, one could, for instance, find a reference to a sublexicon containing comparative endings (ibid. p. 57). An obvious extension is to apply this notion beyond the boundaries of the word. A continuation class of an entry **A** could contain references to the entries that form an idiom with **A**. An example is (1a).

Algorithm A simple algorithm is used to retrieve idioms (in (1b) the relevant fragment of the algorithm is represented in pseudocode). The result of the application of the algorithm is that linguistic information associated with the idioms is supplied to the syntactic/semantic processor. The linguistic information includes the precise form of the idiom, the possibilities for modification etc. Note that conventionality is modeled explicitly. [14]

---

[14] A toy implementation of the lexicon structure and the algorithm has been made in C.

5.3 A connectionist model

The second model we present here is an extension of Cottrell's (1988) localist connectionist model for the resolution of lexical ambiguity. The model (2) consists of four levels. Units at the lowest level represent the smallest units of form. These units activate units on the level that represents syntactic discriminations, which in turn activate units on the semantic level. The semantic features activate relational nodes in the semantic network. *Within* levels, inhibitory links may occur; *between* levels excitatory links may exist, there are, however, no inhibitory links within the semantic network. The meaning of idioms is represented as all other relational nodes in the semantic network. On the level of semantic features, the idiom is represented by a unit that has a *gate* function similar to so-called *SIGMA-PI units* (Rumelhart and McClelland 1986:73): in order for such a unit (**A**) to receive activation, all units activating **A** bottom-up should be active. If one of the units connected to a unit **A** is not active, **A** does not receive activation. Thus when the first word of an idiom is encountered, the idiom is not activated, because the other word(s) is (are) *not* activated. However, once *all* relevant lexemes have been encountered in the input, it becomes active. Note that an external syntactic module is supposed to activate one of the nodes in case of syntactic ambiguity. Since there is more than one syntactic unit activating the idiom, the overall activation of the idiom becomes higher than competing nodes representing non-idiomatic meanings. The idiom is the strongest competitor, and inhibits the non-idiomatic readings. The conventionality principle is thus modeled as a natural consequence of the architecture of the model. Figure (3a) and (3b) show the activation levels of the active units in the model: only activation levels above treshold (500) are displayed. The appendix gives some technical details. The model has been implemented in C with the use of the Rochester Connectionist Simulator (Goddard et al. 1989) by Wessel Kraaij.

Fig (2)   Network representation

Semantic network

Semantic features

Syntactic features

Wordform

kick (object)   kick (action)   die   pail (object)

0.25   0.8

kick as object   -0.8   kick as action   die as idiom   bucket as pail
Gate
-0.8

0.25

kick N   -0.8
-0.8   kick V   bucket N
1.0
III$_a$ III$_b$ III$_c$

0.8

kick   bucket
1.0
I   II

: Unit simulating external input from syntactic module

: Unit simulating external input from sub word level

0.8   : Link with corresponding weight

Fig (3)   Unit structure

Output

Av

Internal
activation

0.6   Sv   0.6   0.6   0.6

Extern   Excit.   Inhib.   Gate

20

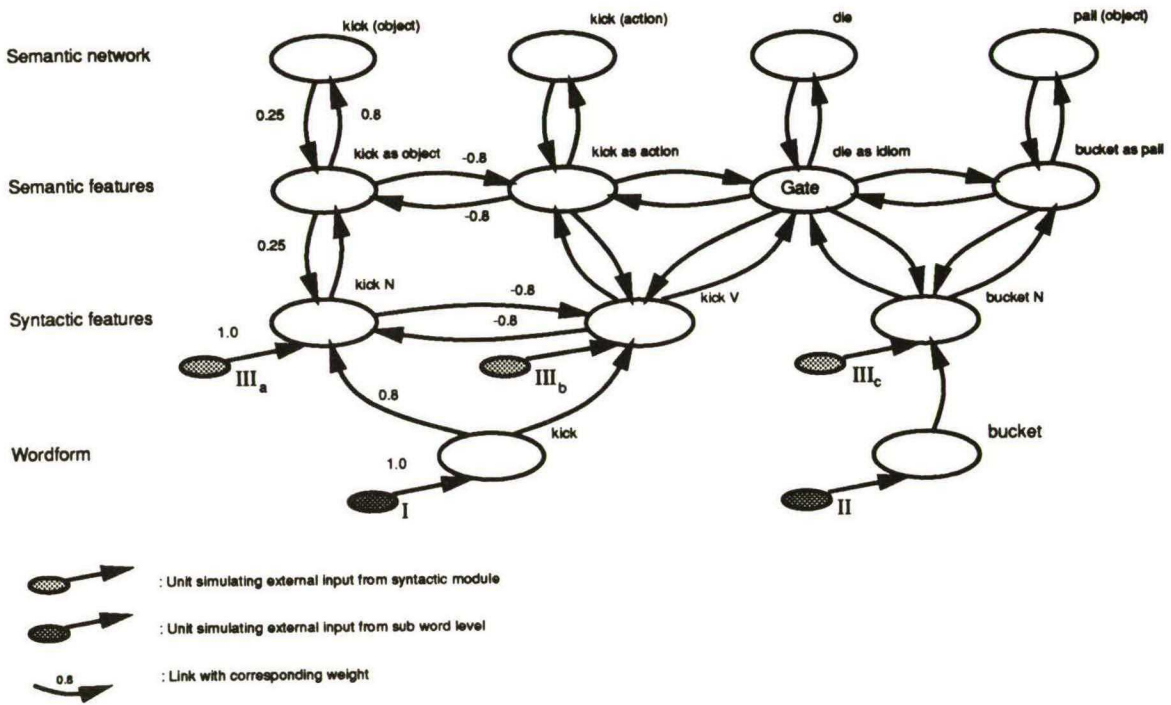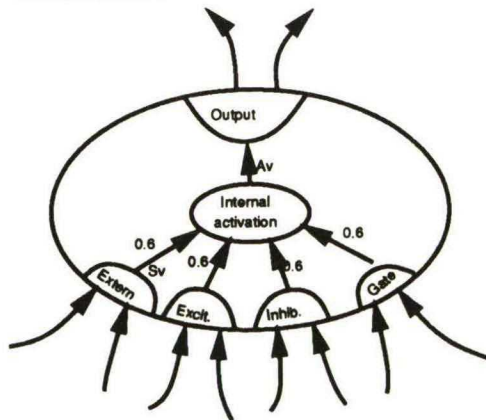## 5.4 Idioms in an inheritance hierarchy

Inheritance mechanisms are becoming increasingly important in the study of natural language processing. [15] A lexicon modeled as an inheritance hierarchy allows for the stipulation of general principles on high and abstract levels of representation, and therefore avoids the stipulation of redundant information. The concept of inheritance can also be applied to a lexicon that contains idioms. The model discussed here, is described in detail in van der Linden (1992). Here, we will concentrate on the structure of the lexicon.

Syntactic information   An idiom and its verbal head (*kick* in the case of *kick the bucket*) maintain an inheritance relation: the idiom can be said to inherit part of its properties from its head. Idioms can be represented as signs that are syntactically viewed as functor-argument structures [16] and have the same format as the verbs that are their heads (see also Zernik and Dyer (1987)). It is therefore possible to relate the syntactic category of the idiom to that of its head. The information that the object argument is specified for a certain string, can be added monotonically. The verb (*kick*) subcategorizes for the whole set of strings with category *np*, whereas the idiom subcategorizes for the subset of that set (*the + bucket*).

The relation between verb and idiom could be specified as KICK $\succ$ KICK_THE_BUCKET, where KICK and KICK_THE_BUCKET are represented as in (11) and $\succ$ denotes an inheritance relation between two signs. $KICK \succ KICK\_THE\_BUCKET$ states that $KICK\_THE\_BUCKET$ is a specialisation of $KICK$.

The grammatical theory for which this lexical structure is designed, is categorial grammar. KICK: $< (np\backslash s)/np >$ denotes a sign named KICK. The sign is an n-tuple $<a_1,.., a_n>$ which in this case only consists of syntactic information. The syntactic category denotes a *functor* which takes an *np* to its right (indicated with the /) and results in a category $(np\backslash s)$, which takes an *np* to its left, and results in a *sentence*. KICK_THE_BUCKET inherits this information, but adds a specific value for the prosody of the argument: the+bucket.

(11) KICK: $< (np\backslash s)/np >$
KICK_THE_BUCKET: KICK $\cup$
$prosody(argument(syntax(KICK\_THE\_BUCKET))) \approx the + bucket$

---

[15] See Daelemans and Gazdar (1992) for recent research and references.

[16] See van der Linden (in prep.) and similar representations in TAG (Abeillé 1990; Abeillé and Schabes 1989) and HPSG (Erbach 1991).

Semantics It follows from the definition of idioms that the meaning of the idiom cannot be inherited from the verb that is its head, but should be added non-monotonically. In (12) the extended representation of the semantics of *kick the bucket* is presented.

(12) b. KICK: $< (np\backslash s)/np, \lambda x \lambda y kick(x)(y) >$

   c. KICK_THE_BUCKET:
KICK $\cup$
$prosody(argument(syntax(KICK\_THE\_BUCKET))) \approx the + bucket \wedge$
$semantics(KICK\_THE\_BUCKET) \approx \lambda x \lambda y die(y)$

As in the model of the lexicon proposed by Zernik and Dyer (1987), the model proposed here puts the syntactic and semantic burden on the lexicon. Also, Zernik and Dyer relate idioms to their heads. Flickinger (1987) presents a hierarchical structure of the lexicon, but does not include idiomatic expressions.

Full specification of signs The full specification of a sign is derived by means of an operation similar to *priority union* (Kaplan 1987:180) or *default unification* (Bouma 1990). The specification operation ($\sqcap$; van der Linden 1992) is defined as a function from pairs of mother and daughter signs to fully specified daughter signs and runs as follows. If unification is successful for the values of a certain property of mother and daughter, the result of specification for that value is the result of unification where unification is understood in its most basic sense: variables unify with constants and variables; constants unify with variables and with constants with an equal value (prosodic information in (13)). If the values do not unify, the value of the daughter is returned (semantic information in (13)). [17]

(13) $(KICK \sqcap KICK\_TV) \sqcap KICK\_THE\_BUCKET$:
$< (np\backslash s)/ < np, the + bucket, \_ >, kick, \lambda x \lambda y die(y) >$

More specific information thus takes precedence over more general information. This is a common feature of inheritance systems, and is an application of the 'principle of priority to the instance' (Hudson 1980) which is acknowledged in knowledge representation and (computational) linguistics (De Smedt 1990; Daelemans 1987).

Not only can this principle be applied to the information that is part of mother and daughter signs, it can also be applied when a choice has to be made between a mother and a daughter sign. In the case of the choice between a literal interpretation of *kick the bucket* and an idiomatic interpretation the principle that the more specific information prevails, can be applied as well. Since the idiom

---

[17]The inheritance networks for which $\sqcap$ is defined are unipolar, non-monotonic and homogeneous (Touretzky et al. 1987). For other networks, other reasoning mechanisms are necessary to determine the properties of a node (Touretzky et al. 1987; Touretzky 1986; Veltman 1990).

inherits from the verb and is thus more specific, it is selected in the case of an ambiguous expression. Conventionality is modelled as a corelate of specificity. Van der Linden (1992) presents a categorial type-logical system which as an effect of the order of the logical rules gives precedence to the idiomatic interpretation over the literal interpretation in case of an ambiguity. The system is also set up in such a way that when the system has encountered only part of the idiom, for instance *kick* the idiom is not taken into account in the analysis. This only happens once all the relevant material that constitutes the idiom has been encountered.

5.5 Comparison of the models

The three models presented here are all able to model the conventionality principle. There are, however, a number of differences between them, that can be used to evaluate them.

- In the two-level model and the connectionist model, in the case of ambiguity the simplest hypothesis that covers the largest part of the input is preferred, and it is assumed that the largest part also constitutes the conventional interpretation. Although this is mostly the case, it does not necessarily have to be so. In the hierarchical model, conventionality is modeled by means of the specialization relation. Specialization seems to be more closely related to conventionality. In PHRAN (Wilensky and Arens 1980), specificity only plays a role in suggesting patterns that match the input, but evaluation takes place on the basis of *length*, and *order* of the patterns. Zernik and Dyer (1987) do not discuss ambiguity.

- In the two-level model conventionality has to be modeled explicitly. In the hierarchical model it is a consequence of the ordering of the rules in the system. In the connectionist model it follows from the architecture of the model.

- The hierarchical model is linguistically motivated, whereas the other models are merely models of the lexical retrieval process.

- The hierarchical model gives a less redundant representation of linguistic information. The other two models could, however, be extended with a hierarchical structure for the representation of syntactic and semantic information. An advantage of the hierachical model will be presented in the next section.

- A disadvantage of the connectionist model is the necessity for parallel processing: in the hierarchical model most processing takes place in serial order, and it therefore demands smaller processing capacity.

On the basis of these considerations, the hierarchical model seems to be the better of the three.

6 Syntactic flexibility

One of the differences between the hierarchical and the other models, is its linguistic motivation and the non-redundant representation of linguistic information. Therefore, the model is that it easily accounts for aspects of the syntactic behavior of idioms. This is the topic of the current section.

Idioms seem to deviate from their literal counterparts with respect to the syntactic constructions idioms can occur in. For instance (14) does not have an idiomatic interpretation.

(14)   # The bucket was kicked by John.

Most research on the flexibility of idioms has been devoted to explanations for this deviation, without firstly assessing the extent to which idioms differ from non-idiomatic expressions. The point to be made here is that for a considerable part idioms do *not* deviate from their literal counterparts: the syntactic flexibility of idioms can for a considerable part be explained in terms of properties of its verbal head, and this behavior can best be explained if the idiom is said to inherit these properties from its head. This thus supplies a further argument in favour of a hierarchical model of the lexicon. In order to illustrate this, the *passive* will be considered in detail here: non-passivizability of a large group of idioms can be explained in terms of properties of its verbal head.

6.1 Passive

Only transitive verbs occur in passive constructions (Bach 1980). [18] Bach mentions a number of classes in which verbs occur that seem to be transitive, but that are in fact complex intransitives, and therefore do not passivize. This classification seems to apply as well to idioms and explains why these do not passivize. A first rather trivial class are idioms which are already in passive form.

(15)   Van de aarde weggenomen worden.
        From the earth away_taken to_be.
        To be dying.

---
[18] For passivization of Dutch intransitives, see van der Linden (in prep.)

24

If the object of an idioms is a lexical reflexive, passivisation is not possible. Reflexivity includes reflexive pronouns and inalienable objects.

(16) Zijn beste beentje voorsetten.
His best leg_[dim] in_front_to_put.
Put one's best foot forward.

If the object of a verb is a lexically stipulated expletive pronoun, passivization is not possible.

(17) Hij sal het niet lang meer maken.
He will it not longer again make.
He will soon die.

The same applies to subjects.

(18) Het loopt af met hem.
It comes to_an_end with him.
He is dying.

Bach mentions a group of verbs that have objects that are no 'true' object NP's. Examples are predicative or copulative verbs, or verbs like *wegen* (*to weigh*) or *spelen* (*to act*).

(19) Hij speelt stommetje.
He plays dumb_[dim].
He keeps his mouth shut.

Verbs of possession are not transitive either

(20) Een bord voor de kop hebben.
A sign in_front_of the head to_have.
To be thick-skinned.

Although for a number of idioms it can thus be argued that the verb it comprises disables passivization, there is still a group of idioms for which this explanation does not do. In van der Linden (in prep.) it is shown that if these idioms do not passivize, this should still be attributed to them being intransitive, but that intransitivity is not inherited from the verbal head, but is a consequence of lack of underlying metaphorical properties: idioms that are neither isomorphic, nor motivated loose their transitivity and cannot be passivized.

Concluding remarks   There is a large group of idioms, the non-passivizability of which should be accounted for in terms of the non-transitivity of the verb that is the head of the expression. The most natural way to represent this, is by means of inheritance: the idioms inherits certain properties from its verbal head, which determine its syntactic flexibility.

7   Concluding remarks

Idioms have a non-literal interpretation that is lexically represented as a property of the expression as a whole, where parts of the expression may have metaphorical referents. As a model of the representation and processing of these expressions, a lexicon structure that is considered as an inheritance hierarchy seems the most viable, at least when the resolution of ambiguity and syntactic flexibility are concerned. When issues outside the scope of this paper are taken into consideration, the comparison becomes slightly different.

- Subsymbolic approaches can model more easily the *interactive* nature of natural language processing.

- With respect to *learning*, here learning idioms, it is clear from recent work in AI and cognitive psychology that distributed subsymbolic representations are promising. Algorithms for learning hierarchical structures exist. An underlying principle of inheritance, structure sharing, goes well with such distributed representations: inheritance hierarchies could be considered a linguistically sufficient generalization of an underlying subsymbolic representation. The symbolic model proposed by Zernik and Dyer (1987) for learning idioms only works in case of detection of a gap in lexical knowledge: bootstrapping in case of an empty lexicon is not possible.

- Upon failure of the principle of conventionality (in the end it is a heuristic) the hierarchical model provides an easy way to model *backtracking* by means of the choice of a different node in the hierarchical structure.

- It is unclear which model is best suited to model the metaphorical properties of idioms, motivation and isomorphism.

The fact that it is easy to model a principle of conventionality, could render the interpretation process of other forms of non-literal language efficient, and it is therefore worth to examine the scope of the principle.

References

ABEILLÉ, A. 1990. Lexical and syntactic rules in a tree adjoining grammar. Proceedings of the Conference of the Association for Computational Linguistics 1990. pp. 292-298.

ABEILLÉ, A. and Y. SCHABES 1989. Parsing idioms in Lexicalized TAGs. Proceedings of the 1989

conference of the European Chapter of the Association for Computational Linguistics 1989. pp. 1-9.

ARONOFF, M. 1976. Word formation in generative grammar. The MIT Press, Cambridge, Massechusetts.

BOBROW, S. and S. BELL 1973. On catching on to idiomatic expressions. Memory and Cognition 3. pp. 343-346.

BOUMA, G. 1990. Defaults in unification grammar. Proceedings of the 1990 Conference of the Association for Computational Linguistics 1990. pp. 165-172.

CACCIARI, C. and P. TABOSSI 1988. The comprehension of idioms. Journal of Memory and Language 27. pp. 668-683.

COTTRELL, G. 1989. A model of lexical access of ambiguous words. In: Lexical Ambiguity Resolution. *Edited by* Small, Cottrell and Tanenhaus. 1988. p. 179-194.

DAELEMANS, W. 1987. Studies in language technology: an object-oriented model of morphophonological aspects of Dutch. PhD-thesis. Univesity of Leuven.

DAELEMANS, W. and G. GAZDAR (Eds.) 1992. Inheritance in natural language processing. Special issue of Computational Linguistics. 18, 2.

DASCAL, M. 1987. Defending literal meaning, Cognitive Science 11, pp. 259-281.

DE SMEDT, K. 1990. Incremental sentence generation, PhD-thesis, University of Nijmegen.

DI SCIULLO A. and WILLIAMS 1987. On the definition of word, MIT press, Cambridge, Massachusetts.

ERBACH, G. 1991. Lexical representation of idioms, IWBS report 169, IBM TR-80.91-023, IBM, Germany.

EVERAERT, M. and E. van der LINDEN (Eds.) 1989. Proceedings of the First Tilburg Workshop on Idioms, Institute for Language Technology and AI.

EVERAERT, M., E. van der LINDEN, A. SCHENK and R. SCHREUDER (Eds.) 1992. Proceedings of IDIOMS, Tilburg, ITK, to appear.

FILLMORE Ch., P. KAY and M. OCONNOR 1988. Regularity and idiomaticity in grammatical constructions: the case of let alone. Language 64. pp. 510-538.

FLICKINGER, D. 1987. Lexical rules in the hierarchical lexicon. PhD-thesis, Stanford University.

FRASER, B. 1970. Idioms within a transformational grammar. Foundations of language 6, pp. 22-42.

GAZDAR, G., E. KLEIN, G. PULLUM and I. SAG 1985. Generalized Phrase Structure Grammar.

Basil Blackwell, Oxford.

GEERAERTS, D. 1992. Specialisation and reinterpretation in idioms. In: Everaert et al. (1992), to appear.

GIBBS, R. 1980. Spilling the beans on understanding and memory for idioms in conversation. Memory and Cognition 8. pp. 149-156.

GIBBS, R., and N. NAYAK 1989. Psycholinguistic studies on the syntactic behavior of idioms. Cognitive Psychology 21, pp. 100-138.

GODDARD, N, K. LYNNE, T. MINTZ, and L. BUKYS 1989. Rochester connectionist simulator. Technical Report. University of Rochester.

GROSS, M. 1984. Lexicon-grammar and the syntactic analysis of French. Proceedings of the International Conference on Computational Linguistics 1984, pp. 275-282.

HERINGER, J. 1976. Idioms and lexicalization in English. In: Syntax and Semantics 6: The grammar of causative constructions. *Edited by* M. Shibatani, Academic Press, New York, pp. 250-216.

HOCKET, Ch. 1958. A course in modern linguistics. Macmillan, New York.

KOLLER, W. 1977. Redensarten: linguistische Aspekte, Vorkommensanalysen, Sprachspiel. Tübingen, Niemeyer.

KOSKENNIEMI, K. 1983. Two-level morphology. PhD-thesis, University of Helsinki.

LAKOFF, G. 1987. Women, fire and dangerous things. Chicago University Press, Chicago.

LAKOFF, G., and M. JOHNSON 1980. Metaphors we live by. Chicago, University Press, Chicago.

LANCKER, D. van and G. CANTER 1981. Disambiguation of ditropic sentences: acoustic and phonetic cues. Journal of Speech and Hearing Research 24, pp. 64-69.

LANCKER, D. van, G. CANTER and D. TERBEEK 1981. Disambiguation of ditropic sentences: acoustic and phonetic cues. Journal of Speech and Hearing Research 24, pp. 330-335.

VAN DER LINDEN, E. 1989. Idioms and flexible categorial grammar. In: Everaert and van der Linden 1989, pp. 127-143.

VAN DER LINDEN, E. 1992. Incremental Processing and the Hierarchical Lexicon. Computational Linguistics 18, 2. Special issue on Inheritance and Natural Lanuage Processing. Edited by Daelemans, W. and Gazdar, G. (to appear).

VAN DER LINDEN, E. in prep. A categorial, computational theory of idioms. PhD-thesis, Tilburg University.

VAN DER LINDEN, E. and Kraaij, W. 1990. Ambiguity resolution and the retrieval of idioms: two approaches. Proceedings of Proceedings of the International Conference on Computational Linguistics 1990. vol 2 , pp. 245-251.

MARTIN, 1989, Representing and acquiring metaphor-based polysemy. Proceedings of the First International Lexical Acquisition Workshop, August 21, 1989, Detroit, Michigan.

NAPOLI, D. 1988. Subjects and external arguments, clauses and non-clauses. Linguistics and Philosophy 11, pp. 323-345.

RUMMELHART, D. and J. MCCLELLAND 1986. Parallel Distributed processing. Explorations in the microstucture of cognition. Volume 1: Foundations; Volume 2: psychological and biological models, MIT press, Cambridge, Massachusetts.

SCHENK, A. 1992. The syntactic behaviour of idioms. In: Everaert et al. 1992, to appear.

SCHWEIGERT, W. 1986. The comprehension of familiar and less familiar idioms. Journal of Psycholinguistic Research 15, pp. 33-45.

STOCK, O. 1989. Parsing with flexibility, dynamic strategies, and idioms in mind. Computational Linguistics 15, 1. pp. 1-19.

SWINNEY, D. 1981. Lexical processing during sentence comprehension: effects of higher order constraints and implications for representation. In: The cognitive representation of speech, *Edited by* T. Meyers, J. Laver and J. Anderson North-Holland.

SWINEY, D. and A. CUTLER 1979. The access and processing of idiomatic expressions. Journal of Verbal Learning and Verbal Behavior 18, pp. 523-534.

THIBADEAU, R., M. JUST and P. CARPENTER 1982. A model of the time course and content of reading. Cognitive Science 6. pp. 157-203.

WASOW, T, I. SAG and G. NUNBERG 1983. Idioms: an interim report. In: Proceedings of the XIIIth international congress of linguists. *Edited by* Shiro Hattori and Kazuko Inoue, Tokyo, CIPL. pp. 102-115.

WILENSKY, R., and Y. ARENS, 1980. PHRAN, A knowledge-based natural language understander. Proceedings of Conference of the Association for Computational Linguistics 1980. University of Pennsylvania, Philedelphia, Pennsylvania.

WOOD, M. MCGEE 1986. A definition of idiom. Masters thesis, University of Manchester 1981. Reproduced by the Indiana University Linguistics Club.

ZERNIK, U. 1987. Strategies in language acquisitions: learning phrases from examples in context. PhD-thesis, UCLA, Los Angeles.

ZERNIK, U. and M. DYER, 1987. The self-extending phrasal lexicon. Computational Linguistics. 13, pp. 308-327.

List of Figures

List of Tables

Appendix

The connectionist model for the retrieval of idioms as presented in section 5.3 is based on the mechanism of interactive activation and competition (IAC). An ideal IAC network consists of nodes that can take on continuous values between a minimum and a maximum. The activation of the units is also supposed to change only gradually in time. This ideal is approximated by dividing time into a series of small steps. If we choose an activation function that cannot change very rapidly this discrete model acts as a good approximization for the ideal IAC-network.

The network (Figure (2)) consists of a set of nodes that are connected with links which can be excitatory or inhibitory (with a negative weight value). Some units can receive external stimuli, e.g. input from the syntactic module. The internal structure of a unit is shown in Figure (3). The input links are connected to a site that corresponds to their type. So each unit has distinct sites for external, excitatory and inhibitory links. The gate unit also offers a separate *gate* site with a special site function.

The site functions for the external, excitatory and inhibitory links simply compute the weighted sum of the input values Iv.

$$Sv = \sum_{i=1}^{n} w_i Iv_i$$

The site function for the gate site is a kind of "weighted AND" function. Its behaviour is similar to the weighted sum function when all input links have a value different from zero. However if one of the input links connected to the *gate* site is zero, the output Sv of the gate site function is also zero. The output of each site is scaled in order to control the influence of the different sites on the activation value.

$$Netinput = Sc_{inh} Sv_{inh} + Sc_{exc} Sv_{exc}$$

$$+ Sc_{ext} Sv_{ext} + Sc_{gate} Sv_{gate}$$

The activation value Av for a new timestamp t can now be computed:

When *Netinput* is larger than zero:

$$Av^t = Av^{t-1} + (max - Av^{t-1}) Netinput$$

$$- decay(Av^{t-1} - rest)$$

32

When *Netinput* is less than zero:

$$Av^t = Av^{t-1} + (Av^{t-1} - min)Netinput$$

$$-decay(Av^{t-1} - rest)$$

We see that the influence of *Netinput* on ΔAv decreases when Av reaches its minimum or maximum value. On the other hand the influence of the decay rate is high in the upper and lower regions. When *Netinput* becomes z ero, the Activation value slowly decreases to its rest value. The output value of the unit is equal to its activation, but only if the activat ion level is above a predefined treshold value. Otherwise the output is zero. So a unit with maximum activation that does not receive input anymore, slowly decreases its output value and than suddenly drops to zero beacuse its activation is below treshold value. This non linear behaviour is an essential property of connectionist models.

The bottom-up links are stronger than the top-down links because a unit may only be activated by bottom-up evidence. Top-down information may however influence the decision process at a lower level.

The values of the parameters in the model are:

Sc_{inh}	0.6
Sc_{exc}	0.6
Sc_{ext}	0.6
Sc_{gate}	0.6
treshold	0.5
decay	0.1
bottom-up weights	0.8
top-down weights	0.25
inhibitory weights	-0.8
external input weights	1.0
max	1.0
min	-1.0
rest	0

A simulation consists of a number of cycles in which activation spreads through the network. In each cycle the output and activation values for a time t are calculated from the values on time *t-1*. Figure (3a) and (3b) show the activation levels of the active units in the model: only activation levels above treshold (500) are displayed. At the beginning of the simulation all units are in rest state. We start the simulation for the disambiguation of "kick (the) bucket" by setting the output value of the

external unit "kick" representing the output of a sub wordform level to 1. After three update cycles, the output of the external unit II (representing the fact that bucket is recognized) is set to 1. The duration of an external input is always one cycle. The availability of syntactic information is simulated by activating III_b and III_c before cycle seven. Figure (3a) shows that the unit representing kick as a verb immediately follows this syntactic information and "kick as a noun" falls beneath activation treshold. After some more cycles a stable situation is reached (Figure (3b)) which represents the best fitting hypothesis: the idiomatic reading.

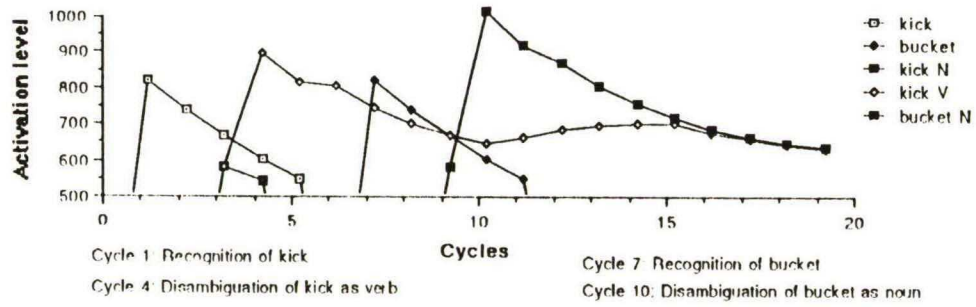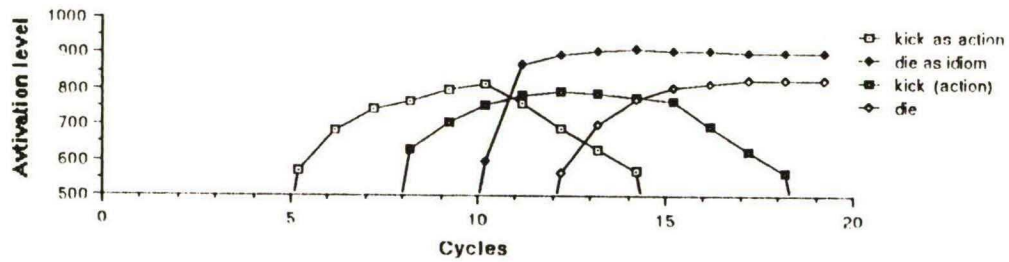**Fig (3a) Activation level of the wordform and syntactic units**

Cycle 1: Recognition of kick
Cycle 4: Disambiguation of kick as verb

Cycle 7: Recognition of bucket
Cycle 10: Disambiguation of bucket as noun

Fig (3b) Activation level of the semantic units and the semantic network

ITK: P.O. BOX 90153 5000 LE TILBURG THE NETHERLANDS