How can metaphors be interpreted cross-linguistically?

Yorick Wilks[1]

Abstract. Research on metaphor as a phenomenon amenable to the techniques of computational linguistics received a substantial boost from a recent US government (the iARPA agency) funding initiative that set up a number of teams in major universities to address the issues of metaphor detection and interpretation on a large scale in text. Part of the stated goal of the project was to detect linguistic metaphors (LMs) computationally in texts in four languages and map them all to a single set of conceptual metaphors (CMs). Much of the inspiration for this funding was the classic work (Lakoff and Johnson, 1980) which posited a set of universal metaphors used across cultures and languages.

I wish to examine the assumptions behind this goal and in particular to address the issue of how and in what representation such CMs can be expressed. I shall argue that a naïve approach to this issue is to make very much the same assumptions as the work of Schank and others in the 1970s (including the present author): namely that there can be a universal language of "primitives" for the expression of meaning, which in practice always turns out to be a form of simple English (or in the case of Schank, atoms like PTRANS, very close to English words). In none of those system was the sense ambiguity of the English-like terms every tackled in a systematic way (though see: Guo 1989). Reviving that assumption for the study of metaphor raises additional issues since, even if the *senses* of the terms in those CM representations could be added, by annotation from a standard lexicon for the CM representations, metaphors often considered to deploy new senses of words which will not be found in existing sense inventories like computational lexicons which, if true, might make such annotation impossible (though later in the paper I shall argue against just that novel deployment of sense in metaphor). This paper is not intended just to present a negative conclusion; I also argue that the representation of metaphors in a range of languages can be brought together within some CM scheme, but that simply reviving the *English-as-interlingua* assumptions of forty years ago is not a good way to make progress in this most difficult area of meaning computation.

In what follows I first discuss first the representation of CMs and ask: in what language are they stated? I argue the need for some inclusion in the representation of the senses of their constituent terms within the CM, or at least a default assumption that the major sense (with respect to some lexicon such as WordNet) is the intended one. I then consider the issue of conventional metaphor and its representation in established lexicons (again such as WordNet) and [1]the effect that can have on detection strategies for metaphor, such as selectional preference breaking. I then argue that the mapping of text metaphors to CMs, as well as the empirical, rather than intuitive, construction of CM inventories requires further use of preference restrictions in lexicons by means of a much-discussed process called projection or coercion. I conclude that only the use of (computable) procedures such as these for metaphor detection and mapping can lead to a plausible program for the large-scale analysis of metaphor in text, and that Lakoff's views on metaphor lack these empirical underpinnings.

[1] Florida Institute of Human and Machine Cognition, 15 SE Osceola, Ocala FL 34471. Email: ywilks@ihmc.us

1 INTRODUCTION

Understanding prose in any natural language rests first on it being in a language one understands, let us say English for the purposes of this paper. But problems in understanding arise even for native speakers of English as well as with translations, human or mechanical, from other languages. One way of capturing the additional understanding needed that goes "beyond knowing the words and the grammar" is expressed by the term "metaphor". This notion conveniently expresses aspects of culture and figurative expression that go beyond literal or ostensive meaning and are crucial to understanding. These phenomena are sometimes opaque even to those who are experts in the language concerned. Metaphor also has the advantage that it has been an area of research in computer language processing for decades, and one that has yielded real results. That research has been driven in part by the writings of George Lakoff at Berkeley [1] who has developed an approach to metaphor that rests on the following assumptions (in my terms, but I think fairly uncontentious):

- There are similar metaphors found in all cultures that are crucial to understanding language.
- These metaphors can be discovered and listed, even if not exhaustively.
- We can proceed with analysis as if these metaphors can be not only paraphrased but expressed in English.

For example, such a universal metaphor might be expressed (in English) as LIFE IS A JOURNEY and we shall refer to items like this as Conceptual Metaphors (CM). There is then an initial analytic question of how to detect metaphors in text, possibly related to or "expressing" that CM such as *The pensioner was nearing the end of his road*. After locating this sentence as a metaphor there is then the task of matching it to such a stored generalized CM form. We shall refer to linguistic strings like the one in italics as Linguistic Metaphors (LM), There may then be the problem, if one believes in the universal nature of CMs, of how to locate expressions of "similar" metaphors in, say, Farsi to that same CM. The capitalised words in the English form of the CM may themselves have many senses and the question

immediately arises as to how an algorithm is to determine which sense is intended by "LIFE" in that CM: that it is not, say, a "a life as in a children's game of hide and seek, a score token".

One problem with metaphor research, at least from a computational or Natural Langauge Processing (NLP) perspective, is that universal theories like the one above (expressed by the three bullets) have proved resistant to computational implementation, which has not been the case with other, quite different, empirical approaches based on bottom-up detection of LMs in text (e.g. [3], [4]), rather than starting from a set of a priori CMs. We shall now turn to questions about *the representational language in which CMs are stated* and how they to be intuitively understood, since their terms (e.g. LIFE) do not disambiguate themselves

2. THE LANGUAGE OF CONCEPTUAL METAPHORS (CMs)

I shall argue that a crucial aspect of the research problem, which many seem to believe is a solution, is that CMs are classically expressed in English words but without any realization of what that entails. When this is pointed out, a frequent response is that this is an accidental fact of no significance and we can just carry on since though they appear to be English words they are not, but rather some form of symbol outside ordinary natural language. I believe this is profoundly inadequate response. It is in fact a recrudescence of the early discussions in AI and NLP in the 1960s and 1970s on the role of interlinguas in machine translation and in cognitive representations generally. There was a fashion at that time for limited languages (expressed by English primitives terms) within systems for the semantic representation of language content (e.g. in the work of Schank [5]; Wilks, [6] and many others). I am not here defendingthat approach, only pointing out that the extended discussion forty years ago (e.g. in [7]) of the adequacy or otherwise of this limited language of (English-like) primitives to carry the general meaning of language expressions has many similarities to what we are discussing now, nearly fifty years later, in regard to CMs.

There was no real resolution to that controversy of long ago: key references are Pulman's [8] attack on the practice from a linguistic perspective, and Lewis [9]from a philosophical one, in the course of which Lewis invented the term "markerese" for the self-description of language in linguistics (e.g. by Fodor and Katz, [10]) by means of word-like *markers* with no illumination or benefit. But the critiques were not heeded and much such representational work continued, simply because researchers in semantics could see no alternative (outside radical connectionism) to continuing to use symbols to represent the meanings of other symbols. Montague [11] was a philosopher who reacted against markerese but his representations of mean, although more replete with logical forms that those of Fodor and Katz, still were expressed in symbols including English-like words, though now usually expressed in lower case and with an apostrophe attached. Language content had to represented somehow, theorists reasoned, so why not in this English-like language? Dictionaries, after all, describe word meanings using the very language they describe, and so the practice has continued, ignoring the waves of philosophical and linguistic criticism, simply because there seemed to be no alternative. What has happened is that the language terms used for representation have been embedded in more logical and formal-seeming structures so as to make them palatable, but the underlying issue has not gone away. That issue is: How can I describe semantic content with a term such as MAN, HUMAN or ANIMATE and be confident I know what it means, and not just "means in English"? I shall now turn to how problems of CM representation problems can be ameliorated with the aid of a sense-lexicon.

3. REPRESENTING CMs UNAMBIGUOUSLY WITH MAJOR WORD SENSES

If we are to use CMs at all, no matter how derived or expressed, they must be in as word-sense-neutral a form as we can manage. To my knowledge this has never yet been fully considered as problem, perhaps an insurmountable problem, let alone a solved problem. We cannot just ignore this as we do when we say, for example, that [POVERTY IS A GAP] is a CM, and underlies the metaphor "poverty gap", and that we just know what the senses of the words in the CM are present in that expression and that they make up a CM. Just suppose that we had two CMs in our inventory of universal metaphors that could be written as:

POVERTY IS A GAP

POVERTY IS AN ABYSS

Now suppose we want to locate Russian metaphors and find the text string (LM) containing the keywords : *бедность провал*, which mean roughly "poverty" and "failure". But, and here is the problem "*провал*" can also means "abyss" and "gap" in English; in which case how do we know which of these two so-called universal CMs to match the Russian LM to? Or should we seek for or construct a third CM [POVERTY IS FAILURE]? It seems clear to me that either:

1) The CMs are in some language other than English, in which case how do we know what English word senses the terms above correspond to, since the English words "poverty", "failure" and "abyss" may all have multiple senses in, say, WordNet [12]. If, however, the terms are not English but some universal language of indeterminate syntax and semantics, how can LMs ever be matched to CMs as any serious theory of metaphor seems to require?
2) If however, the terms in the two CMs above *are* in English, and they certainly appear to be, then we need to know what senses those words have in those particular forms, so as to match any word in an English or Russian LM to them.

A natural way of carrying out the requirement in (2) is to tag the English words in the CMs (and the words in any putative LMs) with WordNet senses. Since the EuroWordNet project [12] in which the present author participated, we now have a convenient

way of setting up such a match since that project took the core Princeton WordNet for English as, essentially, an interlingua, and linked senses in the Wordnets for other languages to those core senses. So, for example (and the correctness of these correspondences does not matter for the argument): there may well be an English WordNet sense of "failure", namely failure#1 that is deemed by a EuroWordNet mapping to be the same sense as Провал#1 in the Russian WordNet. Again, there may be a "Провал#3" that similarly corresponds to "abyss#1".

What do we want to say about universal CMs and their ability to support the analysis of metaphor instances in such a case? The first natural thing to say---given the above WordNet assumptions---- is that the original Russian string "*бедность провал"* can express both CMs and we cannot decide which. But that is only true if we cannot decide which sense the last word bears in the Russian LM. If it bears only one of the two noted senses then the Russian LM matches one and only one of the CMs—assuming now the CM terms are tagged with WordNet senses. Russianists should note here that I am ignoring the case issues for the proper expression of that string in Russian and just concentrating on the main forms of the words. Also, I am not suggesting it would be problematic if a LM were to match to two possible CMs, though I do not believe that need be the case here. It could be that other, perhaps pragmatic, factors outside the text would settle the choice. My only point here is that a systematic empirical account of mapping LMs to CMs should take account of this possibility and standard contemporary metaphor theories do not consider the issue at all.

Now a Russian speaker may take that (LM) phrase to have one and only one of those senses in context—assuming the Russian speaker can understand the distinction we are making with the words "failure" and "abyss" in English—let us assume they can, even though the string may be too short and vague for a wordsense disambiguation program to determine the sense in that LM context.

Or, and this is a quite different possibility, is it the case that, in a metaphorical string such as the LM "Poverty is failure" we cannot rely on the normal psychological or computational methods to resolve a word sense for us. Since the content is, more or less, novel, at least on first encounter, the standard disambiguation techniques may well not work because they are all, to some extent, based on redundancy, which does not apply to novel utterances? So, to use an old and hackneyed example, if someone says *The shepherd swung his crook*, we infer that "crook" is a tool for shepherds not a gangster, simply because of the redundant presence of "shepherd". But in LMs this may not be available, unless the metaphor is dead, or lexicalized or otherwise familiar (in which case wordsense disambiguation hardly applies). What I am suggesting is that perhaps in metaphors, especially novel ones, the words must be taken in their basic senses by default, as it were, *because in a metaphor we lack the familiar context to resolve a participating word to any non-basic sense.*

This conclusion is perhaps not very striking but rather obvious: words of a real language, like English, can only function in an interlingua (such as CMs constitute) on condition that they bear their "basic" senses, which will, in WordNet terms, usually mean #1 for any given word. This implies that in the capitalized English CMs above, each term implicitly has whatever its #1 sense is in WordNet.

So to return to the purported sense correspondence in Euroword-net style:

failure#1 is deemed by a EuroWordNet mapping to be the same sense as Провал#1. Again, there may in addition be a "Провал#3" that similarly corresponds to "abyss#1".

This line of reasoning would imply that we should take the CMs (and LMs, with the caveat above) in their default #1 senses, since we have no information to allow us to do anything else. Hence "Провал" should be taken in the context above to be Провал#1, its first sense, and so as a CM about failure not about an abyss, even thought the latter could conceivably be indicated by another context for the same words. This suggestion that the senses in a CM are major senses of the relevant words also implies that the two CMs above are different from each other, which preserves the insight of the tradition that metaphors are strictly speaking lies (attributed variously to Mark Twain, Nietzsche et al.) rather than the less acceptable alternative that CMs are tautologies, where the constituent senses simply recapitulate each other.

This risk of tautology in the expression of CMs is very real even if we are wary and assign (implicitly as main senses) interpretations to the symbols in CMs. If, in the CM [POVERTY IS A GAP], we allow the first WordNet sense interpretation to "gap" we get:

S: (n) **gap**, spread (a conspicuous disparity or difference as between two figures) "gap between income and outgo"; "the spread between lending and borrowing costs"

Thus, and depending on the sense assigned to "poverty", we have a very real risk of tautology since this sense of "gap" is itself abstract (and not, say, a gap between two pieces of wood) and itself very close to any definition of poverty, or at least "relative poverty" the currently fashionable version. This unfortunate fact can be dismissed, or simply accepted as a weakness or error in WordNet, or, perhaps, as a reason for excluding [POVERTY IS A GAP] as a CM.

One important inference from this discussion, if it has any value, is that we cannot just say, as many researchers in the Berkleyan universal metaphor tradition seem to want to, that some particular metaphor "in one language" is commoner than in another. As we have seen, it is a very sophisticated matter to establish whether LMs in two languages point to a single CM or not, given the problems of how any CM is to be unambiguously represented and, given the need for some lexical resource of at least the size and scope of (Euro)WordNet in order to do that. In the example above, the LM word strings in question in the two languages—Russian and English ---actually point to different CMs in the common interlingua, a conclusion that, we argued, undermines the foundation of the Berkeley approach to understanding metaphor, since the LMs could clearly be interpreted as "meaning the same thing". At this point, let us step

back and review the basic role of "preference" in detecting, then mapping, metaphors.

4. THE ROLE OF PREFERENCE IN DETECTING AND MATCHING METAPHORS

An exception to the "rule of main senses" we have just stated, as far as LMs are concerned, is the situation we have defined elsewhere as one of "conventional metaphor" [13] This is where a lexical resource such as WordNet actually encodes a metaphorical sense as a (dead or) conventional metaphor. Our approach to detecting metaphor has been that an initial *sufficient* criterion for a surface (LM) metaphor to be present is that a verb or adjective "preference" is broken [6] e.g. in the simplest case the verb does not receive the agent or object it expects (whether that last notion is unpacked linguistically or statistically) in a stereotypical case. Verbs and adjectives will, of course, have multiple senses in the lexicon, each with its own preferences. So to write *fall into poverty* is to break the preference for a spatial-container-like object for the basic sense of "fall into". This general criterion reappears frequently in the literature (e.g. the recent work of Shutova [4]) indeed it is not clear there is any alternative to it as a basic criterion for metaphor recognition, unless one believes that metaphors are detected by direct matching to stored CMs. As we have seen above this a notion whose very intellegibilitys dissolves somewhat under scrutiny.

If such preferences, and the associated noun-senses for fillers, are thought of as stored in a respository like WordNet or VerbNet, then what counts as a broken preference depends crucially on the state of lexicon at a given time, since sense inventories extend with time and indeed often come to store senses that were in origin metaphorical. Where that is the case, a dead, or as we would prefer to say conventional, metaphor will not result in a broken preference with respect to WordNet because in such a case the metaphorical sense is itself stored in WordNet and so will fit the demands of the corresponding verb.

So, to take a very simple and uncontentious example:

Public employees' unions have built a fortress around their pension systems

In VerbNet [14] we find the following:

[[VerbNet: build

Member of

§build%2:31:03 (member of VN class base-97.1)

§build-26.1-1

•WordNet Sense 1

•Agent [+animate | +machine]

So "<u>Unions" violates Agent restriction for build</u>

•WordNet Sense 8

•Agent [+animate | +organization]

<u>"Unions" satisfies the Agent restriction ---as an organization—for build</u>]]

The situation is one where the primary sense of "build" is not satisfied by the first sense of the agent the sentence contains but is satisfied by a "lower' (in this case #8) sense. In [13] I proposed that this could serve as a useful heuristic (i.e. main sense failure but some lower sense a successful match) for detecting conventionalized metaphors of the sort this sentence contains, since such metaphors would be missed by any "preference breaking" heuristic for metaphor detection as there is a (lower) sense of "build" available for which the agent preference here is satisfied. The heuristic was that a main sense fails and a lower sense satisfies; and both parts must be true. Its main defect is that it relies on the ordering of senses in WordNet as carrying information, which is generally true but as always with this database has many errors and omissions.

The point here is not to draw attention to this metaphor detection heuristic against a large lexicon for its own sake, but only to show a limitation on the earlier suggestion that metaphor detection (and as we shall discuss below, metaphor mapping to CMs) must depend on the main senses, as listed in a lexicon. Our claim here is that this heuristic for detecting conventional or lexicalized metaphor does not compromise the general value of that rule. In the case of the above example, there are arguably two CM metaphors present: the major one is to do with barriers and the protection of assets, however expressed, and the other is more simply (and even though it is, more strictly, a meronym, though such differences are not crucial here):

ORGANIZATIONS ARE PEOPLE

which is expressed (in major senses of the relevant words) by the process of detection we have described.

The latter move is the basis of how preferences, and their violations in metaphor, are also central to the subsequent process of mapping from a detected metaphor to some stored form, which we are calling CMs. If we were again dealing with "He fell into poverty" we might expect the broken preference for the object of "fall into" to be some coding for hole/abyss/gap/aperture. The inference from that detection to the underlying metaphor in play is generally to assert that the metaphor's object (poverty in this case) is being asserted to be equivalent to the preferred filler that is made available in the lexical coding (e.g. in VerbNet, see [14]) but not in the sentence itself. This would lead directly to some form such as:

POVERTY IS AN ABYSS

as a potential CM, empirically derived from this example text rather than a linguist's intuition. The interesting difficulty is to determine at exactly what level its last term is to be expressed,

since “abyss” is, in general, a very magnified form of hole. The mapping process from a metaphor instance, or LM, to a CM, however expressed, will require an ontology of the kind that underlies WordNet to navigate from what appears in a VerbNet coding (perhaps “hole”) to an item in an already stored CM (perhaps, as here, “abyss”). This method, merely sketched here, can in principle serve to map LMs to CMs, and to create potential CMs from text.

This process, making use of the preferred constituents of lexical codings, has been central to a number of systems based on inferences within lexical semantic structures and under names such as “projection” and “coercion” (e.g. Wilks, [6]; Pustejovsky,[15]; Nirenburg and Raskin, [16] and Hanks [17]) among many others. It provides at least the beginning of a process of determinate empirical construction of CMs from text cases quite different from the intuitive creation of CMs in the Berkeley tradition. Moreover, [22] contains a sophisticated analysis of some of the cross-lingual issues raised here. Further possible examples of the method would be with a failed subject+verb preference in *Israel has inflicted this wound on itself*. There we can get (from the stored VerbNet subject preference for “inflict” as PERSON) we can link the existing target (Israel) to the preferred subject (as source), namely PERSON, and then the WordNet type of “Israel” as COUNTRY to give as a possible CM: COUNTRY IS PERSON. We could do the same for verb+object failure as in: *The bank hyenas are feeding on money*, assuming we have access to “feed on” as a verb with its own preferences FOOD or EDIBLES. Then, using similar reasoning to that for subjects above, and again combining the assigned object and the preferred object, we can derive directly a potential CM: MONEY IS FOOD. For adjective+noun preferences, similar processes are possible, as in *Brazil’s economic muscle will become increasingly important*. If we have a preference established for the preferred type of noun associated with the adjective “economic” as COMPLEX-SYSTEM, then from the existing adjective object “muscle” (and taking its semantic type from WordNet as BODY) we then have directly a CM: COMPLEX-SYSTEM IS BODY. Many metaphor theorists would want to argue that equations of target and source CMs produced by a process such as this must be brought under some higher level generalization on both sides of the assertion in the CM, as we shall now show.

Notice though that no claims here depend on the actual quality or completeness of resources such asVerbNet or WordNet. These are always variable, depending on the language used, and will always contain errors and omissions, as well as being constantly changing with the language itself. The only claim is that some such resource will be needed to carry out the processes described here, even if augmented in practice by statistical corpus computations (some of which augmented these resources in the work described in [13]).

There has been criticism of processes of this sort applied to the empirical construction of CMs in this manner: during a recent large-scale metaphor detection and interpretation project a project manager wrote:

“[CMs that were] proposed….. were inconsistent and generally unmotivated. For the most part, the relationship of an LM (for a Target) and a proposed CM was semantically extremely shallow with generally no mapping at all. This process caused a huge proliferation of “lexical” CMs, often dependent on a synset label from WordNet.”[18]

It is odd, in the current empirical climate, to criticise a linguistic process for being grounded in data, rather than linguistic intuition. One must also respond (a) that there is no known correct *level* for the expression of CMs beyond the intuitions of metaphor theorists, so no level is demonstrably “too lexical” and (b) more fundamentally, the CMs are inevitably in some language (usually English) and require sense disambiguation of their terms, as we argued at length above. They are not in a language that is self-disambiguating, since nothing is. Hence the presence of WordNet labels, even if implicit, so as to indicate main senses as we suggested above, is inevitable. That would be a feature not a bug.

The problems of the appropriate level for the expression of CMs, their distance and separation from LMs and their very origins in intuition, are not ones that preoccupy only NLP researchers, as is clear from Deignan’s:
“.... at some points in the development of CMT [Conceptual Metaphor Theory], there has been a tendency for researchers to propose new conceptual metaphors using limited linguistic evidence. For instance, [19] take the idioms "he really couldn't swallow it" and "[leave] a bad taste in the mouth" as instantiations of a conceptual metaphor termed ACCEPTING SOMETHING IS EATING IT. It is not clear how many other realizations there might be of this conceptual metaphor, and in what way it differs from the more-often cited IDEAS ARE FOOD. Kovecses [20] lists as a conceptual metaphor CONSIDERING IS CHEWING, which again is difficult to separate from IDEAS ARE FOOD. If this tendency becomes widespread, the notion of a conceptual metaphor loses clarity, along with any predictive power it may have had.” ([21] p.105)

I take the force of this comment, from a corpus linguistic standpoint, to be consistent with the NLP processing critique advanced in this paper, and indeed with the internal project critique quoted earlier above. However, there is a difference of emphasis here: Deignan argues that CMT theorists in fact make up CMs from data, no matter what they say about intuition, and I have argued that they should be constructed by a determinate process from data since there is no other reliable route. But the internal project critique earlier seems to say that derivation from data in any such way is a istake and leads to shallow CMs and “real” CMs come only from intuition. I hope I have set out reasons for thinking tis comment profoundly wrong and out of line with all modern thinking on linguistics and data.

5. THE LAKOFF BERKELEY VIEW OF METAPHOR REVISITED

This view, against which I have argued, seems to me to rest on the following, very questionable, assumptions:

1. There is a set of universal CMs, determinable by linguistic intuition and underlying all languages.

There is no suggestion this set should be small, even fixed, as Schankian primitives were once held to be, and certainly some

depend on developments in technology, economics etc. Yet, as I have argued, there is no empirical evidence for their existence or how many of them there are, and intuition as a source of linguistic insight is no longer considered reliable, taken alone. However, there may be a discovery procedure for them from text along the lines suggested here (and in [6]).

2. CMs can be expressed in an English-like language, whatever their real underlying representation.

I have argued that they are in fact in English, as they appear to be, and not as an inevitable approximation; this is made clear by the problem of expressing exactly what senses their constituent words are to be taken in. This situation is only tolerable as a heuristic if some form of cross-lingual sense representation is incorporated into the representation, as suggested here.

3. Surface metaphors (LMs) in languages can be mapped to these CMs in a determinate way.

I have argued that no definitive procedure is ever given, within this tradition, for performing this crucial step and it can only be attempted at all with the aid of some fairly reliable, cross-sense mapping of the languages concerned, such as (Euro)WordNet.

If LMs can be matched bottom-up to CMs in something like the way sketched here---as opposed to being the subject of some direct matching top-down from stored CMs to LMs in text---- it should be possible to count how many LMs correspond to a given CM. That would then make it possible to estimate the frequency of occurrence of CMs in a reliable manner. That analysis could be extended cross-lingually and cross-culturally if parallel text were available. Suppose we had an English-Spanish parallel text in which sentences are aligned. We could then ask whether LMs are detected in parallel (putatively synonymous) sentences and, if so, do they map to the same CMs. If they do, that would be independent confirmation of the utility or universality of such a CM. Quantitative and distributional questions about universal metaphor can only be asked, it seems to me, if procedures of this kind I sketch here are developed, but these are not obviously compatible with standard Lakoffian approaches to metaphor, though there is no reason in principle, or course, why it could not develop so as to incorporate some empirical theory of sense ambiguity like the present one.

My main conclusion is that, for these reasons, Berkeley metaphor theory cannot easily be the basis of an empirical exploration of metaphors in texts in multiple languages, and that any research program aimed at the interpretation and translation of metaphor instances so based will have been mistaken.

ACKNOWLEDGEMENTS

The paper is indebted to comments from Patrick Hanks, Robert Hoffman, Brian MacWhinney, Jaime Carbonnel, Tomas By and Sergei Nirenburg, though the errors are all mine as always.

This research was supported by the Intelligence Advanced Research Projects Activity (IARPA) via Department of Defense US Army Research Laboratory contract number W911NF-12-C-0020. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon. Disclaimer: The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of IARPA, DoD/ARL, or the U.S. Government.

REFERENCES

[1] Lakoff, G., and Johnson, M. 1980 *Metaphors we live by*. University of Chicago Press: Chicago.

[2] Guo, C-M. 1989. *Constructing a machine-tractable dictionary from "longman dictionary of contemporary English*. Doctoral dissertation, Computer Science, New Mexico State University.

[3] Fass, D. and Wilks, Y. 1983. Preference semantics, ill-formedness and metaphor. *Journal of Computational Linguistics,* 9, pp.1798-187.

[4] Shutova, E., Teufel, S. and Korhonen, A. 2012. *Statistical Metaphor Processing*, Computational Linguistics, 39(2).

[5] Schank, R (Ed.,) 1975. *Conceptual Information Processing*. Elsevier: Amsterdam.

[6] Wilks, Y., 1968/2007. Making preferences more active. Reprinted in Ahmad, Brewster and Stevenson (Eds.) *Word and Intelligence I.* Springer: Berlin.

[7] Wilks, Y. 1977/2007. Good and bad arguments for semantic primitives. Reprinted in Ahmad, Brewster and Stevenson (Eds.) *Word and Intelligence I.* Springer: Berlin.

[8]Pulman, S. 1983. *Word meaning and belief*. Croom Helm: London.

[9] Lewis, D., 1972. General Semantics, In: Davidson and Harman (eds.), Semantics of natural language. Reidel: Dordrecht.

[10] Fodor, J. and Katz, J. 1963. The structure of a semantic theory. *Language*, *39*(2), Apr-Jun, 170-210.

[11] Montague, R. 1974. *Formal philosophy : selected papers of Richard Montague* / ed. and with an introd. by Richmond H. Thomason. New Haven: Yale Univ. Press.

[12] Vossen, P. (ed.) 1998. *EuroWordNet: a multilingual database with lexical semantic networks*. Kluwer: Amsterdam.

[13] Wilks, Y., Dalton, A., Allen., J.. Galescu, L. 2013. Automatic Metaphor Detection using Large-Scale Lexical Resources and Conventional Metaphor Extraction. *Proc.1st Workshop on Metaphor in NLP (Meta4NLP 2013).* Atlanta, GA.

[14] Windisch Brown, S., Dligach, D., and Palmer, M. 2011 VerbNet Class Assignment as a WSD Task. In IWSC 2011:

Proceedings of the 9th International Conference on Computational Semantics, January 12 - 14, 2011, Oxford, UK.

[15] Pustejovsky, J. 1995. *The Generative Lexicon.* MIT Press: Cambridge, MA.

[16] Nirenburg, S. and Raskin, V. 2004. *Ontological semantics.* MIT Press: Cambridge, MA.
[17] Hanks, P. 2013. *Lexical analysis: norms and exploitations*. MIT Press: Cambridge, MA.

[18] iARPA: http://www.iarpa.gov/Programs/ia/Metaphor/metaphor.html

[19] Gibbs, R., Bogdonovich, J., Sykes, J., & Barr, D. 1997. Metaphor in idiom comprehension. *Journal of Memory and Language*, 37, pp. 141-154.

[20] Kovecses, Z. 2002. *Metaphor: a practical introduction*. Oxford University Press: Oxford.

[21] Deignan, A. 2005. Metaphor and corpus linguistics. Benjamins: Amsterdam.

[22] Prior, A., Wintner, S., MacWhinney, B., & Lavie, A. (2011). Translation ambiguity in and out of context. *Applied Psycholinguistics, 32*, 93-111.