Natural Language and Knowledge Representation

Diana Santos
INESC, R. Alves Redol,9
P-1000 Lisboa, Portugal
dms@inesc.inesc.pt

INESC Report nr. AR/5-92

1 Introduction

The objective of this paper is to call the attention of knowledge engineers to natural languages, with two distinct kinds of arguments. The first is that the most natural way of encoding, transmitting and reasoning about knowledge is through natural language (NL), and thus repositories of knowledge are almost always in a written form. The second kind of arguments is what one could call "reasoning by analogy": Drawing from work in natural language processing, linguistics, both theoretical and computational, and studies in the philosophy in language, I shall present several cases where a knowledge representation (KR) problem or solution has a close parallel in NL processing (NLP).

My main claim is that if there is a lot to be gained if studies of natural language are not left to linguists alone, NL researchers also have a significant contribution to make as far as KR is concerned. More precisely, knowledge engineering could (and should) give a hand in (at least) most practical tasks that computational linguists (or NLP workers) are involved in (e.g., learning from corpora, generalizing knowledge from examples, defining ontologies and hierarchies for linguistic objects, etc.), but there is a significant insight to be gained for KR if deep studies of NL and ways of dealing with it are taken into consideration.

As a by-product of this paper, I shall point at some actual cases where linguistics and AI knowledge interact, in a short survey of the modern work in the two fields.

2 Interdependency between knowledge and language

2.1 NL as a KR formalism

Natural language is so far the most comprehensive tool for (humans to) encode and reason with knowledge. Thus, it could be viewed as the best knowledge representation formalism available. (See [9] and [28] for a similar point). Intuitively, to master a non-trivial language is one of the more outstanding features of an intelligent being as such.

However, it is rather difficult (to say the least) to build a system with similar capabilities as those demonstrated by the average NL user, and thus most KR work aims at achieving a simpler formalism, as far as the reasoning capabilities and the phenomena describable are concerned, so that a practical implementation can be carried out (see [19] on the tradeoff between expressive power and tractability).

While much of the knowledge that can be described by these artificial formalisms is "translated" or "explained" in the form of NL (patterns of) sentences that the system is (or not) able to deal with (this is particularly clear for example for the non-monotonic reasoning and belief revision camps), it is surprising how few KR researchers have started by studying NL. Especially because NL could constitute an important area of validation of AI progress.

On the contrary, it is often the case that NL is just mentioned only to be discarded by a blunt "our aim is not to model the way NL handles this", instead of being treated as the paradigmatic case. (See e.g. [22, p.108, Note 1] on temporal ontology.) This is most striking, it seems to me, in the development of ontologies for representation of a particular domain. My claim is that while theoretically there is no bound to the partition of a concept space into smaller and smaller units of meaning, NL could yield a pratical boundary if the smaller units or concepts present (in an observable way) in NL were studied in detail.

To be fair, it should be said that not all outstanding researchers in KR share my views of NL. For instance, [16] requires that KR should possess at least the expressive power of natural language (thus implying that it could and should go beyond it). But of course also from this philosophical standpoint, the study of NL properties is a must. On the other hand, [14] is very clear concerning the differences between NL and KR formalisms: In page 78, he states

> *I distinguish sharply between a natural language which is informal and probably not even completely defined means of communication in the real world (...) and a formal deductive scheme for representing knowledge.*

And later he goes on *Natural language is for communication, the internal language*

of thought is for thinking (...) These are vastly different requirements and so the languages can be expected to be very different.[15, page 18]

Incidently, in the same paper (same page) Hayes mentions that looking at the meanings of English spatial prepositions can be useful, thus reinforcing my claim that, no matter one's theoretical viewpoints, to look at NL as a knowledge source is the right thing to do.

Note: One could object to the "NL as KR formalism" argument by claiming that knowledge representation includes not only declarative knowledge but also procedural, i.e., the ways to process and handle such knowledge, and for those, NL does not seem to matter. However, this objection seems to me ill-founded in two respects: 1) NL understanding does require that the inferences illicited by words and structures be studied; 2) There are procedural descriptions in NL (cookbooks, assembly manuals, military orders); which implies that NL is not only important for the declarative part of knowledge storing, but also that it is the language for communication of many other kinds of knowledge.

2.2 Is there knowledge independent of language?

Another point that should be mentioned in this connection is the fact that different languages have different realisations of their lexical fields and different languages may implicitly encode different conceptions of the world. (This is usually referred to as Sapir-Whorf hypothesis[6], actualized recently in [1]). This leads us to conclude that a KR system is not guaranteed to be independent on the mother tongue of the researchers who build it.

A very common example is the distinction animated subject / agent or not. Germanic languages encode it grammatically, while Romance don't. Definiteness is also a concept that is not clear to speakers of languages without articles, while, on the contrary, an average English speaker (without linguistic training) has no clear feeling about aspect.

This fact is particularly clear for machine translation researchers, where it has been claimed that there are meaning distinctions that can only be perceived through other languages [30].

3 Strong similarities between the two domains

3.1 Compositionality

One of the most pervasive phenomena in natural language is the existence of idiomatic expressions, which are roughly defined as those complex linguistic objects (phrases, words formed by two or more identifiable units, etc.) whose meaning is not compositional, in the sense that by knowing the meaning of its parts you cannot arrive at or deduce the meaning of the whole. Typical examples are "the

coast is clear", "to pull strings", "to take the bull by the horns", "to kick the bucket", "to help oneself", "red herring", "software", "false friends". The degree of idiomaticity varies from expression to expression; moreover, it is not true that these phrases have only the idiomatic reading. E.g., one can kick a bucket if s/he likes football and you should not infer that doing this action means instantaneous death.

In knowledge representation in general, it is easy to devise situations where knowledge about, say, baking and cakes does not result in knowledge about how to bake a cake. Or, if a person knows simultaneously how to drive, and how to jump on skis, it does not follow automatically that one can drive in a ski jump. This proves that knowledge is also not compositional in a strict sense, and therefore, ways of organizing idiomatic versus non-idiomatic uses of language in NLP may be useful in the KR domain as well.

3.2 Context sensitivity 1: time and place of utterance

Natural languages are highly sensitive to the context, in several ways. First of all, clearly they can (and do) refer to non-linguistic objects - the referential use of natural language ("Mary", "the king of France"). Then, some words (generally belonging to a closed set) can also refer to parameters of the utterance context: deictics [20] or indexicals, and they can be classified in two groups [18]: pure — e.g. "I", "here", "this", "now"—, or demonstratives, such as "she", "this", "that". Demonstratives are so called because they depend both on the context of the utterance and on an accompanying demonstration (be it a gesture, or previous text). [26] gives yet a more refined picture of these phenomena.

The knowledge engineer should also be interested in at least a partial indexicality of the concepts he is manipulating: They introduce the relationship to an external changing world, and it is not far-fetched to imagine that, for the actual connection of a knowledge based system and the real world, mechanisms parallelling the several kinds of NL indexicals could be necessary.

To be concrete, consider the example of temporal reasoning systems. It is clearly important to be able to state and process differently "Today the temperature is 19 degrees Celsius" and "You cannot change things that occurred before today". In the first case, "today" could and should be replaced by a particular date, but, in the second, "today" represents a true indexical, having a different "value" every day.

3.3 Context sensitivity 2: intentions

NL objects may also change the meaning drastically depending on the intention (and social context) of the utterance. Generally this is considered to be a matter of pragmatics rather than semantics; nonetheless, its correct interpretation is essential to an error-free communication.

This has long been acknowledged in the AI community, and attempts at user modelling and of taking into account the goals of the participants in an intelligent interaction have been pursued for a long time. However, linguistic information should also not be forgotten, e.g. the abundance of linguistic clues for topic/focus relations [13, 12, 7], and intonation information in the study of spoken discourse [29].

3.4 Context sensitivity 3: vagueness

Finally, one problem common to formal analyses of NL and KR, which I will exemplify with two different examples, is that of vagueness of a word or concept. That is, many words depend in a more subtle way on the (linguistic and non-linguistic) context.

Many qualifiers, for instance, are relative to some "fuzzy" notion of average in the domain they qualify (see e.g. [17]), and this clearly has consequences for the subsequent reasoning: for "a big ant" is no contradiction to the common statement that "ants are small".

It is also natural to use "when" for non-exactly overlapping times: "When they built the 39th Street bridge, a local architect drew up the plans / they solved most of their traffic problems" (taken from [24]). In other words, while the informal meaning of "when" could be said to be "at the same time", the times described in the two clauses can be sequenced either way.

3.5 Different levels for different uses

Dictionary developers [10] know that the very same (objective) level of representation (in this case, words) may need to be conceptually partitioned into several levels of use. A very practical example is a word processor with style-critiquing capabilities. For obvious reasons, the spelling checking part has to be based on a list of only correct word forms. The style critiquing function, however, also needs information about common mispelled word forms in order to provide accurate error detection andto furnish correct alternatives. It is clear that the two sets of words cannot be merged, especially because in many instances the same string can be a correct word and a mispelled one. E.g., "it's" is a correct contraction of "it is" and a rather common mispelling of the pronoun "its".

Another clear example is the need to distinguish an inflected verb form from the verb representative (e.g., its "lemma form"): LAY as the lemma form of the verb inflected "lays", "laid", etc. should have a totally different information associated from the one ascribed to the past form of the verb LIE... This should come as no surprise to the KR researcher, who is familiar with the distinction between terminological and assertional knowledge ([21]).

In addition, it has been beautifully shown in [3] that human beings are inconsistent, or better, they possess different sets of consistent beliefs which may or

may not be consistent among them. Moreover, it seems that, again, those sets of beliefs are dependent on the contexts of their use. It is thus acknowledged that no global consistency can model an intelligent agent (see also [25]). Models of inconsistent sets of knowledge about the "same" object are therefore common to both NLP and KR in general.

4 Conclusion

I have presented two lines of argument of why NL should be a main concern for KR researchers, the first philosophical, the second essentially practical. I want now to survey the literature either where related points are made or work using both linguistic and AI knowledge is described.

The claim that AI researchers should learn from linguistics is not new (see [27] for an excellent defence of it) Generally, however, it concerns only the restricted case where AI researchers are using NL. There are nevertheless noteworthy exceptions:

For knowledge representation purposes, the following examples used linguistic theories and/or data: [11] used a theory of natural language meaning as a KR-formalism. [31] (following [2]) use natural language repositories (such as dictionaries and thesauri) to represent background knowledge, presenting moreover a formal representation of this concept for NL understanding.

For linguistic purposes, the following examples use KR techniques (on purpose, no other AI areas will be surveyed): [23, 32] use formalisms with hierarchical ontologies and inheritance inspired by AI work. More particularly, the HPSG (Head-driven Phrase-Structure Grammar) school is investing in KR-formalisms in order to produce the grammar theory's inheritance module (see [5]), while [8] uses classification-based KR systems to understand NL instructions. Last but not least, non-monotonic logics have recently started to be used for linguistic description, see on this respect the papers in [4].

Finally, an historical remark is in place: many, if not most, of the early KR formalisms had natural language understanding as goal, as so notable researchers as Quillian, Woods, S.Shapiro, Schank and Winograd have stated!

Still, I think it is just fair to say that the typical KR researcher and the typical NL researcher never think about each other in the (albeit so similar) problems they have to face everyday.

5 Acknowledgements

I thank JNICT for granting me a PhD scholarship, and IBM Portugal for the general support of my studies. This paper has benefitted from discussions with Jan Engh and José Carlos Medeiros, and also from an anonymous referee of the

ERCIM workshop on Knowledge Representation held in Pisa in May 1992, in which proceedings this paper appears in an extended abstract format.

References

[1] Bach, Emmon. "On Time, Tense, and Aspect: An Essay in English Metaphysics", in Peter Cole (ed.), *Radical Pragmatics*, Academic Press, 1981, pps.63-81.

[2] Binot, Jean-Louis & Karen Jensen. "A Semantic Expert Using an Online Standard Dictionary", IBM Research Report RC 12455, T.J.Watson Research Center, Yorktown Heights, 21 Jan 1987.

[3] Birnbaum, Lawrence. "Rigor Mortis: a response to Nilsson's "Logic and artificial intelligence"", *Artificial Intelligence*, 47, 1991, pp.57-77.

[4] Calder, Jo & Michael Morreau. Reader in "Nonmonotonic Logic in Linguistic Description", The Third European Summer School in Language, Logic and Information, Saarbruecken, August 12-23, 1991.

[5] Carpenter, Bob & Carl Pollard. "Inclusion, disjointness and choice: the logic of linguistic classification", in *Proceedings of the 29th Annual Meeting of the ACL*, 1991, pp.9-16.

[6] Carroll, J.B. *Language, Thought and Reality: Selected Writings of Benjamin Lee Whorf*, MIT Press, 1956.

[7] Davison, Alice. "Syntactic markedness and the definition of sentence topic", *Language*, Volume 60, Number 4, 1984, pp.797-969.

[8] Di Eugenio, Barbara. "Action representation for NL instructions", in *Proceedings of the 29th Annual Meeting of the ACL*, 1991, pp.333-4.

[9] Ejerhed, Eva & Lars-Erik Janlert. "Representing Time in a Natural Language Processing System", Department of General Linguistics, Publication 25, Umea 1985.

[10] Engh, Jan. "Normalisation in Language Industry: Some Normative and Descriptive Aspects of Dictionary Development", unpublished ms, IBM Norway, Oslo, 1991.

[11] Guenthner, Franz, Hubert Lehmann & Wolfgang Schönfeld. "A theory for the representation of knowledge", *IBM J. Res. Develop.*, Vol 30, No.1, January 1986, pp.39-56.

[12] Hajicova, Eva & Petr Sgall. "From Topic and Focus of a Sentence to Linking in a Text", in B.G. Bara and G. Guida (eds.), *Computational Models of Natural Language Processing*, Elsevier (North-Holland), 1984.

[13] Halliday, M.A.K. *An Introduction to Functional Grammar*, Edward Arnold, 1985.

[14] Hayes, Patrick. "Some Problems and Non-Problems in Representation Theory", *Proc. AISB Summer Conference*, Univ. Sussex, 1974, pp.63-79.

[15] Hayes, Patrick J. "The Second Naive Physics Manifesto", in J.R. Hobbs & R.C. Moore (eds.), *Formal Theories of the Commonsense World*, Ablex Publishing Corp., 1985, pp.1-36.

[16] Hirst, Graeme. "Ontological assumptions on knowledge representation", in Brachman, Levesque and Reiter (eds.), *Proceedings of the First International Conference on Principles of Knowledge Representation and Reasoning KR'89*, Toronto, Ontario, Canada, May 15-18, 1989, pp.157-69.

[17] Kamp, Hans. "Events, Instants and Temporal Reference", in Bäuerle, R., Egli, U. and A. von Stechow, *Semantics from Different Points of View*, Springer-Verlag Berlin Heidelberg 1979, pps.376-417.

[18] Kaplan, D. "Demonstratives: An Essay on the Semantics, Logic, Metaphysics and Epistemology of Demonstratives and Other Indexicals", Draft No2, UCLA xerox, 1977.

[19] Levesque, Hector J. & Ronald J. Brachman. "A Fundamental Tradeoff in Knowledge Representation and Reasoning (Revised Version)", in Brachman, Ronald J. & Hector J. Levesque, (eds.), *Readings in Knowledge Representation*, Morgan Kaufmann Pubs, 1985, pp.42-70.

[20] Lyons, John. *Introduction to Theoretical Linguistics*, Cambridge University Press, 1968.

[21] Mac Gregor, Robert. "The evolving technology of classification-based knowledge representation systems", in John Sowa (ed.), *Principles of Semantic Networks: Explorations in the Representation of Knowledge*, Morgan Kaufman, 1991, pp.385-400.

[22] McDermott, Drew. "A temporal Logic for Reasoning About Processes and Plans", *Cognitive Science 6*, 1982, pp.101-55.

[23] Moens, Marc, Jo Calder, Ewan Klein, Mike Reape & Henk Zeevat. "Expressing generalizations in unification-based formalisms", in *Proceedings of the Fourth Conference of the European Chapter of the ACL*, 10-12 April 1989, pp.174-81.

[24] Moens, Marc and Mark Steedman. "Temporal Ontology and Temporal Reference", in *Computational Linguistics*, Vol 14, Number 2, June 1988, pp.15-28.

[25] Nilsson, Nils J. "Logic and artificial intelligence", *Artificial Intelligence*, 47, 1991, pp.31-56.

[26] Oberlander, Jon Reid. "The Semantics of Temporal Indexicals", PhD thesis, University of Edinburgh, 1987.

[27] Sag, Ivan A. "Linguistic Theory and Natural Language Processing", in Klein, E. & F.Veltman, *Natural Language and Speech: Symposium Proceedings, Brussels, November 1991*, Springer-Verlag, 1991, pp.69-83.

[28] Schubert, Klaus. "Implicitness as a guiding principle in machine translation", *Proceedings of COLING'88*, 1988, pp.599-601.

[29] Steedman, Mark. "Surface Structure, Intonation and "Focus"", in Klein, E. & F.Veltman, *Natural Language and Speech: Symposium Proceedings, Brussels, November 1991*, Springer-Verlag, 1991, pp.21-38.

[30] Tsujii, Jun-ichi. "Future Directions of Machine Translation", *Proceedings of COLING'86*, pp.655-68.

[31] Zadrozny, Wlodek & Karen Jensen. "Semantics of Paragraphs", *Computational Linguistics*, Volume 17, Number 2, 1991, pp.171-209.

[32] Zajac, Rémi. "A Transfer Model Using a Typed Feature Structure Rewriting System with Inheritance", in *Proceedings of the 27th Annual Meeting of the ACL*, 1989, pp.1-6.