MACHINE READABLE DICTIONARIES: WHAT HAVE WE LEARNED, WHERE DO WE GO?

Nancy Ide & Jean Véronis

Laboratoire Parole et Langage
CNRS & Université de Provence
29, Avenue Robert Schuman
13621 Aix-en-Provence Cedex 1 (France)

Department of Computer Science
Vassar College
Poughkeepsie, New York 12601 (U.S.A.)

e-mail: {*ide,veronis*}*@univ-aix.fr*

Abstract

Machine-readable versions of everyday dictionaries have been seen as a likely source of information for use in natural language processing because they contain an enormous amount of lexical and semantic knowledge. However, after fifteen years of research, the results appear to be disappointing. No comprehensive evaluation of machine-readable dictionaries (MRDs) as a knowledge source has been made to date, although this is necessary to determine what, if anything, can be gained from MRD research. To this end, this paper provides an overview and assessment of MRD research to date. It then proposes possible future directions and applications that may exploit these years of effort, in the light of current directions in not only NLP research, but also fields such as lexicography and electronic publishing.

1. Introduction

The need for robust lexical and semantic information to assist in realistic natural language processing (NLP) applications is well known. Machine-readable dictionaries (MRDs) have been seen as a likely source of information for use in NLP because they contain an enormous amount of lexical and semantic knowledge collected together over years of effort by lexicographers. Considerable research has been devoted to devising methods to automatically extract this information from dictionaries (see, for instance, [1-11]). Early results in MRD research were promising and led many to feel that large knowledge bases could easily be derived automatically from MRDs (see, for instance, the position papers by Amsler and Boguraev at TINLAP-3 in 1987 [12, 13]). As a result, several large-scale projects and centers (for example, the Electronic Dictionary Research Institute in Japan, ACQUILEX in Europe, and the Consortium for Lexical Research in the U.S.) were established to forward research in the area.

Despite early expectations, fifteen years of work on knowledge extraction from MRDs has made it clear that the information they contain is both too inconsistent and incomplete to provide a ready-made source of comprehensive lexical knowledge. It is now widely acknowledged that automatic knowledge acquisition must necessarily rely on a variety of sources (see, for example, [14]), and research interest has turned toward other sources of linguistic knowledge, most notably corpora . However, there is a danger that the decline in interest in MRDs as a knowledge source will result in the value and contribution of both past and potential research in the area to be ignored. In order to avoid throwing out the baby with the bath water, an assessment of MRD research seems appropriate at this time.

Although MRD research failed to live up to early expectations, the past fifteen years of research have in fact contributed in various ways both to the original goals concerned with constructing knowledge bases, as well as other, unforeseen goals. In particular, and often as a result of the failures in extracting knowledge from MRDs, MRD research has contributed significantly to our understanding of the nature, kinds, and role of semantic information in the processing of natural languages. It has also led to an understanding and formal description of the structure of dictionaries and lexical information in general. In addition, the analysis of the information in dictionaries from the point of view of what is needed for NLP has contributed to lexicographers' own understanding of dictionaries and the dictionary-making process.

We recognize a convergence of interests and goals between NLP and lexicographers and electronic publishers (cf. [15,16]), which may result in a benefit to all of them. A most promising avenue of future activity in MRD research involves collaboration between these communities, resulting in both better dictionaries and exciting new possibilities for electronic, hypertextual dictionary databases, as well as the creation of lexical databases immediately useful for NLP.This paper first examines MRD research and provides an assessment of its value and contribution. We then propose possible future directions and applications that may exploit these years of effort, in the light of current directions in not only NLP research, but also fields such as lexicography and electronic publishing.

2. Assessment of MRD research

Despite early success in extracting information from MRDs, the results of early studies did not scale up to enable fully automatic extraction of knowledge bases from MRDs. In fact, the previous ten or fifteen years of work on fully automatic extraction has produced little more than a handful of limited and imperfect taxonomies. This section summarizes the main reasons why this is so.

2.1. Extraction of information from MRDs is difficult

The work of Chodorow *et al.* [3] and others made the extraction of semantic information from MRDs appear simple. However, claims of high success rates (often, 98% etc.) were misleading, since "success" (in Chodorow *et al.*'s case) meant finding the *head* of a definitions. By far the greatest success was semantic information proved to be much more problematic due to far greater inconsistencies in the ways it was specified, often demanding relatively sophisticated parsing of the definition text (see [17, 18,19]).

Conversion form original formats. MRDs typically come to researchers in unusable formats--most commonly, in the form of typesetter tapes from publishers. To make the MRD usable for research, considerable effort was often required, and in fact the translation of MRDs in typesetter format to something more usable has become an area of study in itself. As a result of ambiguities and inconsistencies in typesetter formats, parsing typesetter tapes requires developing a complex grammar of entries (see, for example, [20]). Even with this, problems still arise because conventions are inconsistent. For example, in the *CED*, the entry *Canopic jar, urn* or *vase* must be interpreted as (*Canopic jar*) or (*Canopic urn*) or (*Canopic vase*), whereas the entry *Junggar Pendi, Dzungaria,* or *Zungaria,* which has the same structure, must be interpreted as (*Junggar Pendi*) or (*Dzungaria*) or (*Zungaria*). Because of the inconsistency, fully automated procedures cannot determine the appropriate interpretations. Most of this work is far outside the realm of NLP research, and in general it is time-consuming and without great intellectual interest. As a result, only a handful of dictionaries are available in a usable format, mainly in English. And now that the magnitude of the task is obvious, researchers may be reluctant to start similar work on dictionaries in other languages.

Inconsistencies in definition format. After the dictionaries are cleaned up, their definitions must be parsed. It is well-known that variations in definition *metatext* (that is, phrases in definition texts that express semantic relations, such as "used in *V*-ing" for instrument, "consisting of a *N*" for parts, etc.) render this process difficult. For example, the definitions in figure 1 demonstrate that the metatextual phrases signalling particular relations ("having a", "with a", "consisting of", etc.) are often inconsistent and cannot be detected with simple Chodorow-like pattern matching techniques alone. Bearing in mind that this example shows only a single case (handle) over a handful of definitions in a single dictionary, it is clear that the amount of work required to determine the possible variants could easily exceed that required to construct the corresponding knowledge base by hand.

jug	a vessel...usually <u>having</u> a handle...
kettle	a metal container <u>with</u> a handle...
ladle	a <u>long-handled</u> spoon...
corkscrew	a device...<u>consisting of</u> a pointed metal spiral <u>attached to</u> a handle...
fork	a small usually metal implement <u>consisting of</u> two, three, or four long thin prongs <u>on the end of</u> a handle...
knife	a cutting instrument <u>consisting of</u> a sharp-edged often pointed blade of metal <u>fitted into</u> a handle...
basket	a container...often <u>carried by means of</u> a handle or handles.

Figure 1. Definitions for objects with a *handle* in the *CED*

The bootstrapping problem. A common strategy in MRD research to use sophisticated syntactic parsers (for example, the Linguistic String Parser as in [21]) to analyze definition texts. But even for the simple examples given above, more than syntactic analysis is required. For example, to determine that a handle needed to differentiate "carried by means of a handle" from, say, "cooked by means of steam". Similarly, in the definition of *ladle*, semantic information would be required in order to determine that the adjective "long-handled" specifies a part of a ladle. In some cases, the resources of a full knowledge base are required to understand the specification--for example, to differentiate "carried by means of a handle" from "carried by means of a wagon", world knowledge is required to determine that "wagon" is not a part of the object carried, although "handle" most likely is. This process is clearly circular, and has led some researchers to attempt MRD analysis using only minimal pre-existing resources that can be constructed by hand and bootstrapping as MRD analysis proceeds. However, it is becoming clear that the difficulties are so considerable that such methods are unlikely to succeed; indeed, to date, none has been convincingly demonstrated.

2.2. The information in MRDs is partial

Inconsistencies. Information in definition texts of dictionaries often varies considerably in terms of content, and consequently, in terms of the kinds of knowledge that can be extracted. This happens because dictionaries are typically the product of several lexicographers' efforts and is constructed, revised, and updated over many years, inconsistencies in the criteria for constructing definition texts necessarily evolve. In addition, space and readability restrictions as well as syntactic restrictions on phrasing may dictate that certain information is unspecified or left to be implied by other parts of the definition.

Ide and Véronis [7] show that even in a small and straightforward sample of extracted information (detection of hypernyms for concrete objects, kitchen utensils) 50-70% of the information is missing or garbled in five major English dictionaries. Morever, the specification of hypernyms is certainly more consistent in dictionaries than that of other semantic relations (e.g., parts, shape, color, smell, etc.), which is given in a much more random way. For example, in the definitions for *abricot* (apricot) and *pêche* (peach) from three major French dictionaries in figure 2, all of the dictionaries indicate that the pit of a peach is hard, but none specifies this property for the pit of the apricot (which is certainly just as hard). Information is given inconsistently even within a given dictionary.

	MICRO-ROBERT	HACHETTE	PETIT LAROUSSE
abricot	Fruit de l'abricotier, à noyau, à chair et peau jaune orangé.	Fruit de l'abricotier, d'une saveur délicate et parfumée, de couleur jaune rosé.	Fruit de l'abricotier, à noyau lisse, à peau et chair jaunes.
pêche	Fruit du pêcher, à noyau très dur et à chair fine.	Fruit comestible du pêcher, au noyau dur, à la chair jaune ou blanche, tendre et sucrée, à la peau rose et duveteuse.	Fruit comestible du pêcher, à chair juteuse et à noyau dur.

Figure 2. Definitions for *abricot* (apricot) and *pêche* (peach)

Lexical vs. other types of knowledge. It is generally acknowledged that some types of information that or world knowledge. For instance, it is interesting to note that there is no direct connection drawn between *lawn* and *house,* or between *ash* and *tobacco* in the *Collins English Dictionary,* although it is clear that this connection is part of human experience. The sense disambiguation strategy described in [22], which applies a connectionnist approach and relies on such connections between words in definition texts, fails in these cases as a result.

3. Have we wasted our time?

It appears that the information in MRDs is probably not consistent or complete enough to provide a substantial basis for a knowledge base. In addition, given that the amount of work required to extract the information that does exist in MRDs has now been demonstrated to be extensive (even requiring the existence of the knowledge bases it is intended to create), the return on investment is clearly very low.

We therefore must ask ourselves: does the past fifteen years of MRD research constitute wasted effort?

3.1. The need for systematic assessment

Before any assessment of the MRD research can be made, we need to know far more about the nature and quality of information in MRDs. Although it is generally known that information is inconsistent and sometimes missing in MRDs, apart from the small and preliminary studies cited above [7], there has been no comprehensive and systematic assessment of the information that is contained in MRDs. Perhaps more importantly, it is impossible to know if information extracted from MRDs is usable for NLP without some assessment of exactly what kinds of knowledge are needed for various NLP tasks. However, no systematic and comprehensive assessment of the kind of knowledge needed in relation to particular NLP tasks (language understanding, translation, etc.) is available.

To provide a simple example, a pervasive and virtually ignored problem in information extracted from MRDs is the attachment of terms too high in automatically extracted hierarchies, which occurs in 21-34% of the definitions in the sample from the five dictionaries cited in [7]. For example, while *pan* and *bottle* are *vessels* in the *CED, cup* and *bowl* are simply *containers,* the hypernym of *vessel.* Because of inconsistencies such as these, semantic networks extracted from different dictionaries look very different. Figure 3 also demonstrates that the problem of too high attachment is particularly pronounced in the *LDOCE,* likely due to the fact that this dictionary uses a controlled vocabulary. This dictionary is among the most widely used in NLP research, and a number of taxonomies have been built by extracting hypernyms from its definition texts. Without some assessment of the level of specificity required in hierarchies in order to be usable for various NLP tasks, it is not possible to determine if too high attachment renders extracted hierarchies unusable or minimally usable for NLP.

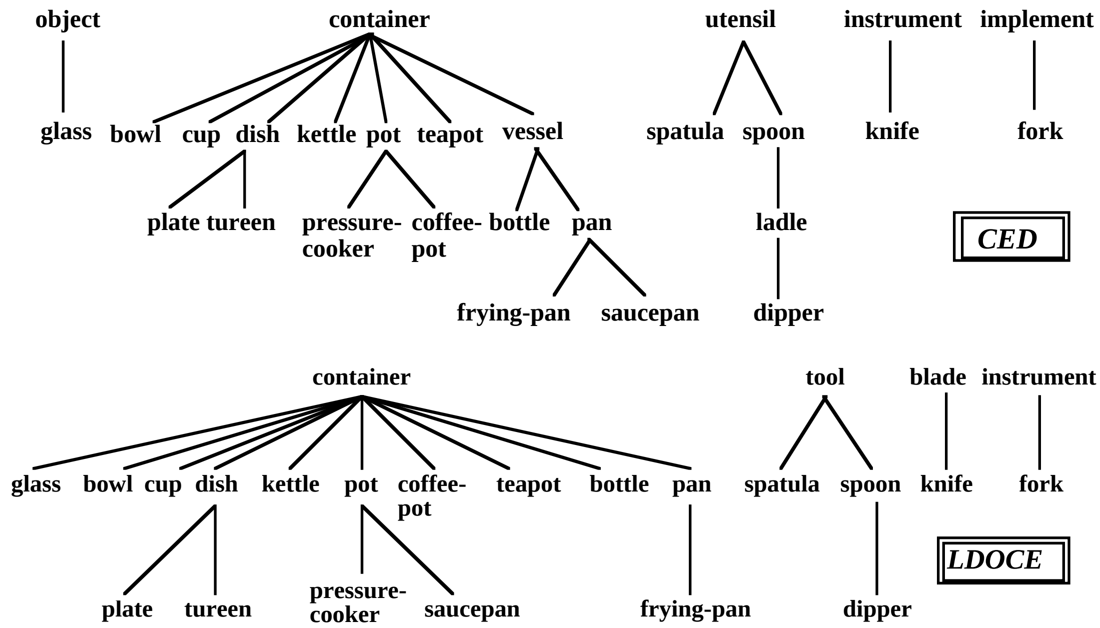

Figure 3. Hierarchies from different dictionaries

Similarly, conceptual divisions in many knowledge bases which rely wholly or partly on MRDs as a source of information are often made on the basis of sense distinctions made within those dictionaries. However, it is not at all clear that the distinctions made for the purposes of lexicography can serve the needs of NLP. For instance, Kilgarriff [23] shows that in a sample of 83 words from the LOB corpus, 69 had at least two senses in LDOCE and were therefore ambiguous. Kilgarriff found that for 60 of these 69 words (about 87%), there was at least one usage in the LOB corpus which could not with any confidence be classified into a single sense. This occurred because, for example, more than one sense was near the meaning of the word as used in the corpus, or because no sense given in the dictionary applied. MRD research has assumed for the most part that sense distinctions in dictionaries correspond to sense distinctions that apply in actual use, and therefore could provide the conceptual divisions that should appear in a knowledge base. However, apart from distinctions between homographs, it is not clear that this assumption holds. The differences in the level of detail and, occasionally, in the ways lines are drawn between senses when one moves from one dictionary to another, already show that sense distinctions in dictionaries are not definitive; studies such as Kilgarriff's bring this fact into focus in relation to real language use, which is obviously the ultimate concern of NLP.

A first step, then, in order to assess the value of MRD research is to systematically assess NLP needs and evaluate the information extracted from MRDs against these needs. We urge this kind of assessment not only for MRDs, but also for other knowledge sources such as language corpora.

3.2. Dictionaries are still among the best repositories of knowledge available

Most of MRD research was directed toward *automatic* extraction of knowledge, based on an underlying postulate MRD research, according to which *large knowledge bases cannot be built by hand.* It is not at lexicographers have been creating knowledge bases by hand for over 200 years. If this postulate is abandoned, we can more realistically assess the potential contribution of MRD research. It is clear that MRDs contain useful data--but most of it is probably usable only with possibly substantial by-hand massaging, requiring human judgement to be incorporated into useful knowledge bases.

In addition, it is the trend in NLP research (possibly in part because of the experience with MRDs) to consider that no single source (dictionaries, corpora, etc.) can provide all or even most of the knowledge required for NLP. Therefore, it is now widely recognized that knowledge base construction requires combining information from multiple sources, especially information provided by corpus analysis, since corpora can provide information such as common collocates, proper nouns, role preference information, frequency of use and similar statistics, etc. However, with corpora as with MRDs, fully automatic extraction is not likely, and it is again unclear what corpora can provide and how valuable the information is for NLP. The assessment called for the in the previous section would provide a valuable basis for work in this area.

Another possibility is to combine information from several dictionaries, since although information derived from individual dictionaries suffers from incompleteness, it is extremely unlikely that the same information is consistently missing from all dictionaries. Therefore, information from several dictionaries can be used to fill in information which is missing or faulty in one or more others. For example, in a small experiment, merging multiple dictionaries produced highly encouraging results: in a merged hierarchy created from five English dictionaries, the percentage of problematic cases was reduced from 55-70% to around 5% [7]. By-hand work is still required, but merging can substantially improve the quality of the extracted information.

The creation of knowledge bases in the future will most likely be accomplished by giving the human knowledge-base-creator access to multiple resources, including MRDs and corpora, together with tools to extract different kinds of information and combine it more or less by-hand. This is, for example, the approach taken by the Japanese EDR Electronic Dictionary project; in addition, projects to develop workstations for this kind of work are underway. This information will be widely varied in kind, including both detailed linguistic information as well as statistics, associational links, etc.

3.3. The structure of dictionaries is much better known

MRD research has necessarily involved considerable work on rendering dictionaries into a usable format for extraction of information. To this end, an encoding format for MRDs has been developed under the aegis of the Text Encoding Initiative (see [25, 26, 27, 28]) which can be applied across mono- and bi-lingual western dictionaries. Such a format must necessarily be both general enough to be applicable across different dictionaries, whose structures often vary widely, and at the same time capture the fundamental structural principles (e.g., hierarchical structure, factoring of information) that underlie dictionaries. A common encoding format enables the application of common software and hence the reusability of MRDs, and is extremely useful in the publishing industry for rendering in-house data in common formats which are directly suitable for typesetting, generation of dictionaries in different forms (e.g., concise, learner's), etc.

The development of an encoding format suitable for MRDs demands identification of the dictionary entry's constituent elements as well as aa deep understanding of the structural principles underlying dictionaries. Thus a related problem is the determination of a database model suitable for representing the information in MRDs. Lexical data, as is obvious in any dictionary entry, is much more complex than the kind of data (suppliers and parts, employees' records, etc.) that has provided the impetus for most database research. Therefore, classical data models (for example, relational models) do not apply very well to lexical data, although several attempts have been made (see for example [6, 29]). Ide, Véronis, and Le Maitre [30] have proposed an alternative feature-based model for lexical databases, which allows for a full representation of sense nesting and defines an inheritance mechanism that enables the elimination of redundant information. The model has been implemented in an object-oriented DBMS [31]. This and other similar work continues to feed the development of database models to represent lexical data and textual information, which is becoming an increasingly active area of research in database design.

4. MRD research and the future of the dictionary

One of the most promising possibilities for the future of MRD research results from a merging of interests among NLP researchers, lexicographers, and electronic publishers. Lexicographers, possibly as a result of MRD research, are increasingly interested in exploiting computer techniques, often for creating lexical data bases containing the kinds of information that NLP research had hoped to extract from MRDs. Some lexicographers are explicitly concerned with creating NLP-like knowledge bases [15, 16], which will in turn provide NLP with more of the resources it needs. However, the most influential force changing traditional lexicography is the advent of media such as CD ROM and handheld devices, which make possible an entirely new kind of dictionary with hypertext facilities, interactive display, etc. Furthermore, the development of the "information superhighway" provides vast potential for new ways to manage, distribute, and update dictionaries and other resources, such as corpora, to which they can be related. As a result, future dictionaries will likely be very similar to linguistic workstations, and provide many of the same facilities. MRD research has much to offer to--and gain from--the development of these future dictionaries.

So far, computerization has been applied to lexicography in only limited ways; the COBUILD project [32] was one of the first to utilize computers to exploit corpora in the creation of lexical entries, and most dictionary publishers now create dictionaries by first creating in-house databases--although such databases typically contain only gross distinctions among information fields (orthographic form, pronunciation, part-of-speech, etymology, definition text, etc.). Almost no work has been done to improve definition texts themselves or to systematize semantic information (apart from occasional

It is clear already that MRD research can provide much input that will be valuable for the development of future dictionaries. Computerization of dictionary-making at the semantic level could involve the following:

(1) *the creation of explicit semantic links* (hypernym, part, color, etc.) between words or entries. This would be especially useful for creating electronic (hypertextual) dictionaries. Such links could lead to the development of precise templates for classes of objects, etc. So far, navigation and query in electronic dictionaries is rather limited and relies on user's judgement and understanding of definitions to be usable. However, an explicit semantic net underlying the dictionary could be very useful for nativation and query. For example, we can envision display to varying levels of detail, depending on user preference, of the information in the template for *fruit*, and even user navigation within the template (click on PARTS-OF "apricot" and get "pit," "flesh," "skin," etc.; click on "pit" and get the properties of apricot pit, etc.). Information could be linked to images and sounds, and displayed in template form; or definitions and sub-definitions could be generated in natural language, in any form (concise, learner's version, full, etc.).

(2) *ensuring consistency of the content of entries.* The templates created for different classes of concepts could be used to ensure that the information given for each entry when it is appropriate to do so. This could eliminate the kinds of inconsistencies demonstrated in the entries for *apricot* and *peach* given in section 3.1.1.

(3) *ensuring consistency of metatext.* As outlined above in section 3.2.2, MRD research has revealed that metatext in dictionaries is highly variable and relatively inconsistent, and, by compiling lists of metatextual specifications for various relations, has identified both the sources and potential solutions to the problem of inconsistency. We can even imagine that metatext could be automatically generated from templates of the type described above.

(4) *ensuring consistency of sense division*. Lexicographers have sought means to remove the arbitrariness of sense divisions.; Computerization has already been helpful, for example, by automatizing and systematizing the use of corpora as a source of information about word senses. Knowledge bases could take this systematization even farther. For instance, a dictionary might define *cup* as "**1.** a container for liquid... **2.** its content", *bowl* as "**1.** a container...; its content", but *glass* only as "**1.** a container...", which ignores the metonymic use of glass ("its content") and is therefore inconsistent. An electronic database containing explicit marking of metonymic links could enable checking that metonymic use is specified where necessary, and in a consistent format.

4. Conclusion

The false expectation that large knowledge bases could be generated automatically from MRDs has led to a perception that the past fifteen years of MRD research has failed to meet the original goals. Indeed, sources, and will require human involvement. From this perspective, it is clear that while they are not the exclusive resource they may have been originally thought to be, MRDs have something to contribute to the creation of knowledge bases. In some instances that contribution is not what had originally been expected, as evident in the fact that MRDs have been found to contain a vast bank of associational information that is useful in many NLP tasks.

It is also clear that MRD research has in fact contributed to other NLP goals. In particular, it has contributed to our understanding of the nature, kinds, and role of semantic information in the processing of natural languages, as well as to the increasingly obvious idea that widely varying types of information--several in addition to the traditionally accepted set--are needed for NLP. This may in turn lead to more systematic assessment of the needs of various NLP tasks, an area which deserves serious attention.

The most promising avenue of activity, however, involves collaboration between the NLP community and lexicographers and electronic publishers. The two communities are already beginning to work with one another; one clear example is a recent survey sent to NLP researchers from Longman publishers, asking for their input in devising new database versions of the *LDOCE*. Obviously, collaboration--both in terms of shared information and shared effort--can benefit both communities. The result could be better dictionaries and exciting new possibilities for electronic, hypertextual dictionary databases, as well as a wealth of material useful for NLP.

References

[1] Amsler, R. A. The structure of the Merriam-Webster Pocket Dictionary. Ph. D. Dissertation, University of Texas at Austin, (1980).

[2] Calzolari, N. Detecting patterns in a lexical data base. *Proceedings of the 10th International Conference on Computational Linguistics, COLING'84* (1984), 170-173.

[3] Chodorow, M. S., Byrd. R. J., Heidorn, G. E. Extracting semantic hierarchies from a large on-line dictionary. *Proceedings of the 23rd Annual Conference of the Association for Computational Linguistics,* Chicago (1985), 299-304.

[4] Markowitz, J., Ahlswede, T., Evens, M. Semantically significant patterns in dictionary definitions. *Proceedings of the 24th Annual Conference of the Association for Computational Linguistics*, New York (1986), 112-119.

[5] Byrd, R. J., Calzolari, N., Chodorow, M. S., Klavans, J. L., Neff, M. S., Rizk, O. Tools and methods for computational linguistics. *Computational Linguistics,* 13, 3/4 (1987), 219-240.

[6] Nakamura, J., Nagao, M. Extraction of semantic information from an ordinary English dictionary and its evaluation. *Proceedings of the 13th International Conference on Computational Linguistics, COLING'88* (1988), 459-464.

[7] Ide, N., Véronis, J. Refining taxonomies extracted from machine-readable dictionaries. In Hockey, S., Ide, N. *Research in Humanities Computing 2,* Oxford University Press (1993).

[8] Klavans, J., Chodorow, M., Wacholder, N. From dictionary to knowledge base via taxonomy. *Proceedings of the 6th Conference of the UW Centre for the New OED,* Waterloo, (1990), 110-132.

[9] Wilks, Y., Fass, D., Guo, C., MacDonald, J., Plate, T., Slator, B. Providing Machine Tractable Dictionary Tools. *Machine Translation*, 5 (1990), 99-154.

[10] Pigamo, F. *Outils de traitement sémantique du langage naturel.* Thèse de l'Ecole Nationale Supérieure des Télécommunications, Paris (1990), 242pp.

[11] Alonge, A. Analysing dictionary definitions of motion verbs. *Proceedings of the 15th International Conference on Computational Linguistics, COLING'92* (1992), 1315-1319.

[12] Amsler, R.A. Words and worlds. *Proceedings of the Third Workshop on Theoretical Issues in Natural Language Processing (TINLAP-3).* Las Cruces, NM (1987).

[13] Boguraev, B. K. The definitional power of words *Proceedings of the Third Workshop on Theoretical Issues in*

[14] McRoy, S. W. Using Multiple knowledge Sources for Word Sense Discrimination. *Computational Linguistics* (1992), 18, 1, 1-30.
[15] Martin, R. Inférences et définition lexicographique. *Colloque "Lexique et Inférences"*, Metz (1991), Proceedings to appear.
[16] Procter, P. Cambridge Language Survey: The development of a non-language specific semantic coding system using multiple inheritance. *Paper presented at International Workshop of the European Association of Machine Translation, "Machine Translation and the Lexicon"*, Heidelberg, 26-28 (April 1993).
[17] Jensen, K., Binot, J.-L. Disambiguating prepositional phrase attachments by using on-line dictionary definitions. *Computational Linguistics* (1987), 13, 3-4, 251-260.
[18] Montemagni, S., Vanderwende, L. Structural patterns vs. string patterns for extracting semantic information from dictionaries. *Proceedings of the 15th International Conference on Computational Linguistics, COLING'92* (1992), 546-552.
[19] Ravin, Y. Disambiguating and interpreting verb definitions. *Proceedings of the 28th Annual Conference of the Association for Computational Linguistics*, Pittsburgh (1990), 260-267.
[20] Neff, M.S., Boguraev, B.K. Dictionaries, dictionary grammars and dictionary entry parsing. *Proceedings of the 27rd Annual Conference of the Association for Computational Linguistics*, Vancouver (1989), 91-101.
[21] Ahlswede, T., Evens, M., Rossi, K. Building a lexical database by parsing Webster's Seventh Collegiate Dicionary. Second Annual Conference of the UW Centre for the NewOED. Waterloo, Canada (1985), 65-76.
[22] Ide, N., Véronis, J. Very large neural networks for word-sense disambiguation. *9th European Conference on Arfificial Intelligence, ECAI'90*, Stockholm (1990), 366-368.
[23] Kilgarriff, A. Dictionary word sense distinctions: An enquiry into their nature. *Computers and the Humanities* (1993), 26 (5-6), 365-388.
[24] Lenat, D.B., Prakash, M., Shepherd, M. CYC: Using common sense knowledge to overcome brittleness and knowledge acquisition bottlenecks. *AI magazine* (1986), 7 (4), 65-85.
[25] Amsler, R. A., Tompa, F. W. An SGML-based standard for English monolingual dictionaries. *Proceedings of the 4th Annual Conference of the UW Centre for the New Oxford English Dictionary*. Waterloo, Ontario (1988), 61-80.
[26] Ide, N., Véronis, J. Print Dictionaries, *TEI Working Paper AI5 D17*, Distributed by the Text Encoding Initiative. Compter Center, University of Illinois at Chicago (1992), 60pp.
[27] Ide, N., Veronis, J., Warwick-Armstrong, S., Calzolari, N. Principles for encoding machine readable dictionaries, *EURALEX'92 Proceedings*, H. Tommola, K. Varantola, T. Salmi-Tolonen, Y. Schopp, eds., in *Studia Translatologica*, Ser. a, 2, Tampere, Finland, (1992), 239-246.
[28] Ide, N., Véronis, J.. Encoding dictionaries. *Computers and the Humanities*, 29, 1-3 (1995) to appear.
[29] Neff, M. S., Byrd, R. J., & Rizk, O. A. Creating and querying lexical databases. *Proceedings of the Association for Computational Linguistics Second Applied Conference on Natural Language Processing*. Austin, Texas (1988), 84-92.
[30] Ide, N., Le Maitre, J., Veronis, J. Outline of a Model for Lexical Databases. *Information Processing and Managment* (1993), 29, 2, 159-186.
[31] Le Maitre, J., Ide, N., Véronis, J. Deux modèles pour la représentation des données lexicales et leur implémentation orientée-objet. *Actes des 9èmes Journées Bases de Données Avancées*, Toulouse (1993), 312-331.
[32] Sinclair, J. M.. *An Account of the COBUILD Project.* London: Collins ELT (1987).