Incorporating syntactic dependency information towards improved coding of lengthy medical concepts in clinical reports

Vijayaraghavan Bashyam, PhD*
Monster Worldwide Inc.
Mountain View, CA 94043
vbashyam@ucla.edu

Ricky K Taira, PhD
Medical Imaging Informatics Group
University of California, Los Angeles
Los Angeles, CA 90024
rtaira@mii.ucla.edu

Abstract

Medical concepts in clinical reports can be found with a high degree of variability of expression. Normalizing medical concepts to standardized vocabularies is a common way of accounting for this variability. One of the challenges in medical concept normalization is the difficulty in comparing two concepts which are orthographically different in representation but are identical in meaning. In this work we describe a method to compare medical phrases by utilizing the information found in syntactic dependencies. We collected a large corpus of radiology reports from our university medical center. A shallow semantic parser was used to identify anatomical phrases. We performed a series of transformations to convert the anatomical phrase into a normalized syntactic dependency representation. The new representation provides an easy intuitive way of comparing the phrases for the purpose of concept normalization.

1 Introduction

A vast amount of electronic information is generated in hospitals as a part of routine clinical care due to the adoption of the electronic medical record by health care centers in the United States (Berner *et al.*, 2005; Jha *et al.*, 2006). A significant portion of this information is in the form of unstructured free-text (Hall, 2000; Tange *et al.*, 1998). A free text representation makes it difficult for applications to accurately extract medical information for generic purposes (Ananiadou *et al.*, 2004). The problem of variability of expression in natural language expression has been well studied (Bates, 1986, 1989, 1998; Blair and Maron, 1985; Funk and Reid, 1983; Furnas *et al.*, 1984; Gomez *et al.*, 1990). In the medical domain in particular, users frequently express the same concept in different ways and different concepts in similar ways (Ananiadou and Nenadic, 2006). To illustrate, the terms *heart attack* and *cardiac attack* both refer to the same concept – *myocardial infarction*. Conversely the term *left lobe* could refer to the *left lobe of lung* or the *left lobe of liver* depending on the context (occurrence in a chest radiology report versus a gastro-intestinal radiology report). Such variability suggests a need to normalize concepts encountered in medical reports to a standard vocabulary in order to ensure interoperability.

Several standardized vocabularies exist in the medical domain such as the *Unified Medical Language System* (Humphreys and Lindberg, 1993), *Systematized Nomenclature of Medicine - Clinical Terms* (College of American Pathologists, July 2003), *Medical Subject Headings* (National Library of Medicine), and the *International Classification of Diseases* (World Health Organization). There have been several attempts in the past (Aronson, 2001; Bashyam and Taira, 2005; Bashyam *et al.*, 2007; Cooper and Miller, 1998; Friedman *et al.*, 2004; Nadkarni *et al.*, 2001; Oliver and Altman, 1994; Ruch *et al.*, 2003; Zou *et al.*, 2003) to map medical concepts to their standardized concept found in these terminologies. These approaches are based on mostly on lexical matching (Bashyam *et al.*, 2007), string matching (Nadkarni *et al.*, 2001), statistical indexing (Cooper and

* formerly with the Medical Imaging Informatics Group, Dept. of Radiological Sciences, University of California - Los Angeles, Los Angeles, CA 90024

Proceedings of the Workshop on BioNLP, pages 125–132,
Boulder, Colorado, June 2009.

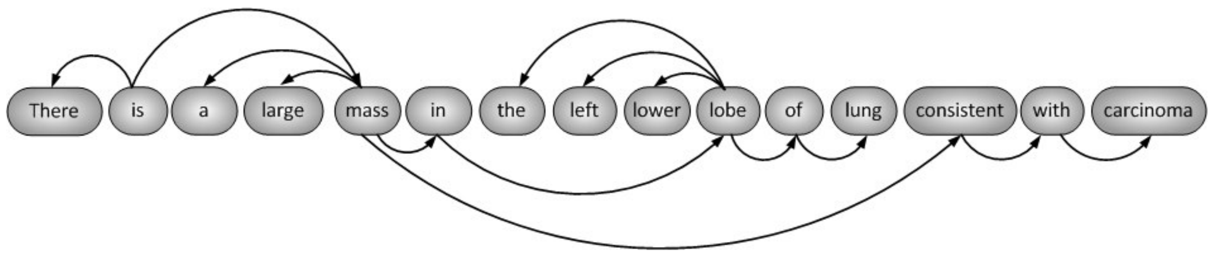

Figure 1. Example of a syntactic dependency parse tree with emphasis towards semantics. Each arc shows a dependency relation between a head and a modifier.

Miller, 1998), natural language processing (Aronson, 2001; Friedman *et al.*, 2004) information retrieval techniques (Bashyam and Taira, 2005; Oliver and Altman, 1994; Ruch *et al.*, 2003; Zou *et al.*, 2003) or a combination of these approaches (Cooper and Miller, 1998). These systems have managed to map a large percentage of medical terms to their respective standard terminologies in their reported experiments. While these systems have managed to perform satisfactorily for the task of normalizing simple expressions, they all acknowledge the larger problem of normalizing lengthy expressions. To illustrate, Nadkarni et al. (2001) mention the mapping of the phrase *spleen rupture and normal stomach* to the concept *stomach rupture* as a possible spurious mapping.

We hypothesize that using deep syntactic information can help in avoiding such spurious mapping. We describe a system which uses information found in syntactic dependencies to help in the coding of lengthy phrases. Preliminary results using this approach are reported as a proof-of-concept.

2 Background

Syntactic dependency parsing has received much focus from the natural language processing community (Eisner, 1996; Kudo and Matsumoto, 2000; Nivre and Scholz, 2004; Yamada and Matsumoto, 2003). A syntactic dependency relation is an asymmetric relation between two words. One word is called the head, and the other word is called the modifier or dependent. A word in the sentence can play the role of the head in several dependency relations (i.e., it can have several modifiers) but each word can play the role of the modifier only once. A special word, named the root, does not play the role of the modifier in any relation. The set of dependency relations that can be defined on a sentence form a tree, called the dependency tree. An example of dependencies in a typical sentence found in a radiology report is shown in Figure 1.

Systems based on syntactic dependencies have been used successfully in several information retrieval experiments with results outperforming traditional retrieval systems (Croft *et al.*, 1991; Gao *et al.*, 2004; Gonzalez *et al.*, 2005; Smeaton, 1986). In particular, this method has been used for word sense disambiguation (Lin, 1997) and thesaurus construction (Lin, 1998). Dependency trees have also been used for medical concept representation in the domains of radiology (Steimann, 1998) and pathology (Romacker *et al.*, 1999).

3 Methods

3.1 Anatomy Phrase Extraction

For identifying anatomy phrases, we use a specialized phrase parser trained to identify anatomy phrases within clinical reports. The input to the parser is a sentence tagged with a part-of-speech tag and a semantic tag. The lexical analyzer module of our NLP system takes a single sentence as the input and produces an output of word tokens tagged with their syntactic and semantic classes. The semantic tag is obtained by mapping tokens in a sentence to a taxonomy handcrafted for the domain of radiology reports custom built from radiology textbooks, radiology review manuals, radiology word compilations and published radiology glossaries apart from actual radiology reports (Taira *et al.*, 2001). Features of our implementation include: 1) a large number (>450) of semantic classes as compared to lexical sources currently available allowing improved discrimination for tasks such as syntactic parsing, semantic interpretation and frame building; 2) the system recognizes special symbols including dates, medical abbreviations, medical coding symbols, numeric measurements, image slice references, and proper names; and 3) the system performs some word sense disambiguation using surrounding syntactic and semantic word features.

Our phrase parsing module currently targets anatomy phrases (e.g., *right upper lobe of lung*), existential relationships (e.g., *there is no evidence of*), and spatial relationships (e.g., *is located 1cm above*). We utilize a supervised learning approach to estimate the feature weights to a maximum entropy model which classifies words as the *start, inside, end, single,* or *outside* of a phrase boundary. A Viterbi dynamic programming algorithm is used to maximize the tag sequence probability. The anatomy phrase chunker has been tested on 4,500 sentences with recall and precision scores of 97.1% and 97.4% respectively.

3.2 Normalized Dependency Representation

We perform a series of transformations to convert an anatomical phrase from a *free-text* representation to a *normalized dependency vector space* representation. The following steps are taken in the representation conversion:

Syntactic Parsing

The anatomy phrase identified by the phrase parser preserves lexical information which is used to obtain a dependency parse tree using a full syntactic parser. This parser is based on a novel *field theory* approach to dependency parsing. The parser is strongly modeled for the radiology domain with performance accuracies of 84.9% and 89.9% for link precision and recall respectively for parsing whole sentences (Taira *et al.*, 2007). In comparison, the state-of-the-art parsers have performance accuracies in the low nineties for link precision and recall in the domain of newspaper text, with performance unknown in the domain of clinical text.

Link Reduction

Our system classifies dependency links into two types – bilexical links and trilexical links. A bilexical link is a strong dependency relation between two words (e.g. determiner←noun) whereas a trilexical link usually has a mediator word in between the two words (e.g. finding→in→location). When possible, a trilexical link is converted to a bilexical link by the elimination of the mediator word and the link type is tagged by the mediator word. The link type can play important roles in certain cases. In cases where the mediator word is also important, the trilexical link is considered as a pair of bilexical links.

Token Level Normalization

Once the parse tree is obtained, the tokens are normalized to their base form. The normalization is an approximate kind of lemmatization. However we also perform word level synonym normalization. For lemmatization, we use the Lexical Variant Generator tools developed by the National Library of Medicine for biomedical text (McCray *et al.*, 1994). For synonyms, we use a handcrafted lexicon built for the domain of radiology. This step helps in avoiding missing a mapping due to lexical differences due to pluralization, abbreviations and acronyms, case differences etc. This representation is referred to as the normalized dependency vector space representation

3.3 Mapping to a Terminology

The normalized dependency parse tree is represented as in a vector space as a *bag-of-links* as analogous to the so-called *bag-of-words* representation in conventional information retrieval. Two phrases can now be compared by using similarity measures such as cosine, dice, jaccard etc. within the dimension-space of dependency-links. One phrase can be the anatomy phrase in a clinical report and the other phrase can be an entry in a standardized terminology.

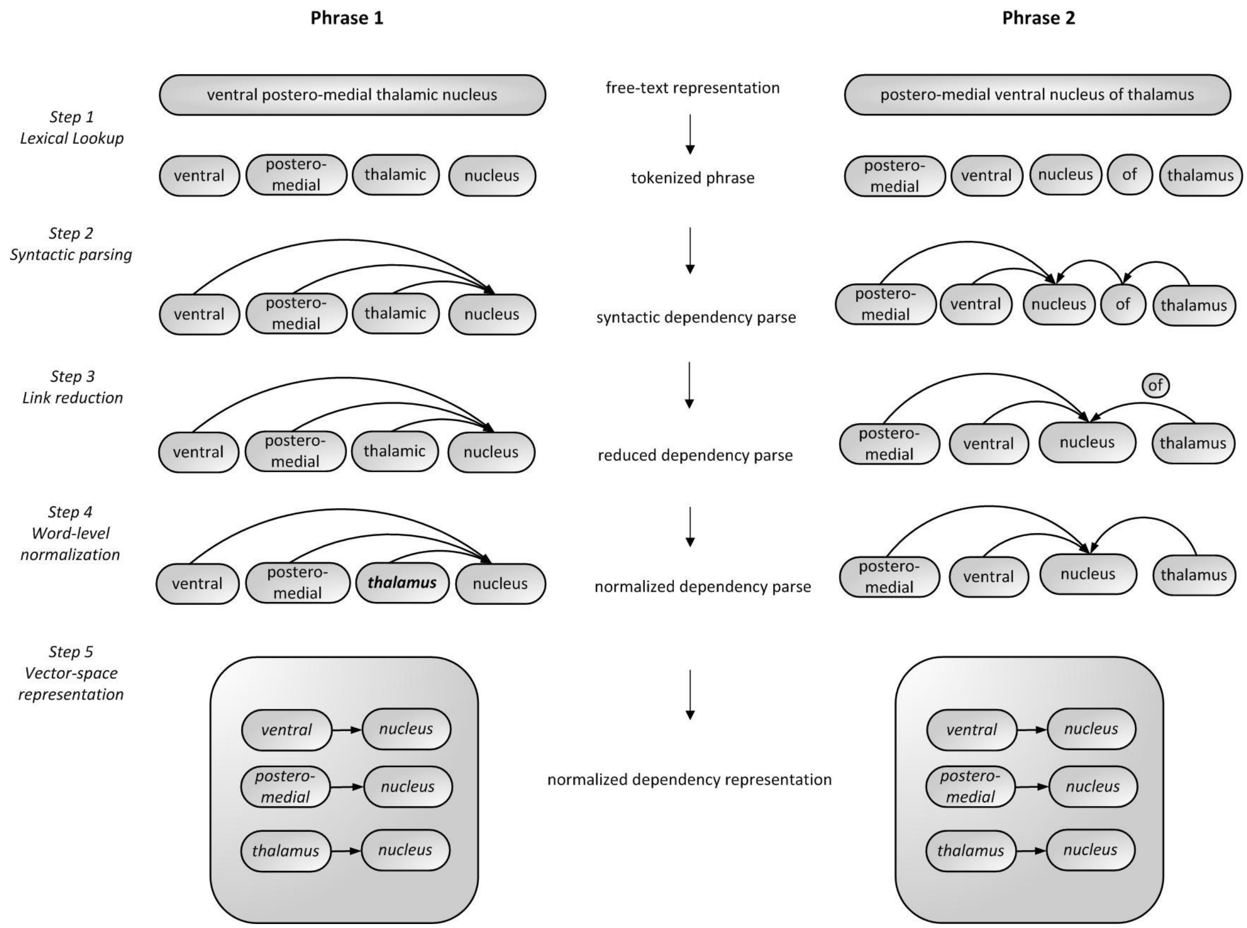

Figure 2. Example illustrating the transformation of a medical phrase from a free-text representation to a normalized syntactic dependency vector space representation.

An exercise in normalization is described in Figure 2 to illustrate how this method works. Consider the following phrase in a neuro-radiology report: *ventral postero-medial thalamic nucleus*. The corresponding concept in the target terminology is the phrase *postero-medial ventral nucleus of thalamus*. These phrases if compared by string matching will not result in direct matches. Permuting words and trying to compare rearrangements is complicated. In our approach, we first preprocess our terminology list and store it in a database. The preprocessing step is described in the right column (Phrase 2) of Figure 2. Starting with the phrase *postero-medial ventral nucleus* of thalamus, we first tokenize the individual words (lexical analysis) in the first step. In the second step, we parse the phrase to arrive at the dependency tree. In the third step, the trilexical link *nucleus←of←thalamus* is converted to a bilexical link by eliminating the word *of* and tagging it as the link type. In the following step, each word is normalized to its base form. In the fifth step, the phrase is represented as a *bag-of-links* and stored in a database. Similarly all the other phrases in our terminology are stored.

When the query phrase *ventral postero-medial thalamic nucleus* is compared against the terminology it undergoes the same processes previously described (Figure 2, Phrase 1). The importance of *word-normalization* can be seen here. In step 4, the word *thalamic* is normalized to *thalamus*. The final output is the *bag-of-links* representation. For convenience of comparison Figure 2 shows together, the query phrase and target phrase undergoing the various steps starting from a *bag-of-words* representation to a *bag-of-links* representation. It is clear that both phrases look identical in the final representation. While a string comparison would have missed equating the two in their original word-level representation, a comparison in the dependency vector space is likely to score them as a perfect match.

4 Experiment and Results

We obtained a set of 2500 neuro-radiology reports from our university medical center. Using the shallow semantic parser, we extracted a set of 2551 unique anatomical phrases. Of the 2551 phrases, 819 phrases were single worded terms. We discarded the single word terms. Single worded phrases do not fall into the difficult-to-map category which this method is specifically aiming to address. Moreover, a minimum of two words are required to define a syntactic dependency and thus the method is irrelevant for single worded terms. Thus we used only the 1732 multi-worded terms in our experiment. The average length of the multi-worded terms was 2.48 words.

We chose the UMLS, a coordinated repository of vocabularies as a target for concept coding. To reduce complexity, we removed non-English concepts and concepts outside the domain of neuro-radiology by filtering out unrelated concepts. Our final terminology had a size of about 100,000 entries. We preprocessed the entire terminology using the above mentioned steps and stored the dependency representation in a database. Every anatomy phrase was queried against this database and cosine similarity was used to measure relevance. No weighting system was employed although it is possible to weight links by their types. A physician domain expert manually evaluated the results of the 1732 queries for performance. Of the 1732 phrases, 1091 phrases (62.9% accuracy, 95% CI ±0.946%) were successfully matched. Since the target set is extremely large in size (as in any IR system), a recall analysis was not performed. A baseline comparison with MMTx (in phrase mode) resulted in 1051 phrases (60.68% accuracy, 95% CI ±0.49%) being mapped by MMTx. Table 1 summarizes the results.

MMTx Matched Phrases	Syn. Dependency Matched Phrases	
1051	1091	n=1732
60.68%	62.99%	
(±0.49%)	(±0.49%)	

Table 1. Overview of Results

5 Discussion

Analysis of the errors showed that the following error types resulted in the inability to match phrases perfectly:

Parsing without context:

A syntactic parser can parse a sentence and identify dependency relations in a sentence. However, when a phrase is given as an input, it is not always easy to parse a phrase and generate a dependency representation. There is context (remaining portions of the sentence) missing which is needed to unambiguously parse the phrase. In the case of anatomical phrases, our system was able to parse it because the source sentences from which they were extracted were available. However, in the case of the UMLS phrases, there is no such available information. Therefore manual parsing of several UMLS phrases had to be performed. One potential solution to this problem could be to identify MEDLINE sentences that contain these UMLS concepts and obtain a dependency parse tree using the context of the sentence.

Modular system architecture:

Since the system is modular, any errors in one of the modules (tokenization, word level normalization etc.) would result in the final dependency representation being imperfect. The specific errors we noticed were:

Parsing Errors:

Our parser has a higher accuracy for parsing phrases than whole sentences. However in this experiment, there were 37 instances where it failed in assigning the correct links. This resulted in partial matches.

Word Normalization Errors:

There is a natural ambiguity introduced when words are normalized to their base forms. Words with completely different senses can have the same root form (e.g. *left←leaves* and *left←left* (spatial direction)). Similarly, a word can have different normalized forms depending on the sense (e.g. *leaf←leaves* and *left←leaves*). A robust method for word-level normalization is desired that can also perform word-sense disambiguation. Currently the NLM's word level normalization tool is being used which is not perfect and therefore errors introduced due to this module result in the entire phrase being transformed incorrectly or ambiguously. The ideal word level normalization will result in the words *cancer, cancerous, carcinoma* all conflating to the same word which is beyond purely morphological analysis.

Link Reduction Errors:

Not all relations manifest as simple bilexical and trilexical links. Some relations are tetralexical and although they can be reduced effectively to bilexical links, the methodology needs to be investigated. To illustrate, consider the phrases '*mass consistent with cancer*' and '*cancerous mass*' parsed as

mass←consistent←with←cancer
cancerous→mass.

The former is parsed as four words with three links. To convert it into a bilexical link, the words '*consistent*' and '*with*' need to be: (1) clustered as a single token and (2) eliminated by transferring it to the link as a label. This is a more complicated process and we still haven't explored such abstractions. A robust rule based link reduction system is desired to handle such cases.

Another limitation of this method is that the heuristic rules for link reduction may not be applicable outside the radiology domain. Finally, syntactic dependency parsers are built using computationally complex algorithms. Thus while using them can result in advanced language understanding, they may not be suitable for real-time applications. There is always a tradeoff between accuracy and speed and it remains to be seen if robust low complexity parsers can be developed.

The inability to perform a recall analysis also make is difficult to judge the theoretical best performance. That is, it is quite likely that there are many phrases in our dataset that do not have a corresponding UMLS concept. Performing a recall analysis would help in determining this.

While we noticed several areas of improvement in our system, we were encouraged by the comparison of the overall results of our system to that of MMTx. We did not do an error analysis of MMTx since several previous publications have documented the various kinds of errors in MMTx (Bashyam *et al.*, 2007; Divita *et al.*, 2004; Meng *et al.*, 2005). Our idea is to provide a baseline comparison showing that our approach performs comparably if not better than MMTx which is the most commonly used[1] tool for concept coding. To our knowledge this the first time syntactic dependencies have been used for this task, Previous attempts have relied purely on shallow parsers.

6 Future Work

Increasing the robustness of the individual modules is a primary requirement for further experiments to prevent the *weakest link effect* cascading to the final output. Specifically we plan to work towards a robust word level normalization system. Additionally, robust evaluation methods including comparisons with other techniques will be investigated.

7 Conclusion

Syntactic dependency based methods for medical concept coding show promise. While some of the described implementations are specific to domain (radiology) and phrase type (anatomy), it is expected that the principle is general enough to be applied in other domains as well.

[1] For an overview of recent applications of MMTx, see (Bashyam *et al.*, 2007)

Acknowledgements

The authors would like to thank Lew Andrada, Gregory Leazer, Jonathan Furner and Christine Borgman for several useful suggestions.

This work was supported in part by the following grants:

1. National Institute of Biomedical Imaging and Bioengineering P01-EB00216
2. National Institute of Health R01-EB002247

References

Ananiadou, S., Friedman, C., and Tsujii, J. (2004). Introduction: named entity recognition in biomedicine. Journal of Biomedical Informatics, 37(6), 393-395.

Ananiadou, Sophia, and Nenadic, Goran. (2006). Automatic Terminology Management in Biomedicine. In Sophia Ananiadou and John McNaught (Eds.), Text Mining for Biology And Biomedicine. Norwood MA: Artech House.

Aronson, A. R. (2001). Effective mapping of biomedical text to the UMLS Metathesaurus: the MetaMap program, AMIA Annu Symp Proc (pp. 17-21). Portland, OR.

Bashyam, V., and Taira, R. K. (2005). Indexing Anatomical Phrases in Neuro-Radiology Reports to the UMLS 2005AA, American Medical Informatics Association Annual Symposium (pp. 26-30). Washington DC.

Bashyam, Vijayaraghavan, Divita, Guy, Bennett, David B, Browne, Allen C, and Taira, Ricky K. (2007, August 2007). A Normalized Lexical Lookup Approach to identifying UMLS concepts in free text. Paper presented at the 12th World Congress on Health (Medical) Informatics (MEDINFO). Brisbane, Australia.

Bates, M. J. (1986). Subject access in online catalogs: A design model. Journal of the American Society for Information Science, 37(6), 357-376.

Bates, M. J. (1989). Rethinking Subject Cataloging in the Online Environment. Library Resources and Technical Services, 33(4), p400-412.

Bates, M. J. (1998). Indexing and access for digital libraries and the internet: Human, database, and domain factors. Journal of the American Society for Information Science, 49(13), 1185-1205.

Berner, E. S., Detmer, D. E., and Simborg, D. (2005). Will the Wave Finally Break? A Brief View of the Adoption of Electronic Medical Records in the United States. J Am Med Inform Assoc, 12(1).

Blair, David C., and Maron, M. E. (1985). An Evaluation of retrieval effectiveness for a full-text document retrieval system. Communications of the ACM, 28(3), 289-299.

College of American Pathologists. (July 2003). SNOMED Clinical Terms ® Technical Reference Guide.

Cooper, G. F., and Miller, R. A. (1998). An Experiment Comparing Lexical and Statistical Methods for Extracting MeSH Terms from Clinical Free Text, Proceedings of the American Medical Informatics Association Annual Symposium.

Croft, W. B., Turtle, H. R., and Lewis, D. D. (1991). The use of phrases and structured queries in information retrieval. Paper presented at the Proc of the 14th Annu Intnl ACM SIGIR Conf on R&D in Information Retrieval.

Divita, Guy, Tse, Tony, and Roth, Laura. (2004). Failure analysis of MetaMap transfer (MMTx). In M Fieschi et. al. (Ed.), Proceedings of the 11 World Congress on Health Informatics (MEDINFO 2004).

Eisner, J. (1996). Three new probabilistic models for dependency parsing: An exploration. Paper presented at the Proc COLING.

Friedman, C., Shagina, L., Lussier, Y., and Hripcsak, G. (2004). Automated Encoding of Clinical Documents Based on Natural Language Processing. Journal of the American Medical Informatics Association, 11(5), 392.

Funk, M. E., and Reid, C. A. (1983). Indexing consistency in MEDLINE. Bulletin of the Medical Library Association, 71(2), 176-183.

Furnas, G. W., Landauer, T. K., Gomez, L. M., and Dumais, S. T. (1984). Statistical semantics: analysis of the potential performance of keyword information systems. In Human factors in computer systems (pp. 187-242): Ablex Publishing Corp. Norwood, NJ, USA.

Gao, J., Nie, J. Y., Wu, G., and Cao, G. (2004). Dependence language model for information retrieval. Paper presented at the Proc ACM SIGIR.

Gomez, L. M., Lochbaum, C. C., and Landauer, T. K. (1990). All the right words: Finding what you want as a function of richness of indexing vocabulary. Journal of the American Society for Information Science, 41(8), 547-559.

Gonzalez, M., Strube De Lima, V. L., and Valdeni De Lima, J. (2005). Binary Lexical Relations for Text Representation in Information Retrieval. In Natural Language Processing and Information Systems (Vol. 3513, pp. 21-31). Berlin / Heidelberg: Springer

Hall, F. (2000). Language of the radiology report. American Journal of Roentology, 175, 1239-1241.

Humphreys, B. L., and Lindberg, D. A. B. (1993). The UMLS project: making the conceptual connection between users and the information they need. Bulletin of the Medical Library Association, 81(2), 170.

Jha, Ashish K., Ferris, Timothy G., Donelan, Karen, DesRoches, Catherine, Shields, Alexandra, Rosenbaum, Sara, and Blumenthal, David. (2006). How Common Are Electronic Health Records In The United States? A Summary Of The Evidence. Health Affairs, doi: 10.1377/hlthaff.25.w496.

Kudo, T., and Matsumoto, Y. (2000). Japanese dependency structure analysis based on support vector machines. Paper presented at the Proc International Conference On Computational Linguistics.

Lin, D. (1997). Using syntactic dependency as local context to resolve word sense ambiguity, Proceedings of the 35th Annual Meeting of the Association for ComputationalLinguistics and 8th Conference of the European Chapter of the Association for Computational Linguistics. Madrid, Spain.

Lin, D. (1998). Automatic retrieval and clustering of similar words. Paper presented at the Proc 36th Annu Meeting of the ACL and 17th Intnl Conf on Comp Ling.

McCray, AT, Srinivasan, S, and Browne, AC. (1994). Lexical methods for managing variation in biomedical terminologies. Paper presented at the Proceedings of the Annual Symposium on Computer Applications in Medical Care.

Meng, F., D'Avolio, L. W., Chen, A. A., Taira, R. K., and Kangarloo, H. (2005). Generating Models of Surgical Procedures using UMLS Concepts and Multiple Sequence Alignment, AMIA Annu Symp Proc (Vol. 520, pp. 4). Washington, DC.

Nadkarni, P., Chen, R., and Brandt, C. (2001). UMLS Concept Indexing for Production Databases: A Feasibility Study. Journal of the American Medical Informatics Association, 8(1), 80.

National Library of Medicine. MEdical Subject Headings. www.nlm.nih.gov/mesh/ meshhome.html Last Accessed 09 March 2007 Bethesda, MD.

Nivre, J., and Scholz, M. (2004). Deterministic dependency parsing of English text. Paper presented at the Proc COLING.

Oliver, D. E., and Altman, R. B. (1994). Extraction of SNOMED concepts from medical record texts, Proc of the 18th Annual SCAMC (Vol. 17983). Washington: McGraw Hill.

Romacker, M., Schulz, S., and Hahn, U. (1999). Streamlining semantic interpretation for medical narratives. Paper presented at the AMIA Annu Symp Proc.

Ruch, P., Baud, R., and Geissbühler, A. (2003). Learning-free Text Categorization, Proc of 9th Conf on Artificial Intelligence in Medicine (pp. 199–204).

Smeaton, A. F. (1986). Incorporating syntactic information into a document retrieval strategy: an investigation. Paper presented at the Proc of the 1986 ACM Conf on R&D in Information Retrieval, Pisa, Italy.

Steimann, F. (1998). Dependency parsing for medical language and concept representation. Artificial Intelligence in Medicine, 12(1), 77-86.

Taira, R. K., Soderland, S. G., and Jakobovits, R. M. (2001). Automatic Structuring of Radiology Free-Text Reports. Radiology, 21, 237-245.

Taira, Ricky K, Bashyam, Vijayaraghavan, and Kangarloo, Hooshang. (2007). A Field Theoretical Approach for Medical Natural Language Processing. IEEE Transactions on IT in Biomedicine, 11(4), 364-375.

Tange, H.J., Schouten, H.C., Kester, A.D.M., and Hasman, A. (1998). The Granularity of Medical Narratives and Its Effect on the Speed and Completeness of Information Retrieval. J Am Med Inform Assoc, 5(6), 571-582.

World Health Organization. http://www.who.int/classifications/icd/en/ Last Accessed 09 March 2007.

Yamada, H., and Matsumoto, Y. (2003). Statistical dependency analysis with support vector machines. Paper presented at the Proceedings of the 8th International Workshop on Parsing Technologies (IWPT).

Zou, Q., Chu, W. W., Morioka, C., Leazer, G. H., and Kangarloo, H. (2003). IndexFinder: A Method of Extracting Key Concepts from Clinical Texts for Indexing, AMIA Annu Symp Proc (Vol. 2003, pp. 763-767).