

Evaluation of Metamap Performance in Radiographic Images Retrieval

Lilac Al-Safadi, Rawan Alomran and Fareeda Almutairi

Department of Information Technology, College of Computer and Information Sciences,
King Saud University

Abstract: A large amount of free text is available as a source of knowledge in the biomedical field. MetaMap is a widely used tool that identifies concepts within the UMLS in English text. In this study, we study the performance of MetaMap. Performance is measured in retrieval speed, precision of results and recall of results. This automated MetaMap indexing is compared with manual indexing of the same text. Results shows that MetaMap by default was able to identify 98.19% of the biomedical concepts occurred in the sample set. MetaMap by default identified 78.79% of the concepts that manually were not identified. MetaMap is tested under four scenarios; the default output, displaying the mapping list, restricting the candidates' score within the candidate list and restricting the candidates' score within the mapping list. This study describes the limitations of the MetaMap tool and ways to improve the performance of the tool and increase its recall and precision.

Keywords: Biomedical, MetaMap, performance, UMLS

INTRODUCTION

Much of the biomedical knowledge is represented in free text form with variation of medical terminologies. Such unstructured representations of information are difficult for computers to process consistently and precisely. Precise identification of biomedical and clinical terms in the text is a crucial step in clinical decisions. The task of automatically determining the concepts referred to in text is a common one (Aronson, 1996a). Gooch and Roudsari (2011) defined term identification as the process of: recognizing a text string as a possible term (candidate term selection), classifying the candidate term (e.g., body part, disease, physiological function) and mapping the term to a single concept (pre-coordination) or to multiple concepts (post-coordination) within a standardized vocabulary or ontology.

Numbers of tools were developed to map free text to a biomedical knowledge source. Examples of such tools include MicroMeSH (Elkin *et al.*, 1998), Metaphrase (Tuttle *et al.*, 1998), CLARIT (Evans *et al.*, 1991), CHARTLINE (Miller *et al.*, 1992), SAPHIRE (Hersh and Leone, 1995; Hersh *et al.*, 1994) and a system developed by Nadkarni *et al.* (2001). The MetaMap tool (Aronson *et al.*, 1994; Rindfleisch and Aronson, 1993, 1994) is distinguished by its reliance on Unified Medical Language System (UMLS¹) knowledge.

MetaMap is a tool that extracts biomedical terms in a free text and maps them to the UMLS source

vocabulary concepts (Aronson and Lang, 2010). It has been used for a variety of tasks including information retrieval (Aronson and Rindfleisch, 1997; Pratt and Wasserman, 2000; Wright, 1998) text mining (Pratt and Wasserman, 2000) and concept indexing (Rindfleisch and Aronson, 2002). Yet, few works have evaluated the ability of MetaMap in identifying biomedical concepts.

Because the concept identification step is critical in converting free-form text into a form understandable by the computer, we need to evaluate how these tools are able to match that concept identification process. This study examines the performance of MetaMap performance under different scenarios; the default output, displaying the mapping list only, after restricting candidate list's scores and after restricting the mapping list's candidate's scores. The purpose of this study is to evaluate the ability of MetaMap to effectively recognize important concepts in radiographic imaging reports and express these concepts as UMLS descriptors. The performance results can help developers in understanding the limitations of MetaMap and accordingly deciding on ways of improving its results. The performance of MetaMap is measured in:

- Response time
- Precision of results
- Recall of results

The experiment is done on a sample set of thirty medical cases with the cooperation of our domain

experts. We calculated the response time of MetaMap to measure its performance and we asked our domain experts to identify the biomedical concepts on the cases and compared it with the MetaMap results to evaluate its ability to identify biomedical concepts.

MetaMap displays a candidate list of all concepts that might relate to the entered text. Additionally, it displays a mapping list that shows the mapping process of the candidates. Candidate score measures how much the candidate relates to the entered text. Experiment results shows that restricting the candidates' score within the mapping list shows 4.5% increase in the MetaMap response time, 9.98% increase in the precision and 31.71% decrease in the recall in compare with the default output. In addition, restricting the candidates' score showed a 20.53% increase in the number of biomedical concepts identified by MetaMap over the manual identification process performed by domain experts.

UMLS AND METAMAP

Ontology knowledge enables building a semantic representation of a document. Biomedical domain ontology can be used as a common framework for knowledge representation and exchange because it can connect patient information to concepts stored in the knowledge base (Rindfleisch and Aronson, 2002). A number of projects focused on translating medical terminologies into medical domain ontologies such as UMLS, MeSH² and Radlex³.

Unified Medical Language System (UMLS) Metathesaurus available from the National Library of Medicine (NLM)⁴ is the largest biomedical domain ontology. The UMLS integrates and distributes key terminology, classification and coding standards and associated resources to promote creation of more effective and interoperable biomedical information systems and services, including electronic health records.

MetaMap (2012) is an application developed by the Lister Hill National Center for Biomedical Communications⁵ at the NLM to map biomedical text to the UMLS Metathesaurus or, equivalently, to identify Metathesaurus concepts referred to in English text. MetaMap employs a knowledge-intensive approach, natural-language processing (NLP) and computational-linguistic techniques and is used worldwide in industry and academia according to the National Library of medicine.

MetaMap maps noun phrases in free text to concepts in UMLS Metathesaurus (Rindfleisch and Aronson, 2002). It breaks text into phrases and generates variants for the words in the phrase and forms a candidate set of all Metathesaurus strings containing one of the variants computing and evaluating their match to the input text. Then candidates are combined

and evaluated to produce best matching results (Aronson, 1996b). The candidate list is a set of number of candidates that might relate to the entered text. Each candidate has a concept name, preferred name, confidence score and semantic type. Other candidate's elements such as match map and sources are not covered in our experiment.

Concept name is the name of the concept that the MetaMap extracts from the phrase. It can be an abbreviation or another synonym for the concept. For instance, if the phrase "lung cancer" was extracted from a text, the MetaMap will display "lung cancer" as a concept name for a candidate.

Preferred name is considered the name that best represent a concept. For each group of synonym concepts, one term should be designated as the preferred name for each concept (Rogers *et al.*, 2012). For instance, if phrase "lung cancer" was extracted from the text, MetaMap will assign "malignant neoplasm of lung" as preferred name for a candidate of lung cancer.

Confidence Score is a measure of how much the candidate relate to the text. It is calculated using an evaluation function that measures the quality of the match between a phrase Metathesaurus candidate as illustrated by Aronson (1996a). The confidence score ranges from 0 to 1000. The MetaMap orders candidates from higher to lower score. The higher the score, the higher is the probability that concepts relate to the phrase. For example, "lung cancer" assigned 1000 mapping score to the single concept "Lung Cancer" while assigned an 861 mapping score for each of the two concepts "Lung" and "Cancer".

Semantic type is displayed in MetaMap as form of code that illustrates the category of which a candidate belongs to. For instance, if phrase "lung cancer" was extracted from the text, MetaMap will assign "neop" as the semantic type, which indicates that lung cancer is classified under Neoplastic Process category.

The mapping process of the candidates is clarified on the mapping list that the MetaMap displays. By default, the mapping list displays a set of candidates with high mapping scores. The mapping score indicates the strength of mapping. It is calculated by re-computing the match strength based on the combined candidates and selecting those having the highest score (Aronson, 1996a).

This study applies MetaMap for automatic indexing of radiographic image reports. It examines MetaMap performance under different scenarios; the default output, displaying the mapping list only, after restricting candidate list's scores and after restricting the mapping list's candidate's scores. This is achieved by performing a pilot test as described in the following section.

EXPERIMENTS AND TESTS

In testing the performance of MetaMap we have developed a samples set of thirty random radiographic

Table 1: MetaMap performance test results

Case #	Size	Indexing using MetaMap default output	Indexing using MetaMap with displaying Mapping list	Indexing using MetaMap after restricting the candidates' score within the candidate list	Indexing of query using MetaMap after restricting the candidates' score within the mapping list
1	17	2.3s	0.8s	0.8s	1.1s
2	116	20.8s	5.7s	4.8s	5.4s
3	82	4.2s	3.3s	2.9s	3.4s
4	35	2.5s	2.1s	2.0s	2.1s
5	25	4.0s	2.3s	2.1s	2.5s
6	7	1.2s	0.4s	0.2s	0.8s
7	12	1.8s	0.8s	0.7s	1.0s
8	33	2.9s	1.3s	1.0s	1.3s
9	158	9.8s	6.5s	5.1s	5.7s
10	11	1.5s	0.7s	0.6s	0.7s

cases. The sample set represents ranges of cases; 3 MRIs, 7 X-rays, 9 CT Scans and 11 Ultrasounds of different body parts: Brain, Spine, Chest, Stomach, Liver, Limb, Kidneys and Pelvis. Within each subset we chose a range of abnormalities (normal to abnormal cases with many findings). The cases have a textual description associated with them. The textual descriptions are generated by domain experts following the current radiographic content reporting used in the healthcare institutes they work in.

To achieve the concept indexing, we have used MetaMap. This was an automated process, requiring no human intervention. The indexing process extracts concepts represented in UMLS Metathesaurus from the textual description. The input to MetaMap is the textual description of the radiographic image entered by the domain expert. The output from MetaMap is a single exact Metathesaurus concept or a list of partially matching concept (similar concepts).

In our experiment, we first analyzed the MetaMap response time. After that, we tested the ability of MetaMap to identify concepts in biomedical cases and compared the MetaMap results to the results of professionals in radiology field. The following sections show the results of ten radiographic imaging cases.

Performance results: To obtain the best performance by MetaMap, we calculated the MetaMap response time that it takes on mapping each medical case to the UMLS Metathesaurus. First, we calculated the Metamap response time of its default output by displaying all the concepts that relate to the text. Then, we compared it with the response time of displaying the mapping list, after restricting the candidates' score within the candidate list and restricting the candidates' score within the mapping list.

Table 1 shows the size of the sentences in terms of number of words and the time took to process the indexing requests in seconds under the four different scenarios; using MetaMap.

The results of the performance test show that the indexing of query using MetaMap's default output responded on average of 5.1 sec. Displaying mapping list will result in faster performance at 2.39 sec. As when restricting the candidates' score within the candidate list; displaying only candidates that score above 751, will fasten the performance to 2.02 sec.

Similarly, restricting the candidates' score within the mapping list; displaying only candidates that score above 751, will result in a performance of 2.4 sec. With the analysis of default output, candidate score 751 was determined as the threshold with acceptable number of mistakes.

Approximately, the performance of the last three scenarios is similar, while the delay in the default scenario is due to the time taken by MetaMap to display the entire candidate list and the mapping list.

We also found that the second run of MetaMap on the same input text responses faster than its first run on that text on an average of 1.5% of a minute using MetaMap's default output. This is because mac OS, Linux and possibly some other operating system cache file Input/Output (I/O) to previous invocations of programs (including MetaMap) According to the National Library of Medicine (2011).

Concept identification results: To examine the ability of the system to recognize the important concepts in the textual description entered by the domain expert, we chose to compare MetaMap's results to the results of manual annotation performed by domain experts. We use the MetaMap mapping list and candidate's score to increase the precision of results. We tested the MetaMap under the four scenarios individually and analyzed the concepts identification process. The first scenario is using MetaMap's default output; displaying a candidate list and a mapping list. The second scenario is displaying the mapping list only. The third is restricting the candidates' score within the candidate list; displaying only those with scores above 751. Finally, we tested MetaMap after restricting the candidates' score within the mapping list; displaying only those with score above 751.

The total size of this test was 496 words. The average number of words per image report is 50. For each imaging report, an average of 74.14% important concepts were identified by default. On average, 74.04% of these concepts occurred in UMLS.

Two measurements are used to evaluate the ability of MetaMap indexing; Precision and Recall. Precision is the percentage of retrieved medical cases that are in fact relevant to the query (i.e., "correct" responses). It is calculated using the following equation:

Table 2: Using MetaMap's default output

Case	No. of concepts retrieved by MetaMap	No. of relevant concepts retrieved by MetaMap	No. of concepts retrieved and not relevant by MetaMap	No. of concepts relevant and not retrieved by MetaMap	No. of concepts not relevant and not retrieved by MetaMap	Precision (%)	Recall (%)	Accuracy (%)
#1	40	31	9	0	4	77.5	100	79.55
#2	201	141	60	2	33	70.15	98.6	73.5
#3	169	118	51	4	18	69.82	96.72	71.2
#4	94	50	44	0	9	53.19	100	57.28
#5	42	35	7	2	7	83.33	94.59	82.35
#6	9	6	3	0	4	66.66	100.0	76.92
#7	26	19	7	1	3	73.08	95.00	73.33
#8	41	35	6	0	17	82.86	100.0	89.66
#9	260	191	69	6	78	73.46	96.95	78.2
#10	23	21	2	0	3	91.30	100.0	92.31

Table 3: Displaying the mapping list

Case	No. of concepts retrieved by MetaMap	No. of relevant concepts retrieved by MetaMap	No. of concepts retrieved and not relevant by MetaMap	No. of concepts relevant and not retrieved by MetaMap	No. of concepts not relevant and not retrieved by MetaMap	Precision (%)	Recall (%)	Accuracy (%)
#1	16	16	0	2	8	100.0	88.88	92.31
#2	115	109	6	11	57	94.78	90.83	90.71
#3	80	61	19	6	34	76.25	91.04	79.16
#4	34	25	9	0	13	73.53	100.0	80.85
#5	23	22	1	6	12	95.65	78.57	82.93
#6	4	4	0	0	5	100.0	100.0	100.0
#7	9	9	0	1	5	100.0	100.0	93.33
#8	24	24	0	0	19	100.0	100.0	100.0
#9	104	93	11	24	97	89.42	79.49	84.44
#10	12	11	1	0	4	91.67	100.0	93.75

Table 4: After setting the candidates' score above 751 within the candidate list

Case	No. of concepts retrieved by MetaMap	No. of relevant concepts retrieved by MetaMap	No. of Concepts retrieved and not relevant by MetaMap	No. of concepts relevant and not retrieved by MetaMap	No. of concepts not relevant and not retrieved by MetaMap	Precision (%)	Recall (%)	Accuracy (%)
#1	18	14	4	3	8	77.77	82.35	75.86
#2	83	53	30	21	66	63.85	71.62	70.00
#3	76	49	27	24	52	64.47	67.12	66.45
#4	45	22	23	9	11	48.88	70.98	50.77
#5	33	30	3	2	10	90.90	93.75	88.89
#6	6	3	3	1	5	50.00	75.00	66.67
#7	12	8	4	2	7	66.66	80.00	71.43
#8	36	32	4	0	18	88.88	100.0	92.59
#9	162	131	31	13	84	80.86	90.97	80.27
#10	21	19	2	0	4	90.47	100.0	92.00

Precision = No. of medical cases Retrieved and Relevant (True positive)/No. of retrieved medical cases (True positive+False positive)

Recall is the number of concepts identified correctly and it represents the percentage of medical cases that are relevant to the query. Recall is calculated using the following equation:

Recall = No. of medical cases retrieved and relevant (True positive)/No. of relevant medical cases (True positive+False Negative)

Usually, precision and recall scores are not discussed in isolation. And often there is an inverse relationship between precision and recall, where it is possible to increase one at the cost of reducing the other. Therefore, we also use accuracy as a weighted arithmetic mean of precision and recall.

The same sample set was manually indexed by the domain experts. The indexers examined the

unformatted free textual description and identified the important concepts in the descriptions that optimally represent the knowledge associated with the medical case.

By default, MetaMap displays a candidate list of all candidates that might relate to the entered text. After analyzing the candidate list, regardless of the mapping list and candidates' score, the indexing experience demonstrated that the system can achieve an average precision rate of 74.14%, an average 98.19% recall and an average accuracy of 77.43% as it is shown in Table 2. As we analyzed the output to see if there is a possibility to increase the preciseness and decrease the number of mistakes, we found that most of the mistakes appeared in the candidate list did not appear in the mapping list. In addition, the candidate score of lots of mistakes was 751 and less.

Table 5: After setting the candidates' score above 751 within the mapping list

Case	No. of concepts retrieved by MetaMap	No. of relevant concepts retrieved by MetaMap	No. of concepts retrieved and not relevant by MetaMap	No. of concepts relevant and not retrieved by MetaMap	No. of concepts not relevant and not retrieved by MetaMap	Precision (%)	Recall (%)	Accuracy (%)
#1	6	6	0	5	8	100.0	54.54	73.68
#2	53	43	10	25	67	94.78	63.24	75.86
#3	35	28	7	32	58	76.25	45.9	68.8
#4	18	11	7	16	22	73.53	40.74	58.93
#5	16	16	0	6	15	95.65	72.72	83.78
#6	5	2	3	1	5	100.0	66.66	66.67
#7	5	5	0	2	9	100.0	71.43	87.50
#8	22	22	0	0	20	100.0	100	100.0
#9	73	65	8	31	101	89.42	67.71	80.98
#10	10	9	1	2	3	91.67	81.81	80.00

MetaMap displays mappings with high score only. Candidates listed in high score mapping are more precise. What we found out is displaying the mapping list with high score mapping will increase the average precision rate, in compare with the default output, up to 92.13%. On the other hand, the average recall will decrease down to 92.83%. Whereas the average accuracy will increase to 89.75% as it is shown in Table 2.

When we used the MetaMap and restrict the candidates' score within the candidate list; displaying only those with scores above 751, we found that the average precision rate is 72.27%, the average recall is 83.18% and the average accuracy is 75.49% as it is shown in Table 3.

Finally, we tried to restrict the candidates' score within the mapping list; displaying only those with score above 751. The indexing experience demonstrated that the system could increase the preciseness, in compare with the default output, by recording an average precision rate of 84.12%. Consequentially, it will decrease the average recall to 66.48%. While the average accuracy rate will slightly increase at 77.62% as it is shown in Table 4 and 5.

DISCUSSION AND CONCLUSION

Large amount of free text that is available as a source of knowledge in biomedical field requires a tool to identify contained concept such as MetaMap. MetaMap is a widely used tool that identifies concepts in the UMLS. UMLS is the largest biomedical domain Ontology.

In this study, we have measured the MetaMap performance in retrieval speed, precisian of results and recall of results. The test was conducted under four scenarios. The total size of this test was 496 words with an average number of 50 words per image report. This study had shown that MetaMap is an effective tool for identifying UMLS Metathesaurus concepts in English text. Our finding denotes that:

MetaMap's default output recorded the highest recall rate at 98.19%. Whereas displaying the mapping list recorded the highest precision rate at 92.13%. Restricting the candidates' score within the candidate

list recorded the fastest response time at 2.02 sec. Most of MetaMap failures are caused of one of the below:

- Most MetaMap failures are related to missing entries in UMLS. 2.68% of concepts in sample set were not expressed in UMLS such as "parapharyngeal" and "intraparotid".
- MetaMap cannot detect relationships between concepts, such as "a cyst formed behind the knee". This would results in cases containing "cyst" and "knee" regardless of their relationship.
- MetaMap cannot recognize multi-words concept entries if they are not expressed as multi-words in the UMLS such as "basal cisterns". This would results in cases containing "basal" and "cisterns" separated.
- MetaMap cannot recognize words with some punctuation such as "M.R.I" and "cardio-thoracic".
- Some concepts are expressed differently in UMLS than the actual free text such as the free text "temporo fronto" is expressed in UMLS as "Temporal frontal".
- MetaMap cannot detect spelling mistakes in text.

MetaMap's recall (i.e., number of concepts identified correctly) is presently sub-optimal. The main reason for this is that many terms in the sample set reports are expressed differently than they are in the Metathesaurus and MetaMap's concept-matching algorithm cannot select them definitively over other partially matched terms. MetaMap's recall performance is determined largely by the coverage of biomedical terms in the UMLS and can only be increased substantially by a corresponding increase in the UMLS vocabulary.

MetaMap's precision (i.e., the percentage of retrieved radiographic images that are relevant to the query) has performed the role effectively. This was shown by the high percentage of correct concept that the MetaMap retrieves in the four scenarios mentioned above. The average precision for the four scenarios is 80.66%, which is a good measure for the respectable performance.

The above issues need to be addressed in the design of a series of filters to allow MetaMap to

increase the precision of returned results. Further research will look at improving MetaMap's partial matching capabilities. The work will continue on fine-tuning these issues to produce an optimal recall-precision rate.

MetaMap is highly configurable across multiple dimensions including data options that choose the vocabularies data model to use, output options that determine the nature format of the output generated by MetaMap and processing options that control the arithmetic computation to be performed by MetaMap. A future work would be to test how these processing options may affect the mapping list.

ACKNOWLEDGMENT

This research project was supported by a grant from the "Research Center of the Center for Female Scientific and Medical Colleges", Deanship of Scientific Research, King Saud University.

REFERENCES

- Aronson, A., T. Rindflesch and A. Browne, 1994. Exploiting a large thesaurus for information retrieval. Proceeding of the RIAO, 94: 197-216.
- Aronson, A., 1996a. MetaMap: Mapping Text to the UMLS Metathesaurus. NLM, NIH, DHHS, Bethesda, pp: 1-26.
- Aronson, A., 1996b. The effect of textual variation on concept based information retrieval. Proceeding of the AMIA Symp, pp: 373-7.
- Aronson, A. and T. Rindflesch, 1997. Query expansion using the UMLS Metathesaurus. Proc. AMIA Symp., 36(1): 485-489.
- Aronson, A. and M. Lang, 2010. An overview of MetaMap: Historical perspective recent advances. J. Am. Med. Inform., 17(3): 229-236.
- Elkin, P., J. Cimino, H. Lowe, D. Aronow, T. Payne, P. Pincetl and G. Barnett, 1988. Mapping to MeSH: The art of trapping MeSH equivalence from within narrative text. Proceeding of the 12th SCAMC, pp: 185-190.
- Evans, D., K. Ginther-Webster, M. Hart, R. Lefferts and I. Monarch, 1991. Automatic indexing using selective NLP and first-order thesauri. Proceeding of the RIAO, 91: 624-644.
- Gooch, P. and A. Roudsari, 2011. A tool for enhancing MetaMap performance when annotating clinical guideline documents with UMLS concepts. Proceeding of the IDAMAP Workshop at 13th Conference on Artificial Intelligence in Medicine (AIME).
- Hersh, W. and T. Leone, 1995. The SAPHIRE server: A new algorithm and implementation. In: Gardner, R.M., (Ed.), Proceeding of the 19th SCAMC, pp: 858-862.
- Hersh, W., D. Hickam, R. Haynes and K. McKibbin, 1994. A performance and failure analysis of SAPHIRE with a MEDLINE test collection. J. Am. Med. Inform. Assoc., 1(1): 51-60.
- MetaMap, 2012. U.S. National Library of Medicine. Retrieved from: http://www.nlm.nih.gov/research/umls/implementation_resources/metamap.html, (Accessed on: June 13, 2012).
- Miller, R., F. Gieszczykiewicz, J. Vries and G. Cooper, 1992. CHARTLINE: Providing bibliographic references relevant to patient charts using the UMLS Metathesaurus knowledge sources. Proceeding of the 16th SCAMC, pp: 86-90.
- Nadkarni, P., R. Chen and C. Brandt, 2001. UMLS concept indexing for production databases: A feasibility study. J. Am. Med. Inform. Assoc., 8: 80-91.
- National Library of Medicine, 2011. MetaMap 2011 Usage. Retrieved from: <http://metamap.nlm.nih.gov>, (Accessed on: June 13, 2012).
- Pratt, W. and H. Wasserman, 2000. QueryCat: Automatic categorization of MEDLINE queries. Proceeding of the AMIA Symp, pp: 655-659.
- Rindflesch, T. and A. Aronson, 1993. Semantic processing in information retrieval. Proceeding of the 17th SCAMC, pp: 611-615.
- Rindflesch, T. and A. Aronson, 1994. Ambiguity resolution while mapping free text to the UMLS metathesaurus. Proceeding of the 18th SCAMC, pp: 240-244.
- Rindflesch, T. and A. Aronson, 2002. Semantic Processing for Enhanced Access to Biomedical Knowledge, In: Kashyap, V. and L. Shklar, (Eds.), Real World Semantic Web Applications. John Wiley & Sons, pp: 195, Amsterdam, ISBN: 427490556X.
- Rogers, W., F. Lang and C. Gay, 2012. MetaMap Data File Builder. U.S. National Library of Medicine, Retrieved from: <http://metamap.nlm.nih.gov>, (Accessed on: June 13, 2012).
- Tuttle, M., N. Olson, K. Keck, W. Cole, M. Erlbaum, D. Sherertz *et al.*, 1998. Metaphrase: An aid to the clinical conceptualization and formalization of patient problems in healthcare enterprises. Meth. Inf. Med., 37(4-5): 373-83.
- Wright, L., 1998. Hierarchical concept indexing of full-text documents in the unified medical language system information sources map. J. Am. Soc. Inform. Sci., 50(6): 514-523.

End notes:

- 1: www.nlm.nih.gov/research/umls
- 2: www.nlm.nih.gov/mesh
- 3: www.rsna.org/redlex
- 4: <http://www.nlm.nih.gov/>
- 5: lhncbc.nlm.nih.gov