

Recent Enhancements to the NLM Medical Text Indexer

James G. Mork¹, Dina Demner-Fushman¹, Susan C. Schmidt¹, Alan R. Aronson¹

¹National Library of Medicine, Bethesda, MD, USA
{jmork, ddemner, schmids, alaronson}@mail.nlm.nih.gov

Abstract. The main goal of the US National Library of Medicine (NLM) Indexing Initiative is to explore indexing methodologies that may help the NLM indexing staff keep pace with the ever increasing challenges of indexing over 700,000 MEDLINE citations each year using a vocabulary of over 27,000 MeSH Descriptors and 220,000 MeSH Supplementary Concept Records. The BioASQ Challenge has been a tremendous benefit by expanding our knowledge of other indexing systems, specifically the technologies used in those systems to identify relevant indexing for biomedical literature. This paper provides an update on improvements to NLM's Medical Text Indexer (MTI) functionality and performance since the first BioASQ Challenge. We have, in a limited way, applied some of the lessons learned from that first Challenge to MTI to assess what performance gains we might see. The research discussed at the 2013 BioASQ Challenge Workshop inspired us to make changes to MTI that have resulted in a 2.69 (4.44%) increase in Precision and very little change in Recall.

Keywords: Indexing methods, Text categorization, MeSH, MEDLINE

1 Introduction

The NLM Medical Text Indexer (MTI) system [1] combines human NLM Index Section¹ expertise and Natural Language Processing technology to curate the biomedical literature more efficiently and consistently. MTI is the main product of the Indexing Initiative [2] and has been providing indexing recommendations based on the Medical Subject Headings (MeSH[®])² vocabulary since 2002. MEDLINE indexers and revisers consult MTI recommendations for approximately 64% of the articles they index. In 2011, NLM expanded MTI's role by designating it as the first-line indexer (MTIFL) for a few journals; today the MTIFL workflow includes over 160 journals and continues to increase. For MTIFL journals, MTI provides the initial indexing for an article that is then reviewed and completed by a human indexer.

Beyond use by the Index Section staff, MTI recommendations have been customized for specific applications in the Cataloging³ and History of Medicine Division (HMD)⁴

¹ <http://www.nlm.nih.gov/bsd/indexhome.html>

² <http://www.nlm.nih.gov/pubs/factsheets/mesh.html>

³ <http://www.nlm.nih.gov/tsd/cataloging/mainpage.html>

systems at NLM. While the main application of MTI remains the generation of MeSH indexing recommendations by processing MEDLINE citations⁵ consisting of identifier, title, and abstract, MTI is also capable of processing any biomedical text. MTI identifies what it calculates as the most relevant MeSH Terms that best describe the biomedical text being processed. This resulting list of MeSH Terms is presented in highest to lowest relevancy order by MTI.

MTI consists of two main methods of identifying potential recommendations for the text being processed:

- **MetaMap Indexing (MMI)**⁶ uses the MetaMap [3] program to identify, summarize, and rank the UMLS[®] Metathesaurus^{®7} concepts in the text to be processed. The UMLS concepts are converted or mapped to potential MeSH Term recommendations using the Restrict to MeSH [4] mapping algorithm.
- **PubMed Related Citations (PRC)**⁸ method [5] uses a modified k-Nearest Neighbors (k-NN) algorithm to identify citations that are closely related to the text being processed. MTI adds some of the indexed MeSH Terms from these related citations to the list of potential recommendations.

In post-processing, MTI combines and ranks the lists of potential MeSH Terms from these two methods, includes recommendations based on various lookup lists, reviews and filters MeSH Terms according to NLM Indexing rules, and finally assigns sub-headings when possible.

2 MTI Enhancements

The Indexing Initiative team explored several different avenues for improving MTI performance this year, mainly focusing on improving Precision. The biggest improvement came from our Vocabulary Density study which looks at the frequency of all MeSH Term usage in MEDLINE over the last five years. Following the Vocabulary Density study, we focused on cleaning up ambiguous and irrelevant MTI recommendations by examining some of the worst performing MeSH Terms.

Vocabulary Density: The inspiration for using journal information to improve MTI performance came from the discussion at the 2013 BioASQ Workshop⁹ by Tsoumakas et al. [6] and one of our senior indexers who recommended that we explore journal-specific indexing and filtering. Tsoumakas et al. used machine learning to train on

⁴ <http://www.nlm.nih.gov/hmd/index.html>

⁵ <http://www.nlm.nih.gov/bsd/mms/medlineelements.html>

⁶ <http://ii.nlm.nih.gov/MTI/Details/mmi.shtml>

⁷ <http://www.nlm.nih.gov/pubs/factsheets/umlsmeta.html>

⁸ <http://ii.nlm.nih.gov/MTI/Details/related.shtml>

⁹ <http://www.bioasq.org/>

only the specific journals that were involved in the BioASQ Challenge and focused on which MeSH Terms and how many MeSH Terms each journal typically used. To explore whether customizing the indexing for a specific journal would be worthwhile, we created the Vocabulary Density study. The study looked at a corpus of 3,401,111 citations that were indexed in the last five years representing 6,606 individual journals from the 2014 MEDLINE Baseline¹⁰. This final, cleaner corpus was the result of filtering out the following list of undesirable citation types from the Baseline.

- Citations without MeSH Terms,
- Citations where automatically assigned MeSH Terms were added without indexer review. This included OLDMEDLINE¹¹ citations (MEDLINE citations indexed prior to 1966) and citations with one or more of the following Comment Types: CommentOn, ErratumFor, PartialRetractionOf, RetractionOf, RepublishedFrom, and UpdateOf.




Fig. 1. Cumulative Distribution of MeSH Heading Use Across Journals

¹⁰ <http://www.nlm.nih.gov/databases/journal.html>

¹¹ http://www.nlm.nih.gov/databases/databases_oldmedline.html

We found that on average, journals used only 999 of the 27,149 potential MeSH Terms (3.68%). The maximum usage of MeSH Terms found was for the *PLoS One (Public Library of Science)* journal which used 17,501 (64.46%). 83.81% of the MeSH Terms were found to have been used by 500 or fewer journals, with 271 MeSH Terms only being used by a single journal. For example, the MeSH Term *Insulin, Lente* has only been used by *The Veterinary clinics of North America. Small animal practice* journal in the corpus. Fig. 1 shows this cumulative distribution of MeSH Term usage across the journals. The most utilized MeSH Term is *Humans* which was used by 6,210 of the 6,606 journals in our study. This selective use of MeSH Terms by the journals confirms the idea that taking into account journal-specific information might lead to improvements in MTI.

There are also 709 MeSH Terms that were found to not have been used in our corpus. These unused Terms were comprised of a combination of new MeSH Terms that have not yet been indexed (e.g., *Anticholinergic Syndrome*), MeSH Terms that are used only by Cataloging (e.g., *Bibliography, National*), Publication Types¹² a special type of MeSH Term that we did not include in this study (e.g., *English Abstract*), MeSH Terms no longer used in current biomedical literature (e.g., *Etioporphyrins*, last indexed in an article published in 1993), MeSH Terms that are strictly category placeholders describing terms below them in the MeSH Tree (e.g., *Hemic and Lymphatic Diseases*), and infrequently used MeSH Terms (e.g., *Swayback*).

To build the current version of the *Vocabulary Density Method*, we removed information on the journals that had fewer than 80 articles over the last five years to ensure we had a baseline level of confidence in the results. We then captured the following information for each MeSH Term for each journal with the requisite number of articles: Frequency of occurrence (freq) and Total citations indexed for the journal (tot). We then calculated a normalized Frequency Factor for each MeSH Term in each journal using the following formula: Frequency Factor = freq / tot. For example, the MeSH Term *Kidney* was found 28 times in the 2,231 citations for the journal *Biochemical Society (Great Britain)* in our corpus. The Frequency Factor for this MeSH Term in this journal is 0.012550 (28/2231).

We are still in the early stages of understanding and using this information, but we have created a simple set of rules to do a preliminary analysis of the effectiveness of the data. We created three rules for removing terms not indexed by a journal over the last five years and for adding terms MTI would not have recommended but which were used by the journal regularly during the same period.

¹² <http://www.nlm.nih.gov/mesh/features2003.html>

The following set of simple rules provided us with a 2.69 (4.44%) improvement in Precision, 1.36 (2.23%) increase in F_1 score, and a 0.05 (0.08%) increase in Recall:

1. If the journal has valid MeSH Term usage and the MeSH Term in question has not been used in the last five years by this journal, we remove the MeSH Term – unless this is a new MeSH Term.
2. For a non-CheckTag MeSH Term with Frequency Factor > 0.74 , we automatically add the term as a MTI recommendation.
3. For a CheckTag¹³ MeSH Term with the Frequency Factor of 1.00 (i.e. used for indexing all of a journal's articles), we automatically add the term as an MTI recommendation. CheckTags are a special type of MeSH Term that are required to be included for each article and cover species, sex, human age groups, historical periods, pregnancy, and various types of research support (e.g., *Male*).

Further work needs to be done to see how we can expand our use of the Frequency Factor in filtering out irrelevant recommendations and adding confidence to recommendations. We also need to decide how to allow MeSH Terms used by a journal for the first time.

Ambiguous Term Identification and Filtering: We invested a considerable amount of effort looking at ambiguous terms causing what we call “Out of the Ballpark” (OOTB) incorrect recommendations. We reviewed over 160 MeSH Terms from across almost all MeSH Tree Categories because of this ambiguity issue. OOTB refers to MTI recommendations that are not closely (within same MeSH Tree Category) related to any of the actual human indexing that was used in an article. For example, if the article is about a *3-arm clinical trial* and MTI recommends *Arm*. *Arm* would be considered an OOTB term since it is completely unrelated to any of the final indexing. Ambiguity is the primary cause of why MTI recommends an OOTB. The types of such ambiguity include:

- **Metaphorical ambiguity** (e.g., *birds of a feather working group* triggering *Birds* and *Feathers*),
- **Brand Name Ambiguity** (e.g., *commit murder* triggering *Tobacco Use Cessation Products* because *Commit* is a brand name),
- **Psychology Term Ambiguity** (e.g., *employee retention* triggering *Retention (Psychology)*), and
- **Body Part/Disease Tree Ambiguity** (e.g., article title says “Ankle joint” triggering *Ankle*, but, the article discusses “sprained ankles” triggering *Ankle Injuries*). The indexer would use the more specific *Ankle Injuries* here and ignore *Ankle*.

During the course of this review, we discovered that many of the terms we were classifying as OOTB were in fact related to this last type of ambiguity, “Body

¹³ <http://www.nlm.nih.gov/mesh/features2003.html>

Part/Disease Tree Ambiguity” and not as egregiously incorrect as the earlier example of *3-arm clinical trial*. We corrected as much of this ambiguity as we could by manually reviewing the text triggers responsible for each of the OOTBs and adding filters to MTI where appropriate (e.g., if the trigger word is *fruit*, make sure text does not contain *fruit fly*, *fruit flies*, *fruit bat(s)*, or *fruit tortrix* before recommending *Fruit*). We also established a series of rules to help with the “Body Part/Disease Tree Ambiguities”. We were able to eliminate 10.92% of the OOTB terms being erroneously recommended with very little loss of Recall in our current test collection.

3 MTI Training and Processing Information

Training or refining the MTI program is an ongoing task. To help verify that any proposed changes to MTI are beneficial, we created the MTI Test Collection. The test collection is completely replaced each year to reduce the tendency to overtrain on the data and to reflect current indexing practices. The current test collection consists of 143,658 citations that were indexed between mid-November 2013 and the end of January 2014.

We process approximately 4,000 new citations each night that we run on our Scheduler¹⁴ pool of 169 Linux clients. The processing takes 10 to 15 minutes depending on what other demands there are on the Scheduler and any problems that arise. We also process approximately 7,000 old and new records for Cataloging and HMD each night which requires around 30 minutes. Overall, MTI processed 45,468,245 items of text in 2013 from our work and from researcher requests around the world.

Training also involves approximately one day of work updating the MTI databases twice a year to incorporate new releases of the UMLS Metathesaurus to verify that we are using the latest data available.

4 MTI Performance in 2014 BioASQ Challenge

The 2014 BioASQ Challenge consisted of a dry run batch and then three batches of five test files made available each Monday morning for a total of 16 files between January 27, 2014 and May 19, 2014. There were between 25 and 45 systems from an unknown number of organizations participating in each of the weekly batch runs with some organizations submitting results for several different systems. MTI performed relatively well and was one of the top tier systems in the first couple of weeks, then dropped down into the middle tier of the systems for the remainder of the Challenge.

We submitted results for two different systems with the primary system being MTI with MTIFL (MTI First Line Index) filtering turned on and the second system using the default settings for MTI (Default MTI). MTIFL filtering uses MTI's *Balanced*

¹⁴ <http://ii.nlm.nih.gov/Scheduler/Scheduler/index.htm>

Recall/Precision Filtering option [1] providing a smaller, more precise indexing list than with Default MTI processing. Table 1 shows preliminary results of the 2014 BioASQ Challenge for our two systems as of May 23, 2014. The table details the results for each of the weekly runs for both of the systems. We include information on Micro Precision (MiP), Micro Recall (MiR), Micro F-Measure (MiF), the number of MEDLINE citations to be processed in each batch (#Cit), the number of citations that were completed (received indexing) as of May 23, 2014 (#Comp), and the percentage completed (%) for each run. The BioASQ team has provided additional results across many more categories that are explained in their Evaluation Framework Specifications [7] document.

Overall, the results are comparable to what we see internally for MTI. For the Challenge, the Default MTI F-measure is slightly higher than the MTIFL F-measure due to the filtering preference of Precision over Recall for MTIFL. Default MTI also has a bias towards Precision over Recall, but, we don't reduce the list of recommendations as much for Default MTI (average 11.30 recommendations) as we do for MTIFL (average 9.51 recommendations). MTIFL is also more customized for use on specific journals.

5 Future Directions

Several research topics that are planned for the future include:

- expand the use of the Vocabulary Density study Frequency Factors,
- identify whether author/publisher supplied keywords might benefit MTI,
- expand machine learning usage to help improve problematic MeSH Headings,
- expand the number of MTIFL journals, and
- extend the Vocabulary Density study to include subheadings assigned to each of the MeSH Terms in each of the Journals.

Acknowledgements

The Medical Text Indexer Team continues to benefit from a very close collaboration with the NLM Index Section as evidenced by one of the authors (SCS) being a senior indexer and reviser. This collaboration provides a deeper understanding of the human indexing process and insights into other possible avenues where MTI might be used to assist in the indexing process at NLM.

This work was partly supported by the Intramural Research Program of the NIH, National Library of Medicine.

Table 1. Preliminary BioASQ Results for Default MTI and MTIFL as of May 23, 2014

Batch	Week	System	MiP	MiR	MiF	#Cit	#Comp	%
Dry Run		Default MTI	0.5682	0.5695	0.5689	3,186	2,515	78.94%
		MTI First Line Index	0.6060	0.5268	0.5636			
1	1	Default MTI	0.5825	0.5574	0.5697	4,440	3,227	72.68%
		MTI First Line Index	0.6128	0.5149	0.5596			
1	2	Default MTI	0.5838	0.5556	0.5694	4,721	3,474	73.59%
		MTI First Line Index	0.6171	0.5080	0.5573			
1	3	Default MTI	0.5930	0.5592	0.5756	4,802	3,643	75.86%
		MTI First Line Index	0.6304	0.5177	0.5685			
1	4	Default MTI	0.5859	0.5658	0.5757	3,579	2,183	60.99%
		MTI First Line Index	0.6232	0.5237	0.5691			
1	5	Default MTI	0.5805	0.5413	0.5602	5,299	3,478	65.64%
		MTI First Line Index	0.6126	0.4982	0.5495			
2	1	Default MTI	0.6028	0.5530	0.5769	4,085	3,250	79.56%
		MTI First Line Index	0.6337	0.5111	0.5658			
2	2	Default MTI	0.5771	0.5698	0.5734	3,496	2,506	71.68%
		MTI First Line Index	0.6097	0.5290	0.5665			
2	3	Default MTI	0.5992	0.5485	0.5727	4,524	3,076	67.99%
		MTI First Line Index	0.6310	0.5087	0.5633			
2	4	Default MTI	0.5950	0.5601	0.5771	5,407	3,635	67.23%
		MTI First Line Index	0.6273	0.5178	0.5673			
2	5	Default MTI	0.5974	0.5555	0.5757	5,454	3,237	59.35%
		MTI First Line Index	0.6273	0.5116	0.5635			
3	1	Default MTI	0.5838	0.5623	0.5729	4,342	2,691	61.98%
		MTI First Line Index	0.6180	0.5220	0.5659			
3	2	Default MTI	0.5848	0.5356	0.5591	8,840	4,394	49.71%
		MTI First Line Index	0.6181	0.4923	0.5481			
3	3	Default MTI	0.6138	0.5712	0.5917	3,702	1,605	43.35%
		MTI First Line Index	0.6434	0.5305	0.5815			
3	4	Default MTI	0.5959	0.5402	0.5667	4,726	918	19.42%
		MTI First Line Index	0.6277	0.4951	0.5535			
3	5	Default MTI	0.5674	0.5967	0.5817	4,533	249	5.49%
		MTI First Line Index	0.5925	0.5384	0.5641			

References

1. J.G. Mork, A. Jimeno Yepes, A.R. Aronson. The NLM Medical Text Indexer System for Indexing Biomedical Literature. BioASQ. 2013.
2. Aronson AR, Bodenreider O, Chang HF, Humphrey SM, Mork JG, Nelson SJ, Rindflesch TC, and Wilbur WJ. The NLM Indexing Initiative. Proc AMIA Symp 2000;:17-21.

3. Aronson AR, Lang FM: An overview of MetaMap: historical perspective and recent advances. *Journal of the American Medical Informatics Association* 2010, 17(3):229-236.
4. Bodenreider O, Nelson SJ, Hole WT, and Chang HF. Beyond Synonymy: Exploiting the UMLS Semantics in Mapping Vocabularies. *Proc AMIA Symp* 1998;:815-9.
5. Lin, J., & Wilbur, W. J. (2007). PubMed related articles: a probabilistic topic-based model for content similarity. *BMC bioinformatics*, 8(1), 423.
6. Tsoumakas G, Laliotis M, Markantonatos N, Vlahavas I. Large-scale semantic indexing of biomedical publications at BioASQ. *BioASQ Workshop*, Valencia, Spain, September 27, 2013.
7. Balikas, G., I. Partalas, A. Kosmopoulos, S. Petridis, P. Malakasiotis, I. Pavlopoulos, I. Androutsopoulos, N. Baskiotis, E. Gaussier, T. Artieres, et al., "Evaluation Framework Specifications", BioASQ, Project deliverable D4.1, 05/2013.