

Medical-Miner at TREC 2011 Medical Records Track

¹J.M. Córdoba, ¹M.J. Maña, ¹N.P. Cruz, ¹J. Mata, ²F. Aparicio,
²M. Buenaga, ³D. Glez-Peña, ³F. Fdez-Riverola

¹Universidad de Huelva

²Universidad Europea de Madrid

³Universidad de Vigo

Abstract

This paper presents work done at Medical Minner Project on the TREC-2011 Medical Records Track. The paper proposes four models for medical information retrieval based on Lucene index approach. Our retrieval engine used an Lucen Index scheme with traditional stopping and stemming, enhanced with entity recognition software on query terms. Our aim in this first competition is to set a broader project that involves the develop of a configurable Apache Lucene-based framework that allows the rapid development of medical search facilities. Results around the track median have been achieved. In this exploratory track, we think that these results are a good beginning and encourage us for future developments.

1 Introduction

Clinical databases have accumulated large quantities of information about patients and their medical conditions. Relationships and patterns within this data could provide new medical knowledge. Unfortunately, few methodologies have been developed and applied to discover this hidden knowledge. In the TREC framework, Medical Records track encourages to research on providing content-based access to the free-text fields of electronic medical records. Related with this same research area, the Medical Minner initiative¹ works for provide new useful tools in this environment, enabling the achievement of a double benefit: first, to provide specific and relevant mechanisms to a professional environment (biomedical), and secondly, obtaining innovative techniques and significant advances in research in the area of information retrieval, particularly on this domain, and based on a significant integration effort. In this context, this paper presents the models developed for our first participation in TREC, particularly in the Medical Records track.

¹<http://www.medicalminer.org>

This paper is organised as follows. In Section 2 we describe some of the ways have been used to represent the documents fields in reports processing, and describe the way the index used in our later experiments is built automatically from document collection. In Section 3 we cover the overall architecture of a indexing application and where Lucene fits. The document retrieval models described in Section 4. This section, also covers the further development and elaboration of the models. Data and results tables are given in Section 5 with key results summarised. Finally, in Section 6, some conclusions are provided.

2 Document Collection Preprocessing

Assuming the system has access to a large document collection as a knowledge resource for answering medical topics, this collection may need to be processed before querying, in order to transform it into a form which is appropriate for topic answering. Thus, in order to effective TREC medical visits retrieval, some issues need to be solved. First, reports need to be grouped into visits. Next, XML Reports must be adapted to be queried in a rapid and accurate way. Additionally, another type of additional processing can be performed. In our case, we used the ICD codes for better retrieval.

2.1 Visit Processing

As the Medical Records track will use the visit as the answer unit then is there need for additional processing in order to group the reports into visits. The Medical Records Track repository contains a simple ASCII table called the Report-to-Visit Mapping Key that specifies which reports belong to the same visit. This table is sorted per visit ID. Once the table is ordered, the visits reports can be processed sequentially. As a precaution, checking procedures have been implemented to control null visits, number of reports per visit, total visits, etc

2.2 XML Parsing

We need to quickly search a collection of XML documents, and, to do this, we need to create an index of terms keeping track of the context in which these terms appear. Our solution uses Apache Lucene and the Java SAX library in order to create an index of Lucene Document objects for the lowest level of granularity to search. In order to have a testing framework as complete as possible, all reports fields have been processed.

2.3 ICD9 Mining

The International Statistical Classification of Diseases and Related Health Problems (most commonly known by the abbreviation ICD) is a medical classification that provides codes to classify diseases and a wide variety of signs, symptoms,

abnormal findings, complaints, social circumstances, and external causes of injury or disease. The Collection reports includes fields with discharge diagnosis in ICD9 codes form. In our approach, codes are processed to include a textual description to gather additional information. Thus, when a ICD9 code is detected in report processing by SAX the code description is searched in a database and automatically added to the index.

3 Indexing

Model main component is a search engine based on Apache Lucene. Lucene is a powerful Java library that lets you easily add document retrieval to any application. In recent years Lucene has become exceptionally popular and is now the most widely used information retrieval library. What we end up with after running Lucene is a directory named index, which contains files used by Lucene to associate terms with documents. To accomplish this, a Lucene index was created with a specific analyzer model-dependent. An Analyzer takes a series of terms or tokens and creates the terms to be indexed. A unique kind of Lucene index has been used for all developed models, or in other words, all models share the same Lucene index.

Documents and fields are Lucene’s fundamental units of indexing and searching. A document is Lucene’s atomic unit of indexing and searching. It is a container that holds one or more fields, which in turn contain the “real” content. Each field has a name to identify it, a text or binary value, and a series of detailed options that describe what Lucene should do with the field value when you add the document to the index. To index our collection sources, we must first translate it into Lucene’s documents and fields. The indexing module takes from preprocessed collection all visit reports into a single Lucene document. Lucene allows duplicate fields to be added to a Document. This can make updating documents easy because it allows to add fields of various reports into a single Lucene document, with a mixing of all visit reports in a single document Lucene. At the end, we have the visit reports sharing the same Lucene document and the fields of the various reports sharing the same Lucene field.

4 Retrieval approaches

This section presents the different models developed for evaluation. Among the models developed four have been selected for submission: Baseline², Medical Face³, BWRTDDD⁴ and BWRTDDD2⁵. From the four topics samples review, we consider using some alternative models to the final runs. Although discard some models, we believe interesting to include a review of these unused models that can have interesting for scientific community.

²UHU1BL TREC run ID.

³UHU2MFB TREC run ID.

⁴UHU3BWRTDDD TREC run ID.

⁵UHU4BWRTDDD2

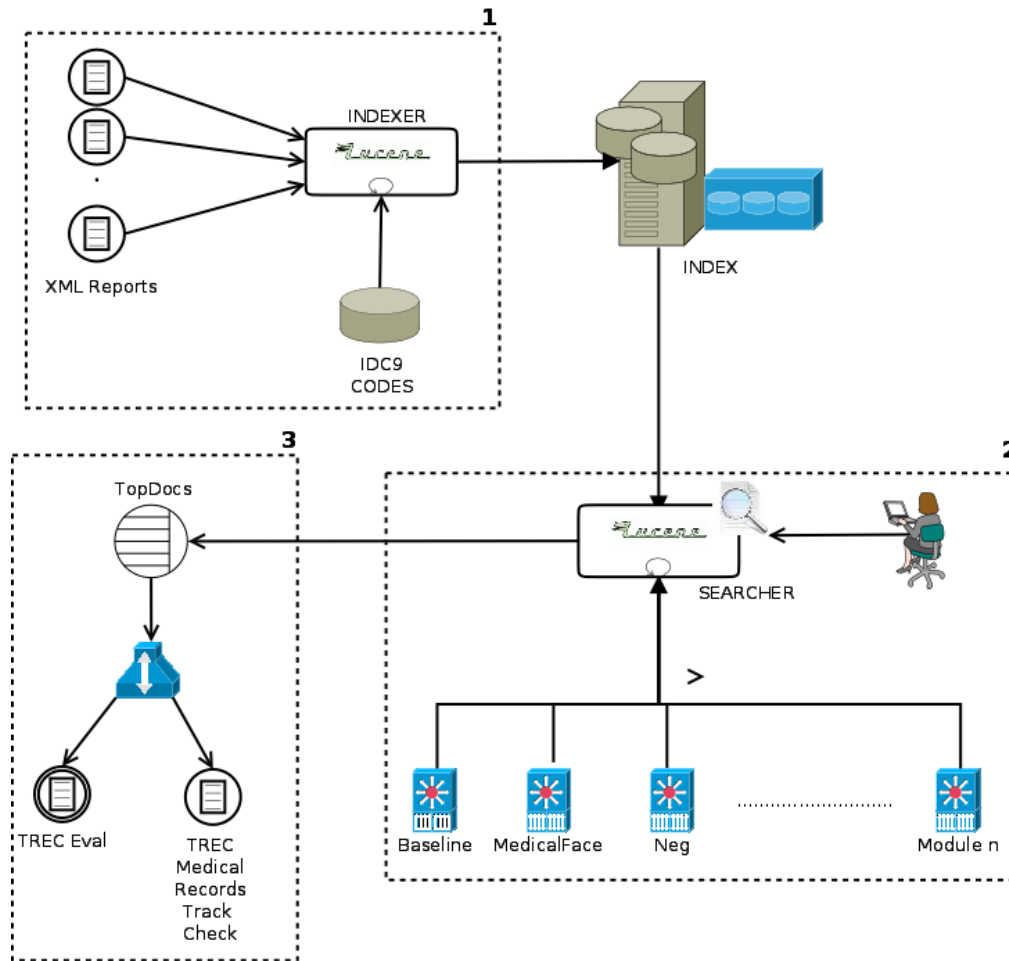


Figure 1: Prototype architecture.

The retrieval module completes the visits retrieval process. The subsequent processes of assessment and results extraction have been implemented according to the TREC eval utilities. The overall process can be seen in Figure 1.

4.1 Baseline

This model has been designed to be the simplest approximation to the task. The Baseline model is based on the method of the bag of words. In this model, topics text is represented as an unordered collection of words, disregarding grammar and even word order. Therefore, the model matches the words in the topic with the words contained in the index. The usefulness of the model is twofold: provides a basis results for comparing and, on the other hand, its code serves as a basis for implementing more complex models. In order to maintain the simplicity, the baseline model makes match the topics words with the words contained in the field “report text”.

To develop this model we used a very simple analyzer. The analyzer takes only the topic and provides a set of terms to be searched in the index. Our base analyzer discards stops words with little semantic value, such as "the", "a", "an", "for",.... Cutting down on the number of terms indexed can save time and space in an index, but it can also limit accuracy.

4.2 Medical Face

MedicalMiner initiative brings together three research groups from Huelva, Vigo and Madrid European Universities. MedicalMiner initiative proposes to analyse, experiment and develop new text and data mining techniques in an interrelated way, in intelligent medical information systems. An intelligent information access system based on them will be developed, offering advanced functionalities able to interrelate medical information, mainly information (text and data) from clinical records and scientific documentation, making use of standard resources of the domain (e.g. UMLS, SNOMED, Gene Ontology). An open source platform will be developed integrating all the elements [2]. A main component of this platform is Medical Face⁶. Medical Face is a Biomedical text annotation system using multiple ontologies. In its simplest form, Medical Face takes a input text and makes medical entity recognition over this text.

For our approach, Medical Face has been used to develop an analyzer that extracts relevant entities in the medical domain from the topics. This relevant entities are used to search for relevant documents in the generated index during indexing process. In this model, only the “report text” field is taken account.

4.3 BWRTDDD

The BWRTDDD model is based on the method of the bag of words in a similar way of the baseline model. Unlike the baseline model this model takes into

⁶<http://orion.esp.uem.es:8080/MedicalFace/mt.html>

account the “report text” and the modified discharge diagnosis fields. The modified discharge diagnosis field is a modified version of the field that takes account a textual description of ICD9 codes generated during preprocessing process. In addition, a specific analyzer has been developed for processing the items.

Like our base analyzer, the BWRTDDD analyzer discards common English words with little semantic value, such as "the", "a", "an", "for",..... Additionally, some irrelevant terms are removed from the topic (patient, for example).

4.4 Exploring Age and Gender

In the first exploration of the sample topics, we think that would be interesting to identify genre and age in the topics and reports. It is easy to detect the occurrence of words that help classify subjects by patient gender and/or age (like woman or elderly). While gender can be detected directly in both the topics and reports, age can be estimated from some labels in the topics (young, adult, elderly,...) and reports (**AGE[in XXs]-year-old). With this simple approach we can match those relevant visits to a topic according to age and gender. Unfortunately, this approach has been unsuccessful with the test topics by the small number of examples that show some of these features (only the adult, children, female and women tags have been found by this model).

4.5 Explore Negation

Another way that looked promising was the negation detection. Usually, narrative reports in medical records contain a wealth of information that may augment structured data for managing patient information and predicting trends in diseases. Pertinent negatives are evident in text but are difficult to automatic processing. In our preliminary study of the sample topics, we saw that one of the examples contained a negation.

Our research group have experience in the development of new text mining techniques adapted to bilingual medical domain. Including basic operations of text categorization and information extraction, capable of addressing specific problems of special relevance in the domain, such as the treatment of negation.

In [1] we present a machine learning system that identify the negation and speculation signals in biomedical texts, in particular, in the BioScope corpus.

Using the [1] approach, a prototype to handle negation was developed. This model takes a topic to detect negation presence. If the topic has a negative part then the medical entities in the negative clause are detected. A visit search are performed deleting those visits with presence of these “negative” entities. Unfortunately, no topic with negative part was given, therefore, no model run was sent.

	bpref	R-prec	P@10
Best	0.7763	0.5917	0.9500
Median	0.4219	0.2924	0.4500
Worst	0.0000	0.0000	0.0000
UHU1BL	0.4534	0.3543	0.4000
UHU2MFB	0.4382	0.3103	0.3000
UHU3BWRTDDD	0.3709	0.1657	0.2000
UHU4BWRTDDD2	0.3709	0.1498	0.1000

Table 1: Medical Minner Medical Records results at TREC 2011.

5 Results

This section shows the results for our four runs presented to the track (table 1). All our models results are around the evaluation median. In this exploratory track, we think that these results are a good beginning and encourage us for future developments. In despite of initial results, more analysis work needs to be developed. A deeply per topic study must be make. Due to the different topic types, a harder work should be performed to classify the behavior of each model depending on the topic type.

About our models, the best results were obtained with the Baseline model. Despite being the best model, there are no major differences with other models. When inspecting in detail the topics we have seen that an important number of them (14-15 depending on model) have better results with the alternative models (UHU2MFB, UHU3BWRTDDD, UHU4BWRTDDD2). This performance in certain topics suggests that a better fit in model parameters can give better results if a larger number of sample topics are available.

6 Conclusions

We have presented a novel runs of looking at the problem of medical text retrieval in the Medical Records Track framework. We feel that our models will provide effective retrieval in the order of track median. In our first TREC participation, we feel that track goal has been met reasonably well by the runs proposed.

In the near future, we will make a more detailed study about results. we hope to get more facts about models performance and alternative ways to improve the models.

Apart from the send runs, new ways to improve the medical text retrieval have been proposed. Mainly we have shown our interest in detection of expressions of negation. We think that the detection of negation must be taken into account in the track future.

References

- [1] N.P. Cruz, M.J. Maña, and J. Mata. Aprendizaje automático versus expresiones regulares en la detección de la negación y la especulación en biomedicina. *Procesamiento de Lenguaje Natural*, 45, 2010.
- [2] M. de Buenaga, F. Fdez-Riverola, M. Maña, E. Puertas, D. Glez-Peña, and J. Mata. Medical-miner: Integración de conocimiento textual explícito en técnicas de minería de datos para la creación de herramientas traslacionales en medicina. *Procesamiento de Lenguaje Natural*, 45, 2010.