RePaLi participation to CLEF eHealth IR challenge 2014: leveraging term variation

Vincent Claveau, Thierry Hamon, Natalia Grabar, Sébastien Le Maguer

▶ **To cite this version:**

HAL Id: hal-01027534

https://hal.archives-ouvertes.fr/hal-01027534

Submitted on 22 Jul 2014

RePaLi participation to CLEF eHealth IR challenge 2014: leveraging term variation

Vincent Claveau[1], Thierry Hamon[2], Natalia Grabar[3], and Sébastien Le Maguer[4]

[1] IRISA - CNRS, Rennes, France
vincent.claveau@irisa.fr
[2] LIMSI - CNRS, Orsay, France
Université Paris 13, Sorbonne Paris Cité, France
thierry.hamon@limsi.fr
[3] STL UMR8163 CNRS, Université Lille 3, France
natalia.grabar@univ-lille3.fr
[4] INRIA - IRISA, Rennes, France
sebastien.le_maguer@irisa.fr

Abstract. This paper describes the participation of RePaLi, a team composed with members of IRISA, LIMSI and STL, to the biomedical information retrieval challenge proposed in the framework of CLEF eHealth. For this first participation, our approach relies on a state-of-the-art IR system called Indri, based on statistical language modeling, and on semantic resources. The purpose of semantic resources and methods is to manage the term variation such as synonyms, morpho-syntactic variants, abbreviation or nested terms. Different combinations of resources and Indri settings are explored, mostly based on query expansion. For the runs submitted, our system shows up to 67.40 p@10 and up to 67.93 NDCG@10.

Keywords: Patient-oriented information retrieval, semantic resources, terminological variation, Indri, UMLS

1 Introduction

Since several years, patients have an increasing access to their Electronic Health Records (EHRs)). However the understanding of clinical texts by laymen is often difficult although they are concerned with such important topics like the health condition of patients. The situation leads to an increasing use of the Internet for searching health information [3, 7], and consequently important changes in the doctor-patient communication [11, 2].

In that respect, it becomes crucial that the patients can use information retrieval systems which are able to make links between the specialized vocabulary of the physicians and the Web documents understandable by patients (in other words, to manage specialized queries and provide answers understandable by patients). Moreover, it is also required that such system provides relevant and

trustworthy information [16]. In the framework of CLEF eHealth 2014 [12], one of the proposed shared tasks (task 3) addresses this challenge [8]: queries have been defined from real patient cases issued from the clinical documents provided by the CLEF eHealth 2014's task 2. The participating systems have to return health related web documents previously collected by the KRESMOI project (`http://www.khresmoi.eu`).

Hence, the purpose of this challenge is to find expert documents that contain answers to non-expert medical users' questions, such as those that can be submitted by patients when reading their EHRs. In this context, this paper describes the participation to this biomedical information retrieval (IR) task of our team composed of members from three French labs: IRISA, LIMSI and STL. For this first participation, our team has focused on using a state-of-the-art IR system and worked on different strategies to expand queries with relevant biomedical terms. Only the English query set was considered. In particular, the objective of the participating systems is to guarantee the semantic interoperability and compatibility between the expert and non-expert language, as the former is used in the documents searched while the latter occurs in the questions processed. The corresponding research topics have been addressed in previous work and mainly are related to the aligning of expert and non-expert terms and expressions. An important amount of this work has been done within the Consumer Health Vocabulary [21, 20, 22, 23] or related initiatives [4, 6, 5]. Most of the alignments produced within the Consumer Health Vocabulary initiative are included in the UMLS [15], which motivates its use in our systems.

The paper is structured as follows. In the next section, we present the IR system at the heart of our participation, which is based on language-modeling techniques implemented in the Indri search engine [18]. In Section 3, we present the different resources that are used to process and expand the queries. The description of the submitted runs is given in Section 4, and their results are detailed in Section 5. The last section is dedicated to some concluding remarks and insights about this first participation to this biomedical IR challenge.

2   IR model

The IR system at the heart of our runs is based on statistical language modeling (LM) as implemented by Indri, a toolkit for LM-based IR [18]. This system has shown high performance in numerous IR tasks; in our biomedical context, it also offers interesting capabilities to express complex queries.

2.1   Markov Random Field

Considering a collection of N random variables $x = (x_1, \ldots, x_N)$, a Markov Random Field (MRF) is a graphical probabilistic model which can be represented by an undirected graph $\mathcal{G}$. This graph is defined by $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ where nodes are the random variables ($\mathcal{V} = x$) and edges represent conditional dependencies

between these variables. Therefore a clique $c \in C(\mathcal{G})$ represent a set of dependent variables. Based on this graph, $P(x)$ is defined as follows:

$$P(x) = \frac{1}{Z_\Lambda} \prod_{c \in C(\mathcal{G})} \psi(c; \Lambda) \tag{1}$$

where $Z_\Lambda = \sum_x \prod_{c \in C(\mathcal{G})} \psi(c; \Lambda)$ is a normalization coefficient, $\psi_c(c; \Lambda)$ the potential functions and Λ the model.

In order to use MRF to information retrieval problems, [14] proposed to considered the graph $\mathcal{G}$ composed of query nodes Q and a document node D. Therefore, nodes involved into a clique c, which contains the document node, are queries for which the document is relevant. As the objective is to rank documents, the following score (Retrieval status value, RSV, which represents the similarity between the query and the document considered) is used:

$$RSV(Q, D) = log(P_\Lambda(D|Q)) = log(\frac{P_\Lambda(Q, D)}{P_\Lambda(Q)}) = \sum_{c \in C(\mathcal{G})} log(\psi(c; \Lambda)) \tag{2}$$

The $\psi(c; \Lambda)$ proposed in [14] is defined as $\psi(c; \Lambda) = exp[\lambda_c * f(c)]$ where $f(c)$ is a real-valued feature function and λ_c corresponding to the weight.

For this challenge task, as a basis for our runs, we used the implementation proposed in [14] which considers 3 kinds of neighborhood: the unigram neighborhood, the bigram neighborhood and the unordered bigram in a user defined size window neighborhood.

First, the unigram neighborhood modeling function $f_T(t, D)$ is defined as follows:

$$f_T(t, D) = \frac{nb(t, D) + \mu \cdot P(t|C)}{|D| + \mu} \tag{3}$$

where t is the current term index of a query q_i and μ the Dirichlet smoothing coefficient.

In addition to the unigram t, the bigram $(t, t+1)$ and the unordered bigram $\{t, t+1\}$ in a user defined window of size w are now considered to get, respectively, $f_{Bi}((t, t+1), D)$ and $f_W(\{t_i, t_{i+1}\}, w, D)$. Therefore, the implemented score function is:

$$\begin{aligned} RSV(Q, D) = \lambda_T * \prod_{t \in Q} f_T(t, D) + \\ \lambda_{Bi} * \prod_{i=1}^{|Q|-1} f_{Bi}((t_i, t_{i+1}), D) + \\ \lambda_W * \prod_{i=1}^{|Q|-1} f_W(\{t_i, t_{i+1}\}, w, D) \end{aligned} \tag{4}$$

In a typical setting, these parameters are set to these default values: $w = 8$, $\mu = 2500$, $\lambda_T = 0.85$, $\lambda_{Bi} = 0.1$ and $\lambda_W = 0.05$. In our case, the parameters were adjusted as explained in Section 2.3.

2.2 Preparing the data

For the indexing step, the collection of web pages has been preprocessed in several ways with Python scripts. First, the HTML formatting marks and scripts were removed. In order to do that, we have tested existing libraries (e.g. Beautiful Soup), but finally end up in writing our own HTML-to-text scripts which allow us to better control the speed vs. quality of this important but time-consuming task. Thus, it is worth noting that we do not exploit the structure (titles, subtitles) or the hypertext links in our approach.

Secondly, the resulting text undergoes several other processing steps. For instance, the most common HTML codes (for example, `é`) are replaced by the corresponding UTF-8 character (é). Since the language model takes into account bigrams of terms in w-sized windows, parts of sentences that were separated by the previous processing steps are collected together.

Finally, when building the index with Indri, stemming is performed on the result text. Porter and Krovetz stemming algorithms were both tested on the 2013 dataset; Krovetz yielded the best results and is thus chosen for every experiment reported in this paper. Also, stop-words are removed at this step.

2.3 Setting the parameters

To set the smoothing parameter μ and the combination parameters λ, we used the 2013 query set and the corresponding (binary) relevance judgment. The objective measure to maximize was the MAP. For μ, we systematically explored values from 0 up to 30,000 by steps of 500; for λs, we explored from 0 to 1 by steps of 0.05.

Table 1 sums up the performance of the system with the default parameters (hereafter default) and those maximizing the MAP on the 2013 query set (best). The significant differences illustrate the importance of this parameter optimization step.

Table 1. Performance (%) of the chosen parameter set on CLEF eHealth 2013 dataset

Parameters	MAP	P@5	P@10	P@100	NDCG@5	NDCG@10	NDCG@100
default	26.27	39.20	40.40	14.14	39.67	40.75	44.20
best	30.53	48.95	46.53	15.58	48.95	47.63	49.43

3 Semantic resources for query expansion

The purpose for the generation of the semantic resources is to prepare knowledge required for the enrichment of the queries. We generate several types of semantically equivalent term sets: known synonyms from the existing resources (e.g.

UMLS, see below) such as {*theophyllamine, ammophyllin*}, morpho-syntactic variants such as {*stenotic aorta, stenosis of the aorta*}, hierarchical relations through lexical inclusions such as {*muscle, muscle pain*}, and abbreviations such as {*AHCD, acquired hepatocellular degeneration*}. Such resources convey more or less close semantics of terms and may play an important role for the enrichment of the queries. Besides, a list of stopwords is used. When needed, the processing of data for the generation of the resources is performed through the Ogmios NLP Platform [9].

3.1 Synonyms from the UMLS

The UMLS (Unified Medical Language System) [15] aggregates several biomedical terminologies. In order to align these terminologies, each term is associated with a CUI (Concept unique identifier). Two terms with a same CUI can then be considered as synonyms. In our experiments, the synonymy relations collected in the UMLS are used in different ways to process the queries.

In Run 5, (simple and multi-word) UMLS terms are searched in the queries, and UMLS synonyms are used to expand the queries. When dealing with multi-word terms, it is very common that these terms also contains other terms: for instance, *Crohn's disease* is a multi-word term, but *disease* also appear as a term in the UMLS. To expand the query, we choose to only consider the longest term; so in the previous example, only synonyms of *Crohn's disease* will be considered. Note that overlapping terms are considered independently.

In Run 6 and 7, only simple (single word) terms are considered. 227,887 synonymy relations between single word terms are extracted from the UMLS, which are latter processed as described below.

3.2 Morpho-syntactic variants

Morpho-syntactic variants convey very close semantics because the only modification is related to the word order, syntactic organisation of words and their morphological modifications. Such variants go beyond the stemming as they take into account complex terms and consider the semantic relations between morphologically modified words. Working with FASTR [10] for the identification of morpho-syntactic variants between the terms, additional rules and transformations can be involved. FASTR uses indeed several transformation rules, such as insertion (*cardiac disease/cardiac <u>valve</u> disease*), morphological derivation (*artery restenosis/arter<u>ial</u> restenosis*) or permutation (*aorta coarctation/coarctation of the aorta*).

Several steps are applied on the set test1 (other sets have not been processed because of the time needed):

- segmentation of the test1 and of queries into words and sentences,
- their part-of-speech tagging with TreeTagger [17],
- their syntactic shallow parsing with YATEA [1] for the acquision of noun phrases and term candidates,

– application of FASTR for the acquisition of variants of term candidates: terms extracted from queries are the reference set of terms, while terms extracted from the test1 set are the variations of the queries terms. In this way, we perform the controled indexing.

We extract 284 terms from the 50 queries processed. On the whole, we generate 575 morpho-syntactic variants for these terms.

3.3 Lexical inclusion and hierarchical relations

The lexical inclusion hypothesis [13] states that when a given term is lexically included in another term there is a semantic relation and hierarchical subsumption between them. The semantics of lexical inclusions may be weaker than the semantic of synonyms. Still, hierarchical subsumption relations may be useful for the information retrieval, especially in the context of non-expert information retrieval, during which the non-expert users may submit queries that are less precise and specific than terms used in the scientific literature.

The processing is done in two steps:

– the terms extracted by YATEA are syntactically analyzed into head and expansion components. For instance, the syntactic analysis of the term *muscle pain* results in two components: head component *pain* and expansion component *muscle*;
– the semantic relation is then established between a given term and its head component. For instance, there is semantic relation between *muscle pain* (the whole term) and *pain* (the head component of the term). With these specifications, the identified relations are hierarchical: the long term *muscle pain* is the hierarchical child of the short term *pain*. Indeed, *muscle pain* conveys more specific information.

This process is applied to the whole set of data (part1 to part8). First, we perform the segmentation into words and sentences. We then apply the part-of-speech tagging with TreeTagger [17] and the syntactic shallow parsing with YATEA [1] for the acquision of noun phrases and term candidates. Finally, we induce the lexical inclusion relations further to the decomposition of complex terms (containing more than one word) in syntactic head and modifier components. As explained earlier, we build pairs {*syntactic head, complex term*} that convey hierarchical relations.

On the whole set of the data, we obtain 1,114,959 such pairs of terms.

3.4 Abbreviations

The abbreviations are very frequent in the biomedical literature, both clinical and scientific) and we expect that both their short and expanded forms may be helpful for information retrieval. The resource with 1,897 abbreviations has been built from information available online.

3.5 English stopwords

We use the list of 627 stopwords in English. These are grammatical words and also words that are very frequent and common in the biomedical documents (e.g., *accordance, amongst, indicate...*).

4 Run description

This section sums up the four runs submitted for the challenge by our team. They all rely on the same index (as described in Section 2.2) since no modification to the document representation is performed. As we were instructed, only the 1 000 first results (highest RSV) of each query is kept.

4.1 Run 1

The goal of this run is to measure the performance of our IR system without any modification or use of biomedical resources. Therefore, it only consists in running the Indri search engine (cf. Section 2.1), with the parameters estimated on the 2013 dataset as described in Section 2.3.

4.2 Run 5

For this run, the queries are expanded with synonyms found in the UMLS. As it was explained in Section 3.1, the expansion is performed as follows:

1. UMLS terms of maximal length are searched within each query;
2. synonyms are searched in the UMLS (terms with the same CUI);
3. synonyms are added to the initial query.

It is important to note that a term can be ambiguous and thus have several CUIs. Since we do not perform a sense disambiguation, we may expand the queries with synonyms that do not correspond to the initial sense of the term. In order to limit the effect of adding unrelated terms, or more generally in order to prevent the expanded query to be too far from the initial information need of the user, the expansion is given a weight much lower than the initial part of the query. Here again, we rely on the 2013 data to set this weight: we found that the expansion weight should not exceed 0.1 (considering 1 as the weight of the initial query).

4.3 Run 6 and 7, and other non submitted runs

Several other runs were prepared, including those submitted as Run 6 and 7. They follow the same principle than Run 5 but the queries undergo a more complex processing in order to parse them and to enrich them semantically with the help of all the semantic resources described in Section 3. As it was previously explained, for these runs, the processing of queries is performed through the Ogmios NLP Platform [9].

More precisely, the following steps are applied to process the queries:

- segmentation of queries in words and sentences, if relevant,
- part-of-speech tagging with TreeTagger [17],
- syntactic shallow parsing with Y$_A$T$_E$A [1] for the acquision of noun phrases and term candidates.

The queries are then enriched with the resources built and prepared for this task. The resources are used in combination with each other. In every run, the stopwords are removed. Several runs have been generated:

- *expanded*: expansion with UMLS synonyms,
- *expanded-ab*: expansion with UMLS synonyms and abbreviations,
- *expanded-ab-terms*: expansion with UMLS synonyms, abbreviations, and the complex terms extracted,
- *expanded-ab*-FASTR: expansion with UMLS synonyms, abbreviations, and FASTR morpho-syntactic variants,
- *expanded-ab-IL*: expansion with the UMLS synonyms, abbreviations, and the lexical inclusion relations in the generalization way, such as {*crohn disease*, *disease*}, where *crohn disease* is reduced to *disease*,
- *expanded-ab-IL-r*: expansion with the UMLS synonyms, abbreviations, and the lexical inclusion relations in the specification way, such as {*disease*, *crohn disease*}, where *disease* is specified to *crohn disease*.

Due to the limitations imposed by the organizers on the number of runs, only two of among these runs were actually submitted: *expanded* as Run 6 and *expanded-ab-IL* as Run 7. However, experiments not reported here on the 2013 dataset show only slight differences in terms of global performance between these runs (except for *expanded-ab-IL-r*, see below). Yet, it should be also highlighted that for some particular queries, the differences may be important. Another lesson learned from the 2013 dataset is that the specification link used in the expansion strategy *expanded-ab-IL-r* does not yield good results, as the expanded may be semantically too far from the original user's information need.

5 Results

5.1 Run 1

In Figure 1, the IR performance measures provided by the 2014 official evaluation are displayed. They are compared with the 2013 query set results. Note that the two systems are identical (in particular, same parameters of the RSV function), as well as the index. Yet, one can see that the 2014 results are significantly better than the 2013 ones. Such a difference is difficult to explain; the 2014 query set may be easier than the 2013 one, or the pooling process used to build the relevance judgment may be more suited this year with more homogeneous systems provided by the participants.
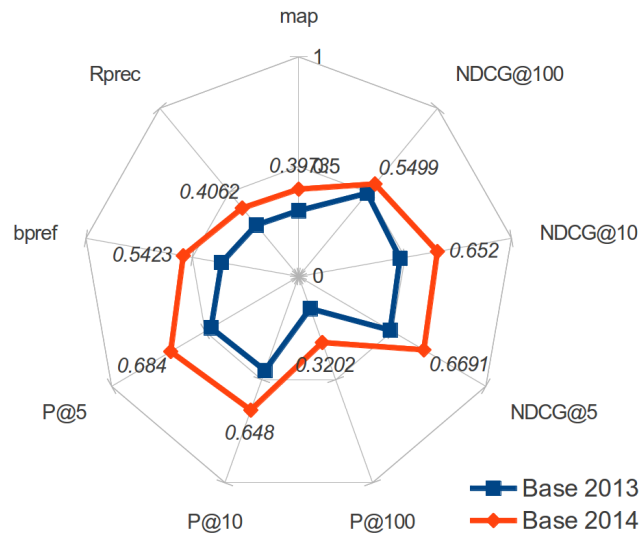
Fig. 1. Performance of Run 1 setting on the 2013 and 2014 query sets

5.2 Run 5, 6 and 7

The official results of Run 5, as well as Run 6 and Run 7 are reported in Figure 2. It appears that all the runs obtain results very similar to those of Run 1. This is expected since the query expansions is down-weighted compared to the initial query (as it was explained in Section 4.2, giving more weight to expansion tends to degrade the results on the 2013 dataset). Yet, Run 5 obtains slightly better results than our other runs. Moreover, the relative ranking of our ranks is identical whatever the evaluation measure used.
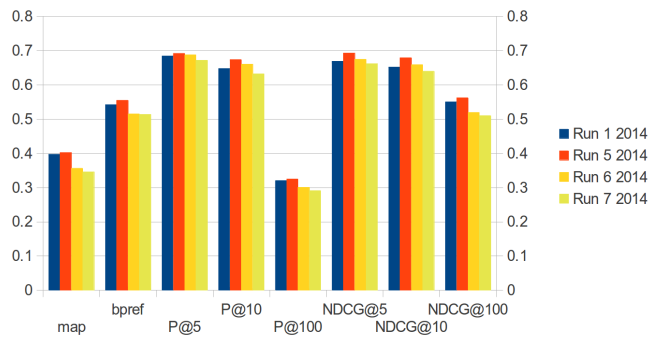

Fig. 2. Performance of Run 5, 6 and 7, compared with Run 1 (official results)

In order to examine the effect of the query expansion strategy used for Run 5, we compare its results to Run 1 in Figure 3. For each query, it represents the gain or loss compared with the median result from all participants. The effect of
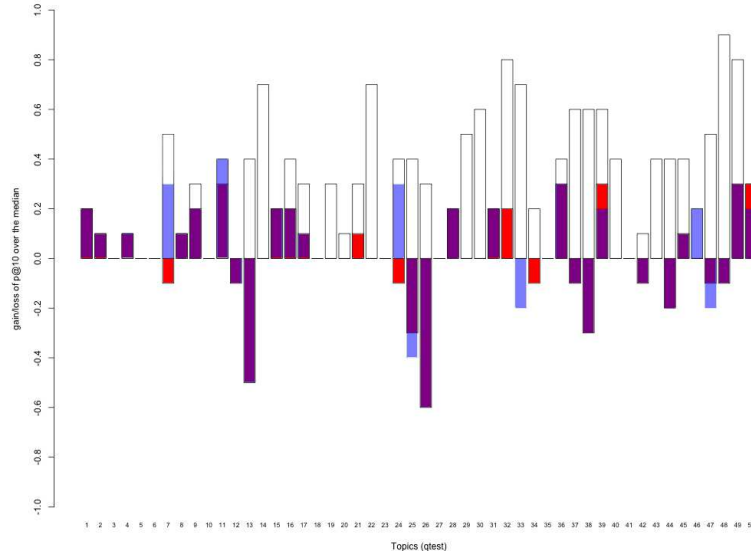

Fig. 3. Query expansion effect of Run 5 (blue) vs. Run 1 (red) vs. best score (white) as gain or loss of P@10 compared with median result.

the query expansion process varies according to the query. For some queries (7, 11 , 24, 34, 46), it improves the precision. For other queries (21, 25, 32, 33, 39, 47, 50), the expansion degrades the results. More interestingly, expansion does not affect P@10 for the 38 remaining queries. Among these queries, a few of them were not expanded (query 2 for example); it means that either no UMLS term was found in the query, or that no synonyms were found in the UMLS for the term. Yet, for most of the queries, some terms were actually added to the initial query, but they do not change the 10 first results.

Similarly, Figure 4 and 5 respectively present the same query by query analysis for Run 6 and 7. Compared to Run 5, these runs rely on more complex processing of the queries. Unfortunately, they perform slightly worse than Run 5. It is difficult to find the precise reasons of this loss of precision. It appears that this is due to only a subset of queries that perform worse than Run 5. In most of the cases, this is due to terms used as expansion in addition to those already used in Run 5. These terms mislead the IR process; for example, for query 48, the term *white blood cell* was expanded with synonyms of *white* (*caucasoid caucasian caucasians occidental*) in Run 6 and 7 but not in Run 5.
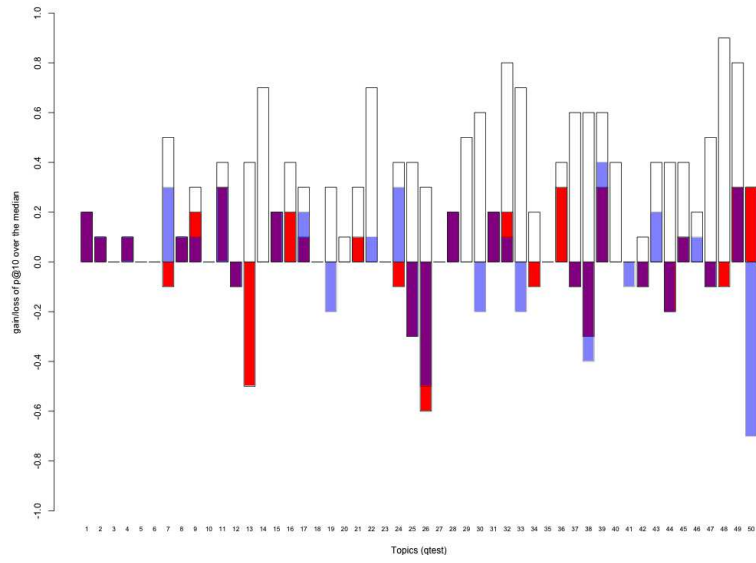
Fig. 4. Query expansion effect of Run 6 (blue) vs. Run 1 (red) vs. best score (white) as gain or loss of P@10 compared with median result.
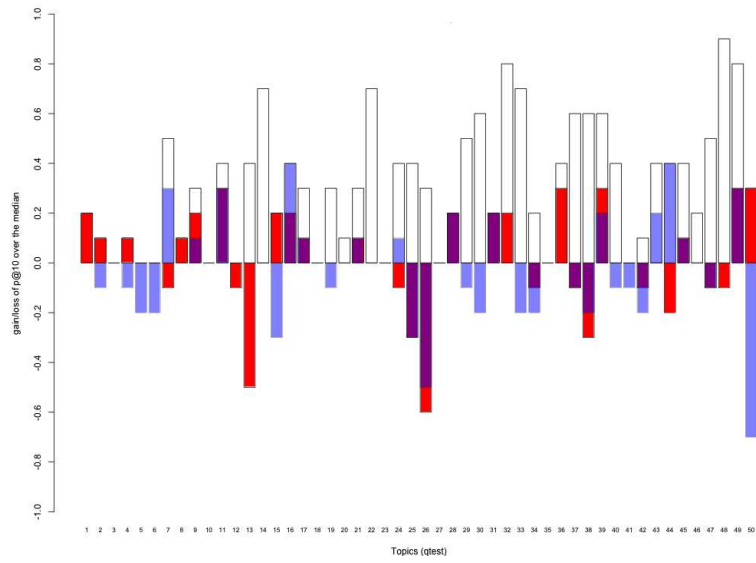

Fig. 5. Query expansion effect of Run 7 (blue) vs. Run 1 (red) vs. best score (white) as gain or loss of P@10 compared with median result.

6 Conclusive remarks and foreseen work

For our first participation, we have based our work on a state-of-the-art IR system and simple techniques to incorporate knowledge in the biomedical IR process. With the help of the 2013 dataset, some important parameters of the IR systems have been set.

The overall results of our submitted runs are good, compared with other IR evaluation campaigns, with P@10 as high as 0.65. Yet, our 3 strategies to incorporate external knowledge have yielded disappointing results. Indeed, the global benefits of the three query expansion strategies are limited, even though it appears as very interesting for particular queries. These mixed results are similar to existing studies on query expansion for general language [19]. Nonetheless, we plan to push further our investigation on how to exploit the biomedical terminologies in IR tasks. A detailed analysis of the results, when the relevance judgment will be released, may lead to better ways to choose which term to consider in the queries, and which synonyms of this term to add to the query. The incorporation of the terminological knowledge during the indexing step is also a promising avenue but raises computational issues.

References

1. Aubin, S., Hamon, T.: Improving term extraction with terminological resources. In: FinTAL 2006. pp. 380–387. No. 4139 in LNAI, Springer (2006)
2. de Boer, M.J., Versteegen, G.J., van Wijhe, M.: Patients' use of the internet for pain-related medical information. Patient Education and Counseling 68(1), 86–97 (2007)
3. Diaz, J.A., Griffith, R.A., Ng, J.J., Reinert, S.E., Friedmann, P.D., , Moulton, A.W.: Patients' use of the internet for medical information. J Gen Intern Med 17(3), 180–185 (2002)
4. Elhadad, N., McKeown, K.: Towards generating patient specific summaries of medical articles. In: Proc of NAACL WS on automatic summarization. pp. 31–39. Pittsburg (2001)
5. Elhadad, N., Sutaria, K.: Mining a lexicon of technical terms and lay equivalents. In: BioNLP. pp. 49–56 (2007)
6. Elhadad, N.: Comprehending technical texts: predicting and defining unfamiliar terms. In: AMIA. pp. 239–243 (2006)
7. Eysenbach, G., Kohler, C.: What is the prevalence of health-related searches on the world wide web? qualitative and quantitative analysis of search engine queries on the internet. In: AMIA Annu Symp Proc. pp. 225–229 (2003)
8. Goeuriot, L., Kelly, L., Li, W., Palotti, J., Pecina, P., Zuccon, G., Hanbury, A., Jones, G., Mueller, H.: Share/clef ehealth evaluation lab 2014, task 3: User-centred health information retrieval. In: Proceedings of CLEF 2014 (2014)
9. Hamon, T., Nazarenko, A.: Le développement d'une plate-forme pour l'annotation spécialisée de documents web: retour d'expérience. TAL 49(2), 127–154 (2008)
10. Jacquemin, C.: A symbolic and surgical acquisition of terms through variation. In: Wermter, S., Riloff, E., Scheler, G. (eds.) Connectionist, Statistical and Symbolic Approaches to Learning for Natural Language Processing. pp. 425–438. Springer (1996)

11. Jucks, R., Bromme, R.: Choice of words in doctor-patient communication: an analysis of health-related internet sites. Health Commun 21(3), 267–77 (2007)
12. Kelly, L., Goeuriot, L., Suominen, H., Schrek, T., Leroy, G., Mowery, D.L., Velupillai, S., Chapman, W.W., Martinez, D., Zuccon, G., Palotti, J.: Overview of the share/clef ehealth evaluation lab 2014. In: Proceedings of CLEF 2014. Lecture Notes in Computer Science (LNCS), Springer (2014)
13. Kleiber, G., Tamba, I.: L'hyperonymie revisitée : inclusion et hiérarchie. Langages 98, 7–32 (juin 1990)
14. Metzler, D., Croft, W.B.: A markov random field model for term dependencies. In: Proceedings of the 28th annual international ACM SIGIR conference on Research and development in information retrieval. pp. 472–479. ACM (2005)
15. NLM: UMLS Knowledge Sources Manual. National Library of Medicine, Bethesda, Maryland (2008), www.nlm.nih.gov/research/umls/
16. Pletneva, N., Vargas, A., Boyer, C.: Requirements for the general public health search. Tech. rep., KHRESMOI project (2011), d8.1.1
17. Schmid, H.: Probabilistic part-of-speech tagging using decision trees. In: Proceedings of the International Conference on New Methods in Language Processing. pp. 44–49. Manchester, UK (1994)
18. Strohman, T., Metzler, D., Turtle, H., Croft, W.B.: Indri: a language-model based search engine for complex queries. Tech. rep., in Proceedings of the International Conference on Intelligent Analysis (2005)
19. Voorhees, E.M.: Query expansion using lexical-semantic relations. In: Proceedings of the 17th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval. pp. 61–69. SIGIR '94, Springer-Verlag New York, Inc., New York, NY, USA (1994), http://dl.acm.org/citation.cfm?id=188490.188508
20. Zeng, Q.T., Kim, E., Crowell, J., Tse, T.: A text corpora-based estimation of the familiarity of health terminology. In: ISBMDA 2006. pp. 184–92 (2005)
21. Zeng, Q.T., Tse, T., Crowell, J., Divita, G., Roth, L., Browne, A.C.: Identifying consumer-friendly display (CFD) names for health concepts. In: AMIA 2006. pp. 859–63 (2005)
22. Zeng, Q.T., Tse, T., Divita, G., Keselman, A., Crowell, J., Browne, A.C.: Exploring lexical forms: first-generation consumer health vocabularies. In: AMIA 2006. pp. 1155–1155 (2006)
23. Zeng, Q., Tse, T.: Exploring and developing consumer health vocabularies. JAMIA 13, 24–29 (2006)