

A Relation Extraction Approach for Clinical Decision Support

Maristella Agosti
maristella.agosti@unipd.it
Stefano Marchesin
stefano.marchesin@unipd.it

Giorgio Maria Di Nunzio
giorgiomaria.dinunzio@unipd.it
Gianmaria Silvello
gianmaria.silvello@unipd.it

Department of Information Engineering
University of Padua, Italy
Via Giovanni Gradenigo, 6/b

Abstract

In this paper, we investigate how semantic relations between concepts extracted from medical documents can be employed to improve the retrieval of medical literature. Semantic relations explicitly represent relatedness between concepts and carry high informative power that can be leveraged to improve the effectiveness of retrieval functionalities of clinical decision support systems. We present preliminary results and show how relations are able to provide a sizable increase of the precision for several topics, albeit having no impact on others. We then discuss some future directions to minimize the impact of negative results while maximizing the impact of good results.

1 Motivation

The volume of medical literature published every year keeps growing at a very fast pace. The time required by clinicians to retrieve relevant information from such an amount of literature using standard systems is often prohibitive. Therefore, there has been a strong interest in Clinical Decision Support (CDS) systems [Ber07] designed to produce effective and timely information that can help clinicians in the decision making process for patient care. Within this context, we focus on case-based retrieval — i.e. given a medical case of interest, the CDS system should retrieve highly related medical

literature from a large collection of medical publications. Due to severe time constraints, clinicians must take fast decisions without having the possibility to thoroughly read the literature; for this reason, case-based retrieval favors precision over recall [BDS⁺04].

A key characteristic of the medical literature is the large use of synonyms and context-specific expressions. To address this term heterogeneity, Knowledge Bases (KBs) have often been exploited by Information Retrieval (IR) systems. The current availability of medical KBs offers us the opportunity to develop techniques that better capture the semantics of medical documents, leading to the following research question:

How can we employ the rich semantic information within medical case reports and related literature to boost retrieval performances and ease the clinical decision process?

Semantic relations are a key aspect within the semantics of a document. They have been mainly used to find relevant concepts to expand a user query, but not as semantic elements to be indexed and retrieved. We hypothesize semantic relations can provide a higher semantic representation of medical cases and literature.

In this work, we present an initial study on the effectiveness of the use of semantic relations for the retrieval of medical literature. We define an approach comprising two methods: a *rule-based* method and a *learning* method. In the rule-based method, we assign a relation to a pair of concepts — contained within the same sentence — when it holds within a reference KB. In the learning method, we train a sentence-level relation extractor that is able to infer relation between a pair of concepts given the sentence's context.

We evaluated our approach by using the publicly shared OHSUMED collection [HBLH94]. OHSUMED

provides rather short queries which represent a hard task for our approach, since limited information — e.g. concepts and relations — can be extracted from them. Testing with OHSUMED allow us to assess the potential and limitations of the approach. The remainder of the paper is organized as follows: Section 2 presents the background and related work, Section 3 describes the proposed approach, Section 4 presents experiments and results and Section 5 draws some conclusions and outlines future work.

2 Related Work

Concept-based IR aims at making use of external sources (like thesauri and ontologies) to provide additional knowledge and context that may not be explicit in a document collection and users’ queries. Concept-based methods can be categorized in two types: (i) methods that use concepts in both indexing and retrieval stages [EMG11], and (ii) methods that apply concept analysis in one specific stage, such as concept-based query expansion [GVDW06]. The approach we adopt extends the use of concepts to relations and uses them in all stages of retrieval. This is more challenging, but it allows for a finer semantic representation of documents and queries.

In the biomedical domain — where there are authoritative and curated ontologies — concept-based approaches demonstrate consistent improvements over classic keyword-based systems. In [KZN⁺12], ‘is-a’ relationships between concepts are used to weight documents containing concepts subsumed by the query’s concepts. [LMO13b] proposes a method to represent medical records and queries by focusing only on medical concepts essential for the information need of a medical search task. In [LMO13a], queries are expanded by inferring additional conceptual relationships from domain-specific resources as well as by extracting informative concepts from the top-ranked medical records.

The field of Biomedical Information Extraction (BioIE) is highly relevant for CDS. [LCJY16] reviews the recent advances in learning-based approaches for BioIE tasks. BioIE tasks comprise entity linking [ZHZ⁺15], event identification [APTK10] and relation extraction [USSD11, WF14]. Being targeted to CDS — i.e. voted to the extraction of key relations that can facilitate clinical decision making — our problem setup is fundamentally different from the conventional biomedical setups. Most of state-of-the-art biomedical relation extraction techniques are developed for specific relations, like protein-protein interactions, gene-disease interactions and so on — which cover only a fraction of the biomedical domain.

Regarding relations in IR, [VMdR17] study the

problem of finding human readable descriptions of a given relationship in a knowledge graph. [SRPD16] apply supervised relation extraction to documents that are relevant for an information need Q and study how many of the extracted relations are indeed relevant for Q . [KD17] explores current state of the art in unsupervised relation extraction (OpenIE) for the task of finding support passages to complement an entity ranking with human-readable explanations of how those retrieved entities are connected to the information need. Conversely, our approach applies supervised relation extraction to extract semantic relations that are used in all stages of retrieval. Hence, relations play a pivotal role in the actual retrieval of documents.

3 Methodology

We present a new approach that uses semantic relations for case-based retrieval. The methodology is composed of the information extraction step (Subsection 3.1) and the information retrieval step (Subsection 3.2).

3.1 Information Extraction

The information extraction step is divided into an entity linking component and a relation extraction component.

The entity linking component extracts entity mentions within the text and links them to a reference KB; this reduces the high number of synonyms, abbreviations and context specific expressions that are present in the medical literature. For entity linking we adopt MetaMap,¹ the most authoritative tool to detect medical entity mentions in free-text. MetaMap analyses biomedical free-text and identifies concepts belonging to the Unified Medical Language System (UMLS), associating each mention with a number of concepts from the UMLS Metathesaurus² — which comprises more than 3 million distinct concepts. Within UMLS, a substantial understanding of the medical domain is included, comprising medical concepts, relations, definitions and so on.

The relation extraction component detects semantic relation between pairs of concepts within a sentence. To be consistent with concepts extracted with MetaMap, we consider semantic relations from UMLS Metathesaurus as well. Furthermore, since our task requires a high coverage of the medical domain, considering UMLS Metathesaurus relations — which are coarse-grained relationships that span to a high number of concepts — allows us to increase the recall of extracted relations.

¹<https://metamap.nlm.nih.gov/>

²https://www.nlm.nih.gov/research/umls/knowledge_sources/metathesaurus/

We define two methods for the extraction of relations from documents and queries: a rule-based method and a learning method. **Rule-based:** a relation is assigned to a pair of concepts if it relates them within UMLS. We assume that a UMLS relation between two concepts always occurs, even when it is not explicitly mentioned in the sentence containing the two concepts.

Learning: we train a distantly supervised [MBSJ09] sentence-level Bidirectional Long Short-Term Memory (BiLSTM) neural network to detect if a relation exists between two concepts based on the context of the sentence. The network architecture is composed of an input (word embedding) layer of concatenated word features and positional features. Words are first converted into pre-trained word embeddings trained on 26 million abstracts and citations in PubMed — released by [PGM⁺13]. Then these word features are concatenated with two sets of positional features — to explicitly account for the pairs of words to which we expect to assign relations [ZLL⁺14]. We apply a max-pooling layer right after the bidirectional recurrent layer and before the output layer — in order to combine segment-level features that, although not very strong in representing the entire sentence, represent local patterns well [ZW15]. In this way, we try to overcome the tendency of recurrent connections to forget long-term information too quickly, leading the supervision at the end of the sentence to be hardly propagated to early steps in model training (due to gradient vanishing [BSF94]).

3.2 Information Retrieval

Documents are indexed by considering all terms as in the Bag-of-Words (BoW) representation. We extend the BoW representation to both concepts (BoC) and relations (BoR) by considering for the indexing all the extracted concepts and relations respectively. The ranking is obtained using Okapi BM25 ranking formula [RW99].

Since relations are extracted at sentence level, we also index passages — i.e. groups of consecutive sentences — by considering all the relations occurring within each group of sentences (passage-level BoR). Relevant passages should contain a higher number of relations related to the information need when compared to non relevant passages — being more similar in their semantic contents to the query. Therefore, documents that contain more relevant passages can be considered more relevant for the query.

We define a weighting scheme such that a document score is computed as the weighted sum of its passage scores, where scores are computed using BM25 as above. The passage-level weighting scheme is as

follows:

$$score(q, d) = \sum_{p \in d} \frac{|R_p \cap R_q|}{|R_q|} BM25(p, q) \quad (1)$$

where d is the document, q is the query, p is a passage belonging to document d , R_q is the set of relations extracted from query q and R_p is the set of relations extracted from passage p .

4 Experiments and Results

We employed the OHSUMED test collection which contains 348,566 references from the on-line medical information database MEDLINE, consisting of titles and/or abstracts from 270 medical journals over a five-year period (1987-1991). The available fields are: title, abstract, MeSH indexing terms, author, source, and publication type. There are 106 queries in the collection. Each query is composed of two sentences: title + description. Title is the brief summary of the medical case at hand, description is the information need required to answer a specific question for the case.

Experimental Setup: We performed two experiments: i) one using the rule-based method to extract relations out of documents and queries; ii) the other using the learning method to extract relations out of documents and queries. We compared the results obtained applying BM25 to the three representations (i.e. BoW, BoC and BoR) and we evaluated the results using the nDCG measure.

Results: i) The rule-based method was able to extract relations from a subset of 44 queries. Therefore, to investigate the effectiveness of relations, we restrict the experiments to this subset only — since the remaining queries lead to no results when considering relations. Of these 44 queries, only 39 have relations matching with some documents. Regarding the relations, we obtained the best results with the passage-level approach. We set the passage length to 2, in order to be compliant with queries' length. Documents' score was computed using the formula shown above (1). The nDCG results on these 39 queries are variable — ranging from 0 (18 cases) to 1 (5 cases), as can be seen in Figure 1.

Such a variance gives us some hints about the informative power of relations. When properly extracted, relations can be highly effective, indeed, we compared the average nDCG values of concepts and relations on only those topics where relations give a result different than 0 and we found a statistically significant average improvement of 20%. A *t-test* was performed to validate the improvement. Regarding the comparison between relations and terms, the behavior of relations is similar to the one of terms (baseline approach), and




Figure 1: nDCG values for topics containing relations. OHSU{56,60,29,46,58} returned NA values for the BoR representation. Queries are sorted in descending order first by BoR (relations) nDCG values, then by BoC (concepts) nDCG values.

there is no statistically significant difference between the two.

ii) The learning method was able to extract relations from a subset of 25 queries. Of these 25 queries, only 12 have relations matching with some documents. The results on these 12 queries are comparable to those presented for the rule-base method, with nDCG values ranging from 0 (in 7 cases) to 1 (in 1 case). The reason for this is two-fold: (a) the shortness of queries that limits the relations that can be extracted; and, (b) the highly different syntactic structure of queries if compared to the sentences within the medical abstracts leading to a mismatch between the query-relations and abstract-relations.

5 Discussion

In this work, we proposed and evaluated the effectiveness of semantic relations as basic constituents for a CDS system. We defined two methods for extracting relations from queries and documents: a rule-based method and a learning method. We found that relations — when pertinent to the initial information need — are highly valuable, outperforming concepts. The challenge lies in how to limit those cases where relations provide no relevant results for the information need. To this end, considering collections where queries present a long and narrative structure (e.g. TREC CDS tracks) might be a possible direction to balance such issue. Furthermore, defining more IR-oriented relation extraction algorithms that are capable of overcoming the high precision-low recall nature of state-of-the-art methods is a direction we will investigate.

Acknowledgements

Supported by the CDC-STARS project of the University of Padua.

References

- [APTK10] S. Ananiadou, S. Pyysalo, J. Tsujii, and D. B. Kell. Event extraction for systems biology by text mining the literature. *Trends in Biotechnology*, 28(7):381–390, 2010.
- [BDS⁺04] D. T. Burke, M. C. DeVito, J. C. Schneider, S. Julien, and A. L. Judelson. Reading habits of physical medicine and rehabilitation resident physicians. *American Journal of Physical Medicine & Rehabilitation*, 83(7):551–559, 2004.
- [Ber07] E. S. Berner. *Clinical decision support systems*, volume 233. Springer, 2007.
- [BSF94] Y. Bengio, P. Simard, and P. Frasconi. Learning long-term dependencies with gradient descent is difficult. *IEEE Transactions on Neural Networks*, 5(2):157–166, 1994.
- [EMG11] O. Egozi, S. Markovitch, and E. Gabrilovich. Concept-based information retrieval using explicit semantic analysis. *ACM Trans. Inf. Syst.*, 29(2):8:1–8:34, April 2011.
- [GVDW06] F. A. Grootjen and T. P. Van Der Weide. Conceptual query expansion. *Data & Knowledge Engineering*, 56(2):174–193, 2006.
- [HBLH94] W. Hersh, C. Buckley, T.J. Leone, and D. Hickam. Ohsumed: an interactive retrieval evaluation and new large test collection for research. In *SIGIR’94*, pages 192–201. Springer, 1994.
- [KD17] A. Kadry and L. Dietz. Open relation extraction for support passage retrieval: Merit and open issues. In *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 1149–1152. ACM, 2017.
- [KZN⁺12] B. Koopman, G. Zuccon, A. Nguyen, D. Vickers, L. Butt, and P. D. Bruza. Exploiting snomed ct concepts and relationships for clinical information retrieval: Australian e-health research centre and queensland university of technology at the trec 2012 medical track. In *The Twenty-First Text REtrieval Conference Proceedings (TREC 2012)[NIST Special Publication: SP 500-298]*, pages 1–8, 2012.
- [LCJY16] F. Liu, J. Chen, A. Jagannatha, and H. Yu. Learning for biomedical information extraction: methodological review of recent advances. *arXiv preprint arXiv:1606.07993*, 2016.
- [LMO13a] N. Limsopatham, C. Macdonald, and I. Ounis. Inferring conceptual relationships to improve medical records search. In *Proceedings of the 10th Conference on Open Research Areas in Information Retrieval*, pages 1–8, 2013.
- [LMO13b] N. Limsopatham, C. Macdonald, and I. Ounis. A task-specific query and document representation for medical records search. In *European Conference on Information Retrieval*, pages 747–751. Springer, 2013.
- [MBSJ09] M. Mintz, S. Bills, R. Snow, and D. Jurafsky. Distant supervision for relation extraction without labeled data. In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP: Volume 2-Volume 2*, pages 1003–1011. Association for Computational Linguistics, 2009.
- [PGM⁺13] S. Pyysalo, F. Ginter, H. Moen, T. Salakoski, and S. Ananiadou. Distributional semantics resources for biomedical text processing. In *Proceedings of the 5th International Symposium on Languages in Biology and Medicine, Tokyo, Japan*, pages 39–43, 2013.
- [RW99] S. E. Robertson and S. Walker. Okapi/keenbow at trec-8. In *TREC*, volume 8, pages 151–162. Citeseer, 1999.
- [SRPD16] M. Schuhmacher, B. Roth, S. P. Ponzetto, and L. Dietz. Finding relevant relations in relevant documents. In *European Conference on Information Retrieval*, pages 654–660. Springer, 2016.
- [USSD11] Özlem Uzuner, Brett R South, Shuying Shen, and Scott L DuVall. 2010 i2b2/va challenge on concepts, assertions, and relations in clinical text. *Journal of the American Medical Informatics Association*, 18(5):552–556, 2011.
- [VMdR17] N. Voskarides, E. Meij, and M. de Rijke. Generating descriptions of entity

relationships. In *European Conference on Information Retrieval*, pages 317–330. Springer, 2017.

- [WF14] C. Wang and J. Fan. Medical relation extraction with manifold models. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 828–838, 2014.
- [ZHZ⁺15] J. Zheng, D. Howsmon, B. Zhang, J. Hahn, D. McGuinness, J. Hendler, and H. Ji. Entity linking for biomedical literature. *BMC Medical Informatics and Decision Making*, 15(1):S4, 2015.
- [ZLL⁺14] D. Zeng, K. Liu, S. Lai, G. Zhou, and J. Zhao. Relation classification via convolutional deep neural network. In *Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: Technical Papers*, pages 2335–2344, 2014.
- [ZW15] D. Zhang and D. Wang. Relation classification via recurrent neural network. *arXiv preprint arXiv:1508.01006*, 2015.