



# DSR: A Collection for the Evaluation of Graded Disease-Symptom Relations

Markus Zlabinger<sup>1</sup>(✉), Sebastian Hofstätter<sup>1</sup>, Navid Rekabsaz<sup>2</sup>,  
and Allan Hanbury<sup>1</sup>

<sup>1</sup> TU Wien, Vienna, Austria

{markus.zlabinger,sebastian.hofstatter,allan.hanbury}@tuwien.ac.at

<sup>2</sup> Johannes Kepler University, Linz, Austria  
navid.rekabsaz@jku.at

**Abstract.** The effective extraction of ranked disease-symptom relationships is a critical component in various medical tasks, including computer-assisted medical diagnosis or the discovery of unexpected associations between diseases. While existing disease-symptom relationship extraction methods are used as the foundation in the various medical tasks, no collection is available to systematically evaluate the performance of such methods. In this paper, we introduce the *Disease-Symptom Relation Collection* (DSR-collection), created by five physicians as expert annotators. We provide graded symptom judgments for diseases by differentiating between *relevant symptoms* and *primary symptoms*. Further, we provide several strong baselines, based on the methods used in previous studies. The first method is based on word embeddings, and the second on co-occurrences of MeSH-keywords of medical articles. For the co-occurrence method, we propose an adaption in which not only keywords are considered, but also the full text of medical articles. The evaluation on the DSR-collection shows the effectiveness of the proposed adaption in terms of nDCG, precision, and recall.

**Keywords:** Disease-symptom relationship · Medical expert data annotation · Disease-symptom information extraction

## 1 Introduction

Disease-symptom knowledge bases are the foundation for many medical tasks – including medical diagnosis [9] or the discovery of unexpected associations between diseases [12, 14]. Most knowledge bases only capture a binary relationship between diseases and symptoms, neglecting the degree of the importance between a symptoms and a disease. For example, abdominal pain and nausea are both symptoms of an appendicitis, but while abdominal pain is a key differentiating factor, nausea does only little to distinguish appendicitis from other diseases of the digestive system. While several disease-symptom extraction methods have been proposed that retrieve a ranked list of symptoms for a disease [7, 10, 13, 14],

no collection is available to systematically evaluate the performance of such methods [11]. While these methods are extensively used in downstream tasks, e.g., to increase the accuracy of computer-assisted medical diagnosis [9], their effectiveness for disease-symptom extraction remains unclear.

In this paper, we introduce the *Disease-Symptom Relation Collection* (DSR-collection) for the evaluation of graded disease-symptom relations. The collection is annotated by five physicians and contains 235 symptoms for 20 diseases. We label the symptoms using graded judgments [5], where we differentiate between: *relevant symptoms* (graded as 1) and *primary symptoms* (graded as 2). Primary symptoms—also called cardinal symptoms—are the leading symptoms that guide physicians in the process of disease diagnosis. The graded judgments allow us for the first time to measure the importance of different symptoms with grade-based metrics, such as nDCG [4].

As baselines, we implement two methods from previous studies to compute graded disease-symptom relations: In the first method [10], the relation is the cosine similarity between the word vectors of a disease and a symptom, taken from a word embedding model. In the second method [14], the relation between a disease and symptom is calculated based on their co-occurrence in the MeSH-keywords<sup>1</sup> of medical articles. We describe limitations of the keyword-based method [14] and propose an adaptation in which we calculate the relations not only on keywords of medical articles, but also on the full text and the title.

We evaluate the baselines on the DSR-collection to compare their effectiveness in the extraction of graded disease-symptom relations. As evaluation metrics, we consider precision, recall, and nDCG. For all three metrics, our proposed adapted version of the keyword-based method outperforms the other methods, providing a strong baseline for the DSR-collection.

The contributions of this paper are the following:

- We introduce the DSR-collection for the evaluation of graded disease-symptom relations. We make the collection freely available to the research community.<sup>2</sup>
- We compare various baselines on the DSR-collection to give insights on their effectiveness in the extraction of disease-symptom relations.

## 2 Disease-Symptom Relation Collection

In this section, we describe the new *Disease-Symptom Relation Collection* (DSR-collection) for the evaluation of disease-symptom relations. We create the collection in two steps: In the first step, relevant disease-symptom pairs (e.g. *appendicitis-nausea*) are collected by two physicians. They collect the pairs in a collaborative effort from high-quality sources, including medical textbooks and an online information service<sup>3</sup> that is curated by medical experts.

<sup>1</sup> MeSH-keywords are meta-data that indicates the core topics of an medical article.

<sup>2</sup> Contact this paper’s first author to gain access.

<sup>3</sup> The website *netdokter.at* which is certificated by the Health on the Net Foundation.

In the second step, the primary symptoms of the collected disease-symptom pairs are annotated. The annotation of primary symptoms is conducted to incorporate a graded relevance information into the collection. For the annotation procedure, we develop guidelines that briefly describe the task and an online annotation tool. Then, the annotation of primary symptoms is conducted by three physicians. The final label is obtained by a majority voting. Based on the labels obtained from the majority voting, we assign the relevance score 2 to *primary symptoms* and 1 to the other symptoms, which we call *relevant symptoms*.

In total, the DSR-collection contains relevant symptoms and primary symptoms for 20 diseases. We give an overview of the collection in Table 1. For the 20 diseases, the collection contains a total of 235 symptoms, of which 55 are labeled as primary symptom (about 25%). The top-3 most occurring symptoms are: *fatigue* which appears for 15 of the 20 diseases, *fever* which appears for 10, and *coughing* which appears for 7. Notice that the diseases are selected from different medical disciplines: mental (e.g. *Depression*), dental (e.g. *Periodontitis*), digestive (e.g. *Appendicitis*), and respiration (e.g. *Asthma*).

**Table 1.** Overview of the DSR-collection. For each disease, we display the number of relevant symptoms (#S), the number of primary symptoms (#P), and the Fleiss’ inter-annotator agreement ( $\kappa$ ).

Disease	#S	#P	$\kappa$	Disease	#S	#P	$\kappa$
Anorexia Nervosa	7	2	1.00	Influenza	11	2	0.57
Appendicitis	7	2	1.00	Measles	9	4	0.38
Asthma	9	4	0.76	Mental Depression	13	3	0.21
Bronchitis	9	1	0.71	Migraine Disorders	12	4	0.37
Cholecystitis	12	1	0.55	Myocardial Infarction	11	4	0.44
COPD	7	3	0.83	Periodontitis	3	4	0.46
Diabetes Mellitus	11	3	0.72	Pulmonary Embolism	13	2	0.83
Epididymitis	8	2	0.67	Sleep Apnea Syndromes	13	2	0.31
Erysipelas	7	3	0.69	Tonsillitis	7	4	0.63
GERD	8	2	0.76	Trigeminal Neuralgia	3	3	0.28

We calculate the inter-annotator agreement using Fleiss’ kappa [2], a statistical measure to compute the agreement for three or more annotators. For the annotation of the primary symptoms, we measure a kappa value of  $\kappa = 0.61$ , which indicates a substantial agreement between the three annotators [6]. Individual  $\kappa$ -values per disease are reported in Table 1. By analyzing the disagreements, we found that the annotators labeled primary symptoms with varying frequencies: The first annotator annotated on average 2.1 primary symptoms per disease, the second 2.8, and the third 3.8.

*Vocabulary Compatibility:* We map each disease and symptom of the collection to the Unified Medical Language System (UMLS) vocabulary. The UMLS is a

compendium of over 100 vocabularies (e.g. ICD-10, MeSH, SNOMED-CT) that are cross-linked with each other. This makes the collection compatible with the UMLS vocabulary and also with each of the over 100 cross-linked vocabularies.

Although the different vocabularies are compatible with the collection, a fair comparison of methods is only possible when the methods utilize the same vocabulary since the vocabulary impacts the evaluation outcome. For instance, the symptom *loss of appetite* is categorized as a symptom in MeSH; whereas, in the cross-linked UMLS vocabulary, it is categorized as a disease. Therefore, the symptom *loss of appetite* can be identified when using the MeSH vocabulary, but it cannot be identified when using the UMLS vocabulary.

*Evaluation:* We consider following evaluation metrics for the collection: Recall@k, Precision@k, and nDCG@k at the cutoff  $k = 5$  and  $k = 10$ . *Recall* measures how many of the relevant symptoms are retrieved, *Precision* measures how many of the retrieved symptoms are relevant, and finally, *nDCG* is a standard metric to evaluate graded relevance [5].

## 3 Disease-Symptom Extraction Methods

### 3.1 Related Methods

In this section, we discuss disease-symptom extraction methods used in previous studies. A commonly used resource for the extraction of disease-symptom relations are the articles of the PubMed database. PubMed contains more than 30 million biomedical articles, including the abstract, title, and various metadata. Previous work [3, 7] uses the abstracts of the PubMed articles together with rule-based approaches. In particular, Hassan et al. [3] derive patterns of disease-symptom relations from dependency graphs, followed by the automatic selection of the best patterns based on proposed selection criteria. Martin et al. [7] generate extraction rules automatically, which are then inspected for their viability by medical experts. Xia et al. [13] design special queries that include the name and synonyms of each disease and symptom. They use these queries to return the relevant articles, and use the number of retrieved results to perform a ranking via Pointwise Mutual Information (PMI).

The mentioned studies use resources that are not publicly available, i.e., rules in [3, 7] and special queries in [13]. To enable reproducibility in future studies, we define our baselines based on the methods that only utilize publicly available resources, described in the next section.

### 3.2 Baseline Methods

Here, we first describe two recently proposed methods [10, 14] for the extraction of disease-symptom relations as our baselines. Afterwards, we describe limitations of the method described in [14] and propose an adapted version in which the limitations are addressed. We apply the methods on the open-access subset of

the PubMed Central (PMC) database, containing 1,542,847 medical articles. To have a common representation for diseases/symptoms across methods (including an unique name and identifier), we consider the 382 symptoms and 4,787 diseases from the Medical Subject Headings (MeSH) vocabulary [14]. Given the set of diseases ( $X$ ) and symptoms ( $S$ ), each method aims to compute a relation scoring function  $\lambda(x, s) \in \mathbb{R}$  between a disease  $x \in X$  and a symptom  $s \in S$ . In the following, we explain each method in detail.

**EMBEDDING:** Proposed by Shah et al. [10], the method is based on the cosine similarity of the vector representations of a disease and a symptom. We first apply MetaMap [1], a tool for the identification of medical concepts within a given text, to the full text of all PMC articles to substitute the identified diseases/symptoms by their unique names. Then, we train a word2vec model [8] with 300 dimensions and a window size of 15, following the parameter setting in [10]. Using the word embedding, the disease-symptom relation is defined as  $\lambda(x, s) = \cos(\mathbf{e}_x, \mathbf{e}_s)$ , where  $\mathbf{e}$  refers to the vector representation of a word.

**COOCCUR:** This method, proposed by Zhou et al. [14], calculates the relation of a disease and a symptom, by measuring the degree of their co-occurrences in the MeSH-keywords of medical articles. The raw co-occurrence of the disease  $x$  and symptom  $s$ , is denoted by  $\text{co}(x, s)$ . The raw co-occurrence does not consider the overall appearance of each symptom across diseases. For instance, symptoms like *pain* or *obesity* tend to co-occur with many diseases, and are therefore less informative. Hence, the raw co-occurrence is normalized by an *Inverse Symptom Frequency (ISF)* measure, defined as  $\text{ISF}(s) = \frac{|X|}{n_s}$ , where  $|X|$  is the total number of diseases and  $n_s$  is the number of diseases that co-occur with  $s$  at least in one of the articles. Finally, the disease-symptom relation is defined as  $\lambda(x, s) = \text{co}(x, s) \times \text{ISF}(s)$ . We compute three variants of the CoOccur method:

- **KWD:** The disease-symptom relations are computed using the MeSH-keywords of the  $\approx 1.5$  million PMC articles.
- **KWDLARGE:** While KWD uses the 1.5 million PMC articles, Zhou et al. [14] apply the exact same method on the  $\approx 30$  million articles of the PubMed database. While they did not evaluate the effectiveness of their disease-symptom relation extraction method, they published their relation scores which we will evaluate in this paper.
- **FULLTEXT:** Applying the CoOccur method only on MeSH-keywords has two disadvantages: First, keywords are not available for all articles (e.g. only 30% of the  $\approx 1.5$  million PMC articles have keywords) and second, usually only the core topics of an article occur as keywords. We address these limitations by proposing an adaption of the CoOccur method, in which we use the full text, the title, and the keywords of the  $\approx 1.5$  million PMC articles. Specifically, we adapt the computation of the co-occurrence  $\text{co}(x, s)$ , as follows: We first retrieve a set of relevant articles to a disease  $x$ , where an article is relevant if the disease exists in either the keyword, or the title section of the article. Given these relevant articles and a symptom  $s$ , we compute the adapted co-occurrence  $\text{co}(x, s)$ , which is the number of relevant articles in that the

symptom occurs in the full text. The identification of the diseases in the title and symptoms in the full text is done using the MetaMap tool [1].

## 4 Evaluation Results and Discussion

We now compare the disease-symptom extraction baselines on the proposed DSR-collection. The results for various evaluation metrics are shown in Table 2. The FULLTEXT-variant of the CoOCCUR method outperforms the other baselines on all evaluation metrics. This demonstrates the high effectiveness of our proposed adaption to the CoOCCUR method.

Further, we see a clear advantage of the CoOCCUR-method with MeSH-keywords from  $\approx 30$  million PubMed articles as the resource (KWDLARGE) – in comparison to the same method with keywords from approximately 1.5 million PMC articles (KWD). This highlights the importance of the number of input samples to the method.

**Table 2.** Comparison of the disease-symptom extraction methods using our proposed DSR-collection. We show significant improvements with: *a* refers to EMBEDDING, *b* to KWD, and *c* to KWDLARGE (two-sided, paired t-test:  $p < 0.01$ ).

Method	nDCG@5	P@5	R@5	nDCG@10	P@10	R@10
EMBEDDING	0.20	0.18	0.08	0.19	0.15	0.13
CoOCCUR-KWD	0.27	0.22	0.10	0.22	0.14	0.12
CoOCCUR-KWDLARGE	0.32 <sup>a</sup>	0.27	0.12	0.28 <sup>ab</sup>	0.19	0.17
CoOCCUR-FULLTEXT	<b>0.41<sup>abc</sup></b>	<b>0.39<sup>abc</sup></b>	<b>0.17<sup>abc</sup></b>	<b>0.36<sup>abc</sup></b>	<b>0.28<sup>abc</sup></b>	<b>0.25<sup>ab</sup></b>

*Error Analysis:* A common error source is a result of the fine granularity of the symptoms in the medical vocabularies. For example, the utilized MeSH vocabulary contains the symptoms *abdominal pain* and *abdomen, acute*<sup>4</sup>. Both symptoms can be found in the top ranks of the evaluated methods for the disease appendicitis (see Table 3). However, since the corpus is not labeled on such a fine-grained level, the symptom *abdomen, acute* is counted as a false positive.

**Table 3.** Top-4 extracted symptoms of each method for the disease *appendicitis*. The retrieved **relevant symptoms** and **primary symptoms** are highlighted.

EMBEDDING	KWD	KWDLARGE	FULLTEXT
Abdomen, Acute	<b>Abdominal Pain</b>	Abdomen, Acute	Abdomen, Acute
<b>Abdominal Pain</b>	Abdomen, Acute	<b>Abdominal Pain</b>	<b>Abdominal Pain</b>
Fever of Unknown Origin	Obesity	<b>Pelvic Pain</b>	<b>Vomiting</b>
Renal Colic	Thinness	Pain, Postoperative	<b>Nausea</b>

<sup>4</sup> Symptom for acute abdominal pain.

Another error source is a result of the bias in medical articles towards specific disease-symptom relationships. For instance, between the symptom *obesity* and *periodontitis*<sup>5</sup> a special relationship exists, which is the topic of various publications. Despite obesity not being a characteristic symptom of a periodontitis, all methods return the symptom in the top-3 ranks. A promising research direction is the selective extraction of symptoms from biomedical literature by also considering the context (e.g. in a sentence) in that a disease/symptom appears.

## 5 Conclusion

We introduced the *Disease-Symptom Relation Collection* (DSR-collection) for the evaluation of graded disease-symptom relations. We provided baseline results for two recent methods, one based on word embeddings and the second on the co-occurrence of MeSH-keywords of medical articles. We proposed an adaption to the co-occurrence method to make it applicable to the full text of medical articles and showed significant improvement of effectiveness over the other methods.

## References

1. Aronson, A.R.: Effective mapping of biomedical text to the UMLS Metathesaurus: the MetaMap program. In: Proceedings of the American Medical Informatics Association Symposium, pp. 17–21 (2001)
2. Fleiss, J.L.: Measuring nominal scale agreement among many raters. *Psychol. Bull.* **76**(5), 378 (1971)
3. Hassan, M., Makkaoui, O., Coulet, A., Toussain, Y.: Extracting disease-symptom relationships by learning syntactic patterns from dependency graphs. In: Proceedings of BioNLP 2015, pp. 71–80. Association for Computational Linguistics, Beijing (2015)
4. Järvelin, K., Kekäläinen, J.: Cumulated gain-based evaluation of IR techniques. *ACM Trans. Inf. Syst. (TOIS)* **20**(4), 422–446 (2002)
5. Kekäläinen, J.: Binary and graded relevance in IR evaluations - comparison of the effects on ranking of IR systems. *Inf. Process. Manag.* **41**(5), 1019–1033 (2005)
6. Landis, J.R., Koch, G.G.: The measurement of observer agreement for categorical data. *Biometrics* **33**, 159–174 (1977)
7. Martin, L., Battistelli, D., Charnois, T.: Symptom extraction issue. In: Proceedings of BioNLP 2014, pp. 107–111 (2014). Association for Computational Linguistics, Baltimore (2014)
8. Mikolov, T., Sutskever, I., Chen, K., Corrado, G.S., Dean, J.: Distributed representations of words and phrases and their compositionality. In: Advances in Neural Information Processing Systems, pp. 3111–3119 (2013)
9. Ni, J., Fei, H., Fan, W., Zhang, X.: Automated medical diagnosis by ranking clusters across the symptom-disease network. In: 2017 IEEE International Conference on Data Mining (ICDM), pp. 1009–1014, November 2017
10. Shah, S., Luo, X., Kanakasabai, S., Tuason, R., Klopper, G.: Neural networks for mining the associations between diseases and symptoms in clinical notes. *Health Inf. Sci. Syst.* **7**(1), 1–9 (2018). <https://doi.org/10.1007/s13755-018-0062-0>

<sup>5</sup> A dental disease where the gum that surrounds the teeth retreats.

11. Shen, Y., Li, Y., Zheng, H.T., Tang, B., Yang, M.: Enhancing ontology-driven diagnostic reasoning with a symptom-dependency-aware Naïve Bayes classifier. *BMC Bioinform.* **20**(1), 330 (2019)
12. del Valle, E.P.G., García, G.L., Santamaría, L.P., Zanin, M., Ruiz, E.M., González, A.R.: Evaluating Wikipedia as a source of information for disease understanding. In: 2018 IEEE 31st International Symposium on Computer-Based Medical Systems (CBMS), pp. 399–404, June 2018
13. Xia, E., Sun, W., Mei, J., Xu, E., Wang, K., Qin, Y.: Mining disease-symptom relation from massive biomedical literature and its application in severe disease diagnosis. In: 45 - AMIA 2018 Annual Symposium, pp. 1118–1126 (2018)
14. Zhou, X., Menche, J., Barabási, A.L., Sharma, A.: Human symptoms-disease network. *Nat. Commun.* **5**(1), 1–10 (2014)