A Semantically Enriched Dataset based on Biomedical NER for the COVID19 Open Research Dataset Challenge

Hermann Kroll
Institute for Information Systems
TU Braunschweig
Braunschweig, Germany
kroll@ifis.cs.tu-bs.de

Jan Pirklbauer
Institute for Information Systems
TU Braunschweig
Braunschweig, Germany
j.pirklbauer@tu-bs.de

Johannes Ruthmann
Institute for Information Systems
TU Braunschweig
Braunschweig, Germany
j.ruthmann@tu-bs.de

Wolf-Tilo Balke
Institute for Information Systems
TU Braunschweig
Braunschweig, Germany
balke@ifis.cs.tu-bs.de

ABSTRACT

Research into COVID-19 is a big challenge and highly relevant at the moment. New tools are required to assist medical experts in their research with relevant and valuable information. The COVID-19 Open Research Dataset Challenge (CORD-19) is a "call to action" for computer scientists to develop these innovative tools. Many of these applications are empowered by entity information, i. e. knowing which entities are used within a sentence. For this paper, we have developed a pipeline upon the latest Named Entity Recognition tools for Chemicals, Diseases, Genes and Species. We apply our pipeline to the COVID-19 research challenge and share the resulting entity mentions with the community.

KEYWORDS

Named Entity Recognition, COVID19 Research Challenge

1 INTRODUCTION

PubMed, the most extensive library for biomedical research, contains nearly 30 million publications. The Allen Institute for AI selects nearly 57,000 documents as relevant for COVID19 research (V9), and around 47,000 full texts are included within this selection. Accessing such an extensive document collection and finding relevant information is a hard task for medical researchers. Especially in times, when results are published within a few days, keeping an overview of the latest research can be exhausting. Novel tools are urgently needed to assist medical researchers in their workflows: novel search engines find relevant information precisely, and new access paths like summarization techniques offer new opportunities to engage the flood of information. These tools are typically empowered by utilizing additional side information like knowledge graphs [1].

Knowledge graphs are structured storages providing fact-style knowledge about entities, e. g. *Simvastatin is used in treatment of hypercholesterolemia*. In the biomedical domain, entities of interest are mainly *Chemicals*, *Diseases*, *Genes* and *Species*. The central problem of utilizing structured information for text retrieval is to detect, which entities are mentioned in the text. This problem is engaged by applying a Named Entity Recognition (NER), i. e. detecting important entities of in arbitrary texts. NER tools like Spotlight (DBpedia)

and WAT (Wikidata) are developed to recognize a variety of different entities in several domains [5, 6]. Unfortunately, the biomedical domain contains a variety of different entities. Dictionary-based recognition tools might fail here because the exact entity mention within a sentence depends on the context. Hence, homonyms must be resolved, e. g. the gene name *CYP3A4* has different ids depending if the sentence talks about mouses or humans. Yet, Named Entity Recognition tools suitable for the biomedical domain have been designed and built by experts already.

In this paper, we utilize two biomedical NER tools, namely TaggerOne [4] and GNormPlus [8], and build a pipeline to annotate arbitrary biomedical texts. Finally, we apply our pipeline to the COVID19 dataset. The detected entity mentions are published in our GitHub[1] repository for free reuse. The code will be published under the MIT license[2]. The data is published for free reuse under the Creative Commons Attribution 4.0 International license (CC BY 4.0)[3]. We hope that this additional entity information can serve as a solid and high-quality platform for novel tools and thus enable more research about COVID19.

2 A BIOMEDICAL NER PIPELINE

First we will introduce a pipeline for biomedical Named Entity Recognition in arbitrary texts. The task of a Named Entity Recognition is to detect entity mentions in texts. An **entity** represents a thing of interest in a specific domain, e. g. Chemicals and Diseases are of interest in the biomedical domain. Further, an entity consists of a unique id and an entity type, e. g. (*Simvastatin*, *Chemical*) is a valid entity. Entities are described by a predefined vocabulary, which is typically build by experts. Entities might be mentioned within a written text. Therefore, we understand text as a sequence of sentences and sentences as a sequence of tokens (single words). A sequence of tokens within an sentence might represent an entity. We call this representation an **entity mention**. Hence, entity mentions consist of an entity and a sequence of corresponding tokens within a sentence.

---

[1]https://github.com/HermannKroll/CORD19BiomedicalNERDataset
[2]https://opensource.org/licenses/MIT
[3]https://creativecommons.org/licenses/by/4.0/

Table 1: Benchmark results of TaggerOne [4]

Corpus	Precision	Recall	F-measure
NCBI Disease	81.5%	80.8%	82.9%
BioCreativeV CD-R	94.2%	88.8%	91.4%

Table 2: Benchmark results of GNormPlus (Human) [8]

Corpus	Precision	Recall	F-measure
BioCreative II GN	87.1%	86.4%	86.7%

Table 3: Document Counts of CORD19 Sources

General	
Number of Documents	57.4K
Number of full texts	43.5K
JSON parses by source	
PubMedCentral (PMC)	49.7K
Elsevier	24.8K
medRxiv	2.3K
ArXiv	1.2K
bioRxiv	1.1K
Chan Zuckerberg Initiative (CZI)	0.2K

Table 4: Number of Detected Entity Mentions for the CORD-19 (Abstracts and Fulltexts)

Corpus	Chemicals	Diseases	Genes	Species
Abstracts	99K	145K	59K	165K
Fulltexts	3,407K	4,039K	2,232K	4,667K

The U.S. library of medicine[4] provides several expert-built tools come with a high quality for detecting entity mentions in text. These tools can be used via command line interfaces and a freely available. We build a pipeline upon these provided tools to automatically detect the following entity types in text: 1. Chemicals, 2. Diseases, 3. Genes and 4. Species. Chemicals are described by the Medical Subject Heading (MeSH) vocabulary[5]. Diseases are either by MeSH terms or by OMIM[6]. The NCBI Gene Vocabulary [7] is utilized for the Genes' NER and the NCBI Species Taxonomy [8] likewise for the Species' NER.

Chemicals and Diseases are detected by TaggerOne [4], which uses a semi-Markov structured linear classifier to run named entity recognition (NER) and normalization simultaneously, thus improving performance compared to other taggers. GNormPlus [8] is used for detecting Genes and Species, which runs NER and normalization as two separate steps. Both NER tools have been evaluated on real-world text corpora to determine the quality of their detected entity mentions. Benchmarks for the relevant corpora can be found in Tables 1 for TaggerOne and 2 for GNormPlus. NCBI Disease corpus is a testset for analysing diseases and the BioCreativeV corpus is a challenge for detecting Chemicals as well as Diseases. The GNormPlus evaluation is done for a Gene Normalisation testset for humans. Besides, GNormPlus is capable of detecting gene families in texts. For more details about both applications, see [4] for TaggerOne and [8] for GNormPlus.

Pipeline. We have developed a pipeline utilizing TaggerOne and GNormPlus for biomedical NER. Our pipeline expects texts in a so-called PubTator format, see [7] and the description on[9]. As an input, the pipeline supports 1. a single PubTator file, 2. a composed PubTator file and 3. a directory of PubTator files. A composed PubTator file consists of the content of two PubTator files separated by two newlines. Besides, we support the tagging of multiple files in parallel. Therefore, we implemented a splitting of the input and parallel working of the underlying tools. The recognition steps stores it's produced data in a relational database. Finally, the pipeline exports the annotated entity mentions in a desired format like PubTator or JSON.

3 THE COVID-19 OPEN RESEARCH DATASET

Research into COVID-19 is a big challenge and highly relevant at the moment. Therefore, scientists in the medical field must be assisted by innovative tools to access the current state of literature efficiently. The COVID-19 Open Research Dataset Challenge (CORD-19) [2] is a "call to action" for computer scientists in the natural language processing (NLP) and data mining field to develop such innovative tools. The dataset in version 9 consists of ca. 57,000 scholarly articles, of which ca. 44,000 have a PDF parse of their full text attached to them. Articles are taken from various sources, most prominently the PubMedCentral collection. The document statistics of the dataset in version 9 can be seen in Table 3. Some documents are accessible in multiple sources and are counted more than once in the statistics. The abstracts and full texts of the documents are given paragraph wise in a JSON-Format, so the texts can easily be extracted and processed. Entity-centric information access plays a key role in the medical domain [3]. Hence, we run our pipeline upon the challenge dataset to assist the community with valuable entity information.

3.1 Detected Entity Mentions

We report the number of the resulting entity mentions for each entity type. We create two different dumps: one dump contains entity mentions within titles and abstracts and the second dump contains entity mentions in the title, abstract and fulltexts of the documents. Table 4 lists the number of entity mention for both dumps grouped by the entity types. Our pipeline detects nearly 99K Chemicals, 145K Diseases, 59K Genes and 165K Species in titles and abstracts. For fulltexts, the pipeline detects around 3.4M Chemicals,

---

4.0M Diseases, 2.2M Genes and 4,7M Species. We estimate the annotation's quality to be comparable to the reported quality in the tools' original publications.

3.2 Dump of the Entity Mentions

We publish the obtained entity mentions as two JSON files. The first file contains the entity mentions for titles and abstracts. The second file contains the entity mentions for titles, abstracts as well as fulltexts. We process the CORD19 fulltexts by selecting the available JSON files. These JSON files contain fulltexts as sequences of body texts. Hence, a fulltext document consists of a title, an abstract and a sequence of body texts. We publish the corresponding entity mentions suitable for the given structure. Therefore, each entity mentions contains an entity location in texts including:

(1) a paragraph representing the position in the text. 0 is an entity mention in the title, 1 is an entity mention in the abstract and 2 is an entity mention in the first body text field and so on.
(2) a start position representing the position of the first entity's character within the corresponding text (title, abstract, body text element).
(3) an end position representing the position of the last entity's character within the corresponding text.

As an example, an entity location with paragraph 5, start 5 and end 10 means that the entity is mentioned in the third body text field starting at character position 5 and ending at character position 10. The first character has the position 0. An entity mention contains the following components:

(1) an entity location,
(2) an entity string representing the entity's token sequence in the text,
(3) an entity type (Chemical, Disease, Gene and Species), and
(4) an entity id corresponding to the previously described vocabularies.

The computed entity mentions are shared within a JSON file. The JSON file consists of a dictionary, where each CORD19 document id is mapped to a list of entity mentions. A short prototypical snapshot of the exported JSON file is shown below:

```
[
  <paper_id: str>: [ #For every JSON-parse of the dataset
    {  # For every entity mention
      "location": {
        "paragraph": <int> # 0 = title, 1 = abstract
                           # > 1 = body text
        "start": <int> # 0 = first character of paragraph
        "end": <int>
      },
      "entity_str": <str> # entity mention in source text
      "entity_type": <"Chemical"|"Disease"|"Gene"|"Species">
      "entity_id": <str> # e.g. MESH-Identifier
    },...
  ],...
]
```

More details can be found in our regularly updated GitHub repository.

4 SUMMARY AND OUTLOOK

In this paper, we discussed the importance and usefulness of entity mentions for retrieval applications. We developed an effective pipeline to automatically annotate biomedical entity mentions in arbitrary texts. Moreover, we built our pipeline on top of the latest available biomedical NER tools to ensure the quality of our entity mentions.

Applying our pipeline to the COVID-19 open research dataset, we published the resulting entity mentions as a semantically enriched dataset for free reuse on GitHub. We will continuously update our GitHub repository whenever new versions of the COVID-19 dataset are published.

REFERENCES

[1] Laura Dietz, Alexander Kotov, and Edgar Meij. 2018. Utilizing Knowledge Graphs for Text-Centric Information Retrieval. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval* (Ann Arbor, MI, USA) *(SIGIR '18)*. Association for Computing Machinery, New York, NY, USA, 1387–1390. https://doi.org/10.1145/3209978.3210187

[2] Alan Institute for AI, Anthony Goldbloom et al., and The White House. 2020. COVID-19 Open Research Dataset Challenge (CORD-19), Version 9. Retrieved April 27, 2020 from https://www.kaggle.com/dataset/08dd9ead3afd4f61ef246bfd6aee098765a19d9f6dbf514f0142965748be859b/version/9.

[3] Jorge R. Herskovic, Len Y. Tanaka, William Hersh, and Elmer V. Bernstam. 2007. A Day in the Life of PubMed: Analysis of a Typical Day's Query Log. *Journal of the American Medical Informatics Association* 14, 2 (03 2007), 212–220.

[4] Robert Leaman and Zhiyong Lu. 2016. TaggerOne: joint named entity recognition and normalization with semi-Markov Models. *Bioinformatics* 32, 18 (06 2016), 2839–2846. https://doi.org/10.1093/bioinformatics/btw343 arXiv:https://academic.oup.com/bioinformatics/article-pdf/32/18/2839/24406872/btw343.pdf

[5] Pablo N. Mendes, Max Jakob, Andrés García-Silva, and Christian Bizer. 2011. DBpedia Spotlight: Shedding Light on the Web of Documents. In *Proceedings of the 7th Int. Conf. on Semantic Systems* (Graz, Austria) *(I-Semantics '11)*. Association for Computing Machinery, New York, NY, USA, 1–8.

[6] Francesco Piccinno and Paolo Ferragina. 2014. From TagME to WAT: A New Entity Annotator. In *Proceedings of the First Int. Workshop on Entity Recognition & Disambiguation* (Gold Coast, Queensland, Australia) *(ERD '14)*. Association for Computing Machinery, New York, NY, USA, 55–62.

[7] Chih-Hsuan Wei, Hung-Yu Kao, and Zhiyong Lu. 2013. PubTator: a web-based text mining tool for assisting biocuration. *Nucleic Acids Research* 41, W1 (05 2013), W518–W522. https://doi.org/10.1093/nar/gkt441 arXiv:https://academic.oup.com/nar/article-pdf/41/W1/W518/3859973/gkt441.pdf

[8] Chih-Hsuan Wei, Hung-Yu Kao, and Zhiyong lu. 2015. GNormPlus: An Integrative Approach for Tagging Genes, Gene Families, and Protein Domains. *BioMed research international* 2015 (09 2015), 918710. https://doi.org/10.1155/2015/918710