

---

# ELSEVIER OA CC-BY CORPUS

---

**Daniel Kershaw**  
Elsevier Ltd  
London, UK  
d.kershaw@elsevier.com

**Rob Koeling**  
Elsevier Ltd  
London, UK  
r.koeling@elsevier.com

August 6, 2020

## ABSTRACT

We introduce the Elsevier OA CC-BY corpus. This is the first open corpus of Scientific Research papers which has a representative sample from across scientific disciplines. This corpus not only includes the full text of the article, but also the metadata of the documents, along with the bibliographic information for each reference.

## 1 Introduction

We are pleased to release the Elsevier OA CC-BY corpus for NLP and AI research. This is a corpus of 40k (40,001) open access (OA) CC-BY articles from across Elsevier's journals representing a large scale, cross-discipline set of research data to support NLP and ML research.

Research into the application of NLP and Machine Learning to scholarly content has attracted considerable attention in recent years. However, progress has been held back because of limited availability of large, cross-discipline datasets. Through releasing this dataset we hope to aid the research community in their work to expand the understanding of commonalities and differences between processing of scientific text and text of a different nature (e.g. news text). Moreover, this dataset allows research on challenges in processing scientific text that do not exist for other types of data.

In this article we want to report on details of the dataset being released, focusing on the structure of the data, coverage of features and tools which can be used with it.

## 2 Related Work

Open source datasets of academic articles are not new, notable corpora are Pub-Med OA corpus<sup>1</sup>, CiteSeerX [1], ACL [2] and arXiv [3]. The corpora mentioned above represent millions of academic articles freely available online in either PDF,  $\LaTeX$  or free text. However, for these articles meta-data usually have not been resolved, meaning that establishing the connection between documents through citations is challenging. To overcome this, corpora such as the Literature Graph in Semantic Scholar have been released [4], which has recently been superseded by S2ORC [5]. These corpora contain documents from multiple datasets (ACL, arXiv, PubMed · · ·) which have then been resolved by extracting references, venues and other metadata from the articles. However, the original datasets are biased towards certain academic disciplines and domains such as computing. The resulting literature graphs allow for the high resolution mapping of connections between entries found in academic articles, and therefore for the mapping of relations between authors, academic concepts and publication venues, to name a few.

Similar to S2ORC are the datasets which have been compiled in response to world events such as COVID-19 [6]. The COVID-19 dataset was derived from a number of already existing datasets and supplemented with additional published articles. As this dataset was designed for a particular cause, it is biased towards the biological and medial fields. The research output of the COVID19 dataset is extensive, including document summerisation [7], drug re-purposing [8]

---

<sup>1</sup><https://www.ncbi.nlm.nih.gov/pmc/tools/openftlist/>

and risk factor identification [9]. A comprehensive overview will take a review article in itself to summarise this contribution. This in essence shows the usefulness of domain specific corpora.

The release of these types of datasets opened the doors for the development of task-specific datasets focused on curating data for one or more specific tasks. For example, [10] released labelled data for developing models for citation intent classification. Whereas [11] released annotated articles to develop document level information extraction systems. Specific task datasets for co-reference resolution and entity extraction usually come from shared tasks from workshops such as SemEval, TREC and BioNLP. High-quality annotations are required for these tasks, which can be challenging (meaning time consuming and expensive) to produce. The corresponding datasets are generally small, highly curated and often focus on one particular phenomenon, topic domain or text genre. Very good overviews of available datasets for typical tasks are presented in [12] (for Named Entity Recognition) and in [13] (for co-reference resolution). These overviews show that datasets containing academic writing are mostly limited to the bio-medical domain. Moreover, the limited size of these data sets constitute a serious limitation. Even when models developed on the basis of this data show promising results, there are limited options to validate the results across domains or using larger corpora.

### 3 Dataset

This section details how the data set was constructed, which fields are included and coverage of fields in the dataset.

#### 3.1 Data Sampling

The corpus consists of scientific articles published by Elsevier since the beginning of 2014 which are open access (OA) and covered by a CC-BY 4.0 licence.<sup>2</sup> In order to create a balanced dataset from across academic disciplines, a stratified sampling method was used to achieve equal representation from across Elsevier disciplines as represented by ASJC (All Science Journal Classification) codes. ASJC codes represent the scientific discipline of the journals in which the article was published. To simplify this, the 334 ASJC codes were grouped into their 27 top-level subject classifications.<sup>3</sup>

2000 documents were sampled from each of the 27 top-level ASJC classes. Each article can have multiple ASJC codes, if the article was selected for one class then it was removed from the remaining pool. The resulting sample of documents is balanced across disciplines (see Table 3 for the break down of document per discipline)

#### 3.2 Data Structure

Each document within the corpus is contained within its own JSON file. The name of the file is the ID of the article. The data for each article is structured as described in the JSON scheme and field descriptions below.

```

1 {
2   "docId": <str>,
3   "metadata": {
4     "title": <str>,
5     "authors": [
6       {
7         "first": <str>,
8         "initial": <str>,
9         "last": <str>,
10        "email": <str>
11      },
12      ...
13    ],
14    "issn": <str>,
15    "volume": <str>,

```

<sup>2</sup>This license allows users to copy, to create extracts, abstracts and new works from the article, to alter and revise the article, and to make commercial use of it (including re-use and/or resale of the article by commercial entities), provided the user gives appropriate credit (with a link to the formal publication through the relevant DOI), provides a link to the license, indicates if changes were made and the licensor is not represented as endorsing the use made of the work. <https://www.elsevier.com/about/policies/open-access-licenses/user-licences>

<sup>3</sup>Serial titles are classified using the ASJC scheme. This is done by in-house experts at the moment the serial title is set up for Scopus coverage; the classification is based on the aims and scope of the title, and on the content it publishes. [https://service.elsevier.com/app/answers/detail/a\\_id/15181/supporthub/scopus/](https://service.elsevier.com/app/answers/detail/a_id/15181/supporthub/scopus/)

```

16     "firstpage": <str>,
17     "lastpage": <str>,
18     "pub_year": <int>,
19     "doi": <str>,
20     "pmid": <str>,
21     "openaccess": "Full",
22     "subjareas": [<str>],
23     "keywords": [<str>],
24     "asjc": [<int>],
25 },
26 "abstract": [
27     {
28         "sentence": <str>,
29         "startOffset": <int>,
30         "endOffset": <int>
31     },
32     ...
33 ],
34 "bib_entries": {
35     <str>: {
36         "title": <str>,
37         "authors": [
38             {
39                 "last": <str>,
40                 "initial": <str>,
41                 "first": <str>
42             },
43             ...
44         ],
45         "issn": <str>,
46         "volume": <str>,
47         "firstpage": <str>,
48         "lastpage": <str>,
49         "pub_year": <int>,
50         "doi": <str>,
51         "pmid": <str>
52     },
53     ...
54 },
55 "body_text": [
56     {
57         "sentence": <str>,
58         "secId": <str>,
59         "startOffset": <int>,
60         "endOffset": <int>,
61         "title": <str>,
62         "refoffsets": {
63             <str>: {
64                 "endOffset": <int>,
65                 "startOffset": <int>
66             }
67         },
68         "parents": [
69             {
70                 "id": <str>,
71                 "title": <str>
72             },
73             ...
74         ]

```

```

75   },
76   ...
77   ]
78 }

```

**docId** The docID is the identifier of the document. This is unique to the document, and can be resolved into a URL for the document through the addition of <https://www.sciencedirect.com/science/pii/<docId>>

**abstract** This is the author provided abstract for the document

**body\_text** The full text for the document. The text has been split on sentence boundaries, thus making it easier to use across research projects. Each sentence has the title (and ID) of the section which it is from, along with titles (and IDs) of the parent section. The highest-level section takes index 0 in the parents array. If the array is empty then the title of the section for the sentence is the highest level section title. This will allow for the reconstruction of the article structure. References have been extracted from the sentences. The IDs of the extracted reference and their respective offset within the sentence can be found in the “refoffsets” field. The complete list of references are can be found in the “bib\_entry” field along with the references’ respective metadata.

**bib\_entities** All the references from within the document can be found in this section. If the meta data for the reference is available, it has been added against the key for the reference. Where possible information such as the document titles, authors, and relevant identifiers (DOI and PMID) are included. The keys for each reference can be found in the sentence where the reference is used with the start and end offset of where in the sentence that reference was used.

**metadata** Meta data includes additional information about the article, such as list of authors, relevant IDs (DOI and PMID). Along with a number of classification schemes such as ASJC and Subject Classification.

**Author\_highlights** Author highlights were included in the corpus where the author(s) have provided them. The coverage is 61% of all articles. The author highlights, consisting of 4 to 6 sentences, is provided by the author with the aim of summarising the core findings and results in the article.

### 3.3 Data Coverage and Statistics

Table 1: JSON Field coverage

Field	Number of Articles
Abstract	98.81
Body_text	100.00
Author_highlight	62.12
Metadata	100.00
Metadata - issn	100.00
Metadata - firstpage	88.41
Metadata - lastpage	88.18
Metadata - pub_year	100.00
Metadata - doi	100.00
Metadata - openaccess	100.00
Metadata - subjectareas	100.00
Metadata - keywords	100.00
Metadata - asjc	100.00
Bib_entries	97.20

Table 2: Distribution of Publication Years

Publication Year	Number of Articles
2014	3252
2015	4556
2016	5939
2017	6684
2018	7780
2019	9503
2020	2292

Table 3: Distribution of Articles Per Mid Level ASJC Code. Each article can belong to multiple ASJC codes.

<b>Discipline</b>	<b>Number of Articles</b>
General	310
Agricultural and Biological Sciences	3985
Arts and Humanities	1014
Biochemistry, Genetics and Molecular Biology	8417
Business, Management and Accounting	1002
Chemical Engineering	2196
Chemistry	2749
Computer Science	3004
Decision Sciences	530
Earth and Planetary Sciences	2764
Economics, Econometrics and Finance	1081
Energy	2845
Engineering	5962
Environmental Science	6241
Immunology and Microbiology	3258
Materials Science	4008
Mathematics	1561
Medicine	9225
Neuroscience	3277
Nursing	310
Pharmacology, Toxicology and Pharmaceutics	2233
Physics and Astronomy	3927
Psychology	1796
Social Sciences	3623
veterinary	1010
Dentistry	43
Health Professions	774

## 4 Use of Dataset

The dataset [14] can be downloaded from Mendeley Data.<sup>4</sup> The dataset contains a formatted JSON version of the raw XML which can be accessed through the Elsevier APIs.<sup>5</sup> The original XML files can be processed with the Annotation-Query<sup>6</sup> framework released by Elsevier Labs.

Please use the following reference when using the dataset:

```
@dataset{https://10.17632/zm33cdndx.2,
  doi = {10.17632/zm33cdndx.2},
  url = {https://data.mendeley.com/datasets/zm33cdndx/2},
  author = "Daniel Kershaw and Rob Koeling",
  keywords = {Science, Natural Language Processing, Machine Learning, Open Dataset},
  title = {Elsevier OA CC-BY Corpus},
  publisher = {Mendeley},
  year = {2020},
  month = aug
}
```

<sup>4</sup><https://data.mendeley.com/datasets/zm33cdndx/2>

<sup>5</sup><https://dev.elsevier.com/>

<sup>6</sup><https://github.com/elsevierlabs-os/AnnotationQuery>

## 5 Conclusion and Future Work

This dataset was released to support the development of ML and NLP models targeting science articles from across all research domains. While the release builds on other datasets designed for specific domains and tasks, it will allow for similar datasets to be derived or for the development of models which can be applied and tested across domains.

## 6 Acknowledgements

We would like to thank all our colleges at Elsevier for helping to pull together this dataset. Especially Darin McBeath from Elsevier Labs, and Anita Dewaard and Noelle Gracy from the Research Collaboration Office.

## References

- [1] C. Lee Giles, Kurt D. Bollacker, and Steve Lawrence. Citeseer: An automatic citation indexing system. In *Proceedings of the Third ACM Conference on Digital Libraries*, DL '98, page 89–98, New York, NY, USA, 1998. Association for Computing Machinery.
- [2] Dragomir R. Radev, Pradeep Muthukrishnan, and Vahed Qazvinian. The acl anthology network corpus. In *Proceedings of the 2009 Workshop on Text and Citation Analysis for Scholarly Digital Libraries*, NLP4DL '09, page 54–61, USA, 2009. Association for Computational Linguistics.
- [3] Tarek Saier and Michael Färber. unarxive: a large scholarly data set with publications' full-text, annotated in-text citations, and links to metadata. *Scientometrics*, pages 1–24, 2020.
- [4] Waleed Ammar, Dirk Groeneveld, Chandra Bhagavatula, Iz Beltagy, Miles Crawford, Doug Downey, Jason Dunkelberger, Ahmed Elgohary, Sergey Feldman, Vu Ha, and et al. Construction of the literature graph in semantic scholar. *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 3 (Industry Papers)*, 2018.
- [5] Kyle Lo, Lucy Lu Wang, Mark Neumann, Rodney Kinney, and Daniel Weld. S2orc: The semantic scholar open research corpus. *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 2020.
- [6] Lucy Lu Wang, Kyle Lo, Yoganand Chandrasekhar, Russell Reas, Jiangjiang Yang, Doug Burdick, Darrin Eide, Kathryn Funk, Yannis Katsis, Rodney Kinney, Yunyao Li, Ziyang Liu, William Merrill, Paul Mooney, Dewey Murdick, Devvret Rishi, Jerry Sheehan, Zhihong Shen, Brandon Stilson, Alex Wade, Kuansan Wang, Nancy Xin Ru Wang, Chris Wilhelm, Boya Xie, Douglas Raymond, Daniel S. Weld, Oren Etzioni, and Sebastian Kohlmeier. Cord-19: The covid-19 open research dataset, 2020.
- [7] Dan Su, Yan Xu, Tiezheng Yu, Farhad Bin Siddique, Elham J. Barezi, and Pascale Fung. Caire-covid: A question answering and multi-document summarization system for covid-19 research, 2020.
- [8] Qingyun Wang, Manling Li, Xuan Wang, Nikolaus Parulian, Guangxing Han, Jiawei Ma, Jingxuan Tu, Ying Lin, Haoran Zhang, Weili Liu, Aabhas Chauhan, Yingjun Guan, Bangzheng Li, Ruisong Li, Xiangchen Song, Heng Ji, Jiawei Han, Shih-Fu Chang, James Pustejovsky, Jasmine Rah, David Liem, Ahmed Elsayed, Martha Palmer, Clare Voss, Cynthia Schneider, and Boyan Onyshkevych. Covid-19 literature knowledge graph construction and drug repurposing report generation, 2020.
- [9] Francis Wolinski. Automatic extraction of risk factors from covid-19 literature. 2020.
- [10] Arman Cohan, Waleed Ammar, Madeleine van Zuylen, and Field Cady. Structural scaffolds for citation intent classification in scientific publications. *Proceedings of the 2019 Conference of the North*, 2019.
- [11] Sarthak Jain, Madeleine van Zuylen, Hannaneh Hajishirzi, and Iz Beltagy. Scirex: A challenge dataset for document-level information extraction. *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 2020.
- [12] Archana Goyal, Vishal Gupta, and Manish Kumar. Recent named entity recognition and classification techniques: A systematic review. *Computer Science Review*, 29:21 – 43, 2018.
- [13] Rhea Sukthanker, Soujanya Poria, Erik Cambria, and Ramkumar Thirunavukarasu. Anaphora and coreference resolution: A review. *Information Fusion*, 59:139 – 162, 2020.
- [14] Daniel Kershaw and Rob Koeling. Elsevier oa cc-by corpus, August 2020.