

# unarXive 2024: A Large-Scale Scientific Corpus for Citation-Aware Retrieval and Generation

Ines Besrou and Michael Färber

ScaDS.AI, Technische Universität Dresden, Germany  
besrouines@gmail.com, michael.farber@tu-dresden.de

## Abstract

Full-text collections of scientific papers are essential for NLP research and the training of language models. However, existing resources remain incomplete: they often lag behind the fast-paced growth of scientific publishing, lack comprehensive citation networks, and discard essential structural elements. In this work, we introduce *unarXive 2024*, a large-scale, richly structured corpus containing every arXiv submission from January 1991 to December 2024 – over 2.28 million documents across physics, mathematics, computer science, and other fields. Our release enhances each paper with detailed metadata, reconstructs a substantially more complete citation network than existing datasets, and preserves fine-grained structural information, including section boundaries, mathematical notation, and non-textual elements. Beyond the corpus itself, we provide dense and sparse indexes optimized for retrieval-augmented generation (RAG) over the full arXiv archive. All resources, including code and data, are publicly available: <https://github.com/faerber-lab/unarXive-2024>

**Keywords:** Scientific Corpus, Information Retrieval, Open Science

## 1. Introduction

Full-text datasets from academic publications are becoming increasingly important for research in NLP and scholarly analysis (Lo et al., 2020). While metadata-level resources (e.g., OpenAlex (Priem et al., 2022) and SemOpenAlex (Färber et al., 2023)) have supported research in bibliometrics, citation analysis, and trend detection, recent efforts have aimed to enrich datasets with structured full-text content to enable more advanced natural language processing (NLP) tasks such as scientific summarization (Cohan et al., 2018), citation recommendation (Bhagavatula et al., 2018), and knowledge graph construction (Ammar et al., 2018). Fundamental elements of such datasets are consistent full-text formatting, comprehensive and well-structured metadata, and the availability and structure of citation data.

Despite their importance, available datasets reveal consistent deficiencies across several essential criteria, as shown in Table 1. For example, there are (1) limitations in size and scope in *SciXGen*, (2) no citation network data in *arXMLiv*, (3) no or only limited handling of mathematical notation in *S2ORC*, and (4) outdated coverage, incomplete citation linking, and inefficient full-text structuring in *unarXive 2022*. In particular, the corpus only includes publications up to December 2022 and has not been updated to reflect the rapidly growing number of new submissions to arXiv. Moreover, while *unarXive 2022* includes citation links, coverage remains incomplete, with many references unresolved, which reduces reliability for citation-based tasks.

To address key limitations in existing datasets, we present the *unarXive 2024* dataset. Covering

over 2.28 million arXiv papers published between 1991 and 2024, it introduces several enhancements to existing scholarly corpora: improved metadata schemas, a substantially expanded citation network with 39.5 million resolved citations extracted from 77.4 million citation contexts (51.2% resolution rate), and a structured representation of paper bodies organized into logical sections. We also release ready-to-use dense and sparse retrieval indexes, making the dataset directly applicable for retrieval-augmented generation (RAG) and related IR/NLP tasks. In line with best practices, both the dataset and source code are openly distributed under a permissive license, adhering to the FAIR data principles. Constructing this corpus involved large-scale  $\LaTeX$  parsing, robust section segmentation, metadata enrichment, language identification, and the resolution of citation contexts using LLM-based BibTeX extraction and OpenAlex matching. The resulting dataset supports a wide range of downstream applications, including section-aware scientific text analysis (Achakulvisut et al., 2021), citation graph modeling and influence estimation (Cohan et al., 2020), and the training of domain-specific language models such as SciBERT (Beltagy et al., 2019) and SciNCL (Ostendorff et al., 2022). Its structured format and citation-resolved full text also make it ideally suited for retrieval-augmented generation pipelines, including scientific question answering systems (Auer et al., 2023) that require verifiable, citation-aware outputs.

Overall, we make the following contributions:

- **Comprehensive corpus coverage:** We provide a structured dataset of 2.28 million arXiv papers (1991–2024) with full text, metadata, and 39.5 million resolved citation links.

Table 1: Comparison of large data sets derived from paper full-texts

Data Set	Source Data	Citation Network	Structured Doc.	Math.	# Docs	Disciplines	Purpose
CORE (Pontika et al., 2016)	multiple	-	×	×	> 100M	various	general NLP
S2ORC (PDF) (Lo et al., 2020)	multiple	69.4%	✓	×	12M	various	general NLP
unarXive 2020 (Saier and Färber, 2020)	arXiv.org	42.6%	×	✓	1.2M	phys., maths, CS	general NLP
CORD19 (Wang et al., 2020)	multiple	-	✓	✓	1M	Med., Bio., Chemistry	Pandemic research
arXMLiv (Ginev, 2020)	arXiv.org	-	✓	✓	1.6M	phys., maths, CS	math linguistics
SciXGen (Chen et al., 2021)	arXiv.org	41.6%	✓	✓	205k	CS	text generation
unarXive 2022 (Saier et al., 2023)	arXiv.org	44.4%	✓	✓	1.9M	phys., maths, CS	general NLP
<b>unarXive 2024 (ours)</b>	arXiv.org	51.2%	✓	✓	2.28M	phys., maths, CS	general NLP

- **Retrieval-ready infrastructure:** We release dense and sparse indexes to support retrieval-augmented generation (RAG) and scientific information retrieval tasks.
- **Open and FAIR data practices:** All data and code are publicly available under FAIR principles to enable reuse and transparency.

## 2. Related Work

Large-scale corpora derived from scientific full-texts differ significantly in structure, citation completeness, and temporal scope (see Table 1). *CORE* (Pontika et al., 2016) emphasizes scale but lacks structured full text, limiting discourse-aware modeling. *S2ORC (PDF)* (Lo et al., 2020) retains section boundaries and partial citation resolution, but relies on PDF-based parsing, leading to noisy citation contexts, loss of structural fidelity, and absence of math formulas. It has also remained frozen since 2020, rendering it increasingly outdated. Domain-focused datasets like *CORD19* (Wang et al., 2020) offer structured full texts with math content but are restricted to biomedical literature. *arXMLiv* (Ginev, 2020) preserves LaTeX structure and mathematics but lacks citation links and complete paper coverage. *SciXGen* (Chen et al., 2021) delivers high-quality structured content for text generation but is limited in scale and scope. *unarXive 2022* (Saier et al., 2023) combined structured full text, math notation, and citation data, but its 44.4% citation resolution, outdated 2022 cutoff, and lack of retrieval indexes limit its suitability for citation analysis and RAG-based applications.

In contrast, *unarXive 2024* addresses several important limitations: It extends coverage through December 2024, improves citation resolution to 51.2% across 39.5M links, and leverages native LaTeX sources to preserve structure and mathematics across 2.28M documents. Moreover, it introduces dense and sparse retrieval indexes tailored for large-scale retrieval-augmented generation (RAG), enabling robust scientific QA, citation-grounded text generation, and influence modeling without relying on external databases or toolchains.

## 3. Methodology

Conceptually, our data creation pipeline consists of three main stages: (1) document parsing, (2) metadata augmentation, and (3) reference linking, each of which is described in detail below.

### 3.1. Full-Text Parsing and Structuring

Most arXiv papers are available as LaTeX source files, which makes them uniquely suitable for structured large-scale processing. To build our dataset, we downloaded all available source packages directly from arxiv.org. Each package typically contains multiple .tex files, figures, and bibliographies. We begin by flattening each paper’s LaTeX source into a single .tex file using the latexexpand tool. The flattened file is then converted to XML using *Tralics*, from which we generate a structured, machine-readable JSON representation. During this process, we preserve and annotate the document’s logical structure by grouping content into sections. This grouping enables consistent access to standard parts of a paper (e.g., *Introduction*, *Related Work*, *Conclusion*), which is particularly important for training section-aware models. We also retain URLs from embedded links, which aids in more accurate reference resolution (see Section 3.3). Additionally, we mark the in-text positions of citation markers, tables, figures, and mathematical expressions. These are then explicitly linked to their respective references, captions (as textual surrogates), and original  $\LaTeX$  code to ensure semantic fidelity.

### 3.2. Metadata Augmentation

Adding comprehensive metadata facilitates more effective analyses, both within and across disciplines. For each paper, we include core metadata fields such as `paper_id`, `title`, `authors`, `submitter`, `license`, `doi`, and `categories`. As an enhancement over previous versions, we additionally provide a language field (e.g., "en") to indicate the primary language of the paper. We first infer the document language using metadata provided by the OpenAlex API, which includes language annotations derived from the full text. If this metadata is

unavailable or missing, we fall back to automatic detection using `langdetect`, a lightweight language identification library (Mimino, 2023).

We also provide a `cited_by_count` field, capturing how often each paper has been cited within our dataset up to May 2025. Citation counts are obtained via the OpenAlex API, one of the most comprehensive and current open repositories of scholarly metadata, using DOI-based lookup when available, and exact title matching as a fallback. This approach yielded citation counts for 2,225,191 papers (79.4%). The inclusion of this field enables bibliometric analysis and citation impact modeling.

### 3.3. Reference Linking

To construct a reliable citation network, raw bibliographic references, typically stored as plain text in `.bib` files, must be accurately linked to the documents they cite. This enables key downstream applications such as citation graph modeling, influence estimation, and scientific recommendation systems. To complete the citation network, each reference is linked to a unique identifier: either a DOI or an OpenAlex ID. Earlier efforts, such as *unarXive 2022*, relied on rule-based tools like Neural Parscit and heuristic methods to extract identifiers (e.g., DOIs, arXiv IDs) and match references against large metadata corpora. While partially effective, these approaches struggled with format variability and incomplete coverage. By integrating both DOIs and OpenAlex IDs, we improve recall, support interoperability across bibliographic systems, and ensure compatibility with both publisher metadata and open citation graphs.

In contrast, our approach extracts raw BibTeX strings for each citation and employs the open-source Falcon-3-10B language model (Almazrouei et al., 2023) to identify cited paper titles, even under non-standard formatting. Falcon-3-10B was selected for its strong performance on information extraction tasks, efficient inference, and permissive licensing, making it well-suited for scalable and transparent citation processing (Almazrouei et al., 2023). The extracted titles are then matched to OpenAlex records via API, yielding OpenAlex IDs and DOIs when available. To quantify citation coverage, we consider all in-text citation markers in the dataset and attempt to link each cited reference to a persistent identifier, i.e., a DOI or an OpenAlex ID. The citation resolution rate is defined as the proportion of citations successfully linked in this way, including both links newly added by our pipeline and links already present in *unarXive 2022*. To reduce false positives, we additionally verify the list of authors, since title-based matching alone can introduce noise for papers with highly similar titles. Typical failure cases include incomplete or noisy bibliographic entries, ambiguous titles, and miss-

ing or inconsistent external metadata. Overall, this process resolves 51.2% of all citations.

### 3.4. Schema Design

Processed arXiv papers in the *unarXive 2024* dataset are stored as structured JSONL files, following a standardized schema specifically designed to support a variety of downstream NLP and IR tasks. This structure balances human readability and machine efficiency, and is particularly well-suited for applications such as document retrieval, citation analysis, and question answering.

The JSONL representation of each paper captures three core components: metadata, full-text content, and supplementary data. In Listing 1, we provide an excerpt from the structured output for a sample paper to illustrate how plain text and structured fields are organized. A sample of our dataset is also available in our GitHub repository.

## 4. Final Dataset

Our data creation pipeline yields a comprehensive corpus of 2,283,685 arXiv papers. In total, 39,456,244 in-text citation markers were resolved to persistent identifiers (DOIs or OpenAlex IDs), corresponding to a citation resolution rate of 51.2% over all citation markers in the dataset, and resulting in one of the most complete and up-to-date citation networks.

To support retrieval-augmented generation (RAG) and related downstream tasks, we provide a structured representation of each paper in the form of key-value dictionaries, where each arXiv ID maps to a cleaned and segmented set of paper sections (e.g., title, abstract, introduction). This facilitates efficient access to semantically meaningful content blocks and supports use cases such as scientific question answering and summarization.

In addition, we provide pre-built retrieval indexes: a dense vector index using the `e5` model and a sparse BM25 index, both built over abstracts for fast neural and lexical retrieval. Unlike existing scholarly QA stores, which are often small, outdated, and inconsistently structured, our infrastructure supports large-scale, up-to-date retrieval across millions of papers, making it well-suited for modern RAG pipelines.

In the sections that follow, we outline how the dataset is distributed to ensure ease of access and adoption by the broader research and practitioner communities, and we provide detailed statistics and usage trends to support a deeper understanding of scholarly publishing patterns. Lastly, we outline how the dataset can be applied in various settings.

## 4.1. Distribution

In alignment with the FAIR principles, we selected established distribution platforms and licenses to make our dataset and source code widely accessible:

- **Dataset:** The dataset is available at [https://huggingface.co/datasets/ines-besroun/unarxive\\_2024](https://huggingface.co/datasets/ines-besroun/unarxive_2024).
- **Indexes:** The retrieval indexes are available at:
  - **BM25:** [https://huggingface.co/datasets/ines-besroun/unarxive\\_BM25](https://huggingface.co/datasets/ines-besroun/unarxive_BM25)
  - **E5:** [https://huggingface.co/datasets/ines-besroun/unarxive\\_E5](https://huggingface.co/datasets/ines-besroun/unarxive_E5)
- **Source Code:** The full codebase for dataset construction is available on GitHub under the permissive MIT License: <https://github.com/faerber-lab/unarXive-2024>.

## 4.2. Data Analysis

Over the past three years, unarXive has seen an average annual increase of approximately 204,064 papers, bringing the total to over 2.28 million by the end of 2024. As shown in Figure 1, the number of submitted papers is growing exponentially.

While the vast majority of submissions are in English (2,267,464 papers), our dataset also includes several hundred papers in other languages—such as French (560), German (47), Catalan (40), and 25 additional languages. Although limited in size, these multilingual entries offer valuable opportunities for tasks such as language identification, filtering, and benchmarking cross-lingual model behavior on real-world scholarly content. They can also serve as test cases for translation, and cross-lingual citation analysis (Reimers and Gurevych, 2020).

The disciplinary distribution of the unarXive 2024 dataset is shown in Figure 2. Among the papers, 1,209,261 (52.95%) are in physics, followed by 494,924 (21.67%) in computer science, and 451,214 (19.76%) in mathematics. The remaining publications span other scientific disciplines. This distribution reflects the subject coverage of arXiv and may introduce disciplinary biases into downstream retrieval and generation tasks.

Figure 3 shows the total number of publications per year. The continued growth in publication volume reflects the global expansion of research activity, with computer science now surpassing physics in output and approximately doubling every four years (Douglas and Verstyuk, 2025).

```
1 { "paper_id": "2412.00056",
2   "_pdf_hash": null,
3   "_source_hash": "99670c45c2f72706f15c0a
4     0f9b8a7d4e12a34567",
5   "_source_name": "2412.00056.gz",
6   "metadata": { ... },
7   "abstract": {
8     "section": "Abstract",
9     "text": "In recent years, vision...",
10    "cite_spans": [ ... ],
11    "ref_spans": [ ... ]
12  },
13  "bib_entries": { ... },
14  "ref_entries": { ... },
15  "sections": { ... } }
16
17 # zoom on metadata
18 "metadata": {
19   "id": "2412.00056",
20   "submitter": "...",
21   "authors": "...",
22   "title": "Improving Medical ...",
23   "comments": "15 pages",
24   "journal-ref": null,
25   "doi": null,
26   "report-no": null,
27   "categories": "cs.CV cs.AI",
28   "license": "http://creativecommons.org/licenses/by/4.0/",
29   "abstract": "In recent years, vision...",
30   "versions": [ ... ],
31   "update_date": "2024-12-03",
32   "authors_parsed": [ ... ],
33   "language": "en",
34   "cited_by_count": 5,
35   "discipline": "Computer Science" }
36
37 # zoom on sections
38 "sections": {
39   "Introduction": {
40     "text": "In recent years...",
41     "cite_spans": [ ... ],
42     "ref_spans": [ ... ]},
43   "Uncertainty and Consistency": {
44     "text": "Uncertainty quantification
45     ...",
46     "cite_spans": [ ... ],
47     "ref_spans": [ ... ]
48   }
49 }
```

Listing 1: Excerpt of our document representation




Figure 1: Exponential growth of paper submissions from 1991 to 2024.

## 4.3. Applications

The structured full text, enriched metadata, and 39.5 million linked citations in *unarXive 2024* provide a powerful foundation for a broad spectrum of NLP and IR tasks in the scholarly domain. By retaining section boundaries, citation markers, reference mappings, and mathematical expressions,




Figure 2: Disciplinary distribution of the unarXive 2024 dataset.




Figure 3: Annual number of publications in unarXive.

the dataset enables semantically rich modeling at both the sentence and document level.

Its fine-grained citation spans facilitate both *in-line* and *document-level citation recommendation*. Models trained on in-text citation contexts – linked to OpenAlex identifiers – can be readily adapted for tasks such as citation intent detection (Devlin et al., 2019), citation polarity analysis (Li et al., 2013), or integration into systems like C-REX (Färber et al., 2021b) and NeuralCite (Ebesu and Fang, 2017).

The dataset also supports the generation of contextualized paper embeddings. By replacing citation spans with resolved identifiers, researchers can train models like SciBERT (Beltagy et al., 2019), SciNCL (Ostendorff et al., 2022), or SciFive (Phan et al., 2021), effectively capturing how papers are cited in context (Nishikawa and Koshiba, 2024). These embeddings enable downstream applications such as clustering, topic modeling, and personalized paper recommendation (Kanakia et al., 2019).

Additional use cases include large-scale extraction of scientific methods, datasets, and tasks – powering impact analyses at the dataset or method level (e.g., h-indices) (Färber et al., 2021a) and enabling mining of hyperparameter usage patterns at scale (Saier et al., 2024).

Unlike previous scholarly QA resources – which are often PDF-based, limited in scope, or outdated

– *unarXive 2024* enables citation-grounded generation and supports end-user applications such as scientific question answering (Auer et al., 2023), long-form retrieval-based answering (Ajith et al., 2024), and longitudinal analyses.

## 5. Conclusion

We presented *unarXive 2024*, an updated and enriched version of the unarXive dataset, which extends coverage through December 2024 and includes over 2.28 million papers and 39.5 million resolved citation links – significantly improving the citation network. Key features include: (1) size and scope with comprehensive metadata and full-text coverage; (2) an improved and robust citation network enabled by extensive linking of in-text citations to OpenAlex identifiers; and (3) enhanced document structure through grouping and normalization of related sections for better usability.

Future work will focus on enhancing semantic linking by connecting citations not only to papers but also to related datasets and code repositories, thereby enriching the contextual information available. Additionally, we aim to incorporate temporal dynamics to track how papers and their citation networks evolve over time, enabling more detailed trend analysis and insight into the progression of scientific fields.

## 6. Ethical Considerations

**Licensing and compliance.** *unarXive 2024* is derived from publicly available arXiv submissions. We preserve per-paper license information (`license` field) and redistribute structured text and metadata only, not publisher PDFs or other proprietary assets.

**Privacy and personal data.** Author names and affiliations appear as part of the public scholarly record. We do not enrich or infer sensitive attributes.

**Bias and representational equity.** The dataset is dominated by papers from physics, mathematics, and computer science and is overwhelmingly English-language. These disciplinary and linguistic imbalances may affect downstream models, analyses, and evaluations.

**Misuse and safety.** The corpus can be used in retrieval-augmented systems. We encourage citation-grounded generation and do not recommend high-stakes use without expert oversight and domain-specific validation.

**Transparency and accountability.** We provide clear data provenance (arXiv IDs), persistent identifiers (DOIs and OpenAlex IDs where available), and versioned releases.

**Environmental considerations.** To reduce redundant compute, we release ready-to-use indexes.

## 7. Limitations

**Scope.** The dataset covers arXiv submissions from 1991 to 2024 and therefore underrepresents disciplines and languages that are less active on arXiv. Fields such as the humanities, social sciences, and non-English research are only marginally included.

**Parsing noise.** Unconventional or complex  $\text{\LaTeX}$  source files can occasionally lead to imperfect parsing, which may affect section boundaries, mathematical notation, or figure and table references.

**Citation linking.** Not all bibliographic references could be resolved to persistent identifiers such as DOIs or OpenAlex IDs. Unresolved citations should be treated as unknown rather than as missing or incorrect data.

## 8. Bibliographical References

- Titipat Achakulvisut, Daniel E Acuna, Atipat Ruangkittisakul, et al. 2021. Towards democratizing and automating scientific writing with natural language processing. *Nature Machine Intelligence*, 3(3):191–198.
- Anirudh Ajith, Mengzhou Xia, Alexis Chevalier, Tanya Goyal, Danqi Chen, and Tianyu Gao. 2024. [LitSearch: A retrieval benchmark for scientific literature search](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 15068–15083, Miami, Florida, USA.
- Ebtesam Almazrouei, Hamza Alobeidli, Abdulaziz Alshamsi, Alessandro Cappelli, Ruxandra Cojocaru, Mérouane Debbah, Etienne Goffinet, Daniel Hesslow, Julien Launay, Quentin Malaric, Daniele Mazzotta, Badreddine Noune, Baptiste Pannier, and Guilherme Penedo. 2023. [The falcon series of open language models](#). *arXiv preprint arXiv:2311.16867*.
- Waleed Ammar, Dirk Groeneveld, Chandra Bhagavatula, Iz Beltagy, Matt Crawford, Doug Downey, Jason Dunkelberger, Ahmed Elgohary, Sergey Feldman, Vu Ha, et al. 2018. [Construction of the literature graph in semantic scholar](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Demonstrations*, pages 84–91. Association for Computational Linguistics.
- Sören Auer, Dante A C Barone, Cassiano Bartz, Eduardo G Cortes, Mohamad Yaser Jaradeh, Oliver Karras, Manolis Koubarakis, Dmitry Mourmstsev, Dmitrii Pliukhin, Daniil Radyush, Ivan Shilin, Markus Stocker, and Eleni Tsalapati. 2023. The SciQA scientific question answering benchmark for scholarly knowledge. *Scientific Reports*, 13(1):7240.
- Iz Beltagy, Kyle Lo, and Arman Cohan. 2019. [Scibert: A pretrained language model for scientific text](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019*, pages 3613–3618.
- Chandra Bhagavatula, Sergey Feldman, Russell Power, and Waleed Ammar. 2018. Content-based citation recommendation. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Demonstrations*, pages 93–98.
- Hong Chen, Hiroya Takamura, and Hideki Nakayama. 2021. [Scixgen: A scientific paper dataset for context-aware text generation](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 1491–1501.
- Arman Cohan, Franck Dernoncourt, Doo Soon Kim, Trung Bui, Seokhwan Kim, Walter Chang, and Nazli Goharian. 2018. [A discourse-aware attention model for abstractive summarization of long documents](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 615–621.
- Arman Cohan, Sebastian Feldman, Iz Beltagy, Doug Downey, and Daniel S Weld. 2020. [Specter: Document-level representation learning using citation-informed transformers](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2270–2282.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [Bert: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the*

- Association for Computational Linguistics: Human Language Technologies*, NAACL-HLT 2019, pages 4171–4186.
- Michael R. Douglas and Sergiy Verstyuk. 2025. [Progress in artificial intelligence and its determinants](#). *arXiv preprint arXiv:2501.17894*.
- Travis Ebesu and Yi Fang. 2017. [Neural citation network for context-aware citation recommendation](#). In *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR 2017, pages 1093–1096.
- Michael Färber, Alexander Albers, and Felix Schüber. 2021a. [Identifying used methods and datasets in scientific publications](#). In *Proceedings of the Workshop on Scientific Document Understanding co-located with 35th AAAI Conference on Artificial Intelligence*, SDU@AAAI 2021.
- Michael Färber, David Lamprecht, Johan Krause, Linn Aung, and Peter Haase. 2023. [Semopenalex: The scientific landscape in 26 billion RDF triples](#). In *Proceedings of the 22nd International Semantic Web Conference*, ISWC'23, pages 94–112.
- Michael Färber, Vinzenz Zinecker, Isabela Bragaglia Cartus, Sebastian Celis, and Maria Duma. 2021b. [C-rax: A comprehensive system for recommending in-text citations with explanations](#). In *Proceedings of the Companion of The Web Conference*, pages 441–445.
- Deyan Ginev. 2020. [arxmliv:2020 dataset](#), an html5 conversion of arxiv.org. <https://sigmathling.kwarc.info/resources/arxmliv-dataset-2020/>.
- Arjun Kanakia et al. 2019. [Scispacy: Fast and robust library for biomedical-nlp](#). In *BioNLP*.
- Xiang Li, Yifan He, Adam Meyers, and Ralph Grishman. 2013. [Towards fine-grained citation function classification](#). In *Proceedings of Recent Advances in Natural Language Processing*, RANLP 2013, pages 402–407.
- Kyle Lo, Lucy Lu Wang, Mark Neumann, Rodney Kinney, and Daniel S. Weld. 2020. [S2orc: The semantic scholar open research corpus](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*.
- Ilya Mimino. 2023. [langdetect: Language detection library for python](#). <https://github.com/Mimino666/langdetect>.
- Kai Nishikawa and Hitoshi Koshiba. 2024. [Exploring the applicability of large language models to citation context analysis](#). *arXiv preprint arXiv:2409.02443*.
- Malte Ostendorff, Nils Rethmeier, Isabelle Augenstein, Bela Gipp, and Georg Rehm. 2022. [Neighborhood contrastive learning for scientific document representations with citation embeddings](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, EMNLP 2022.
- Long N. Phan, James T. Anibal, Hieu Tran, Shaurya Chanana, Erol Bahadroglu, Alec Peltekian, and Grégoire Altan-Bonnet. 2021. [Scifive: a text-to-text transformer model for biomedical literature](#). *CoRR*, abs/2106.03598.
- Nancy Pontika, Petr Knoth, Matteo Cacellieri, and Samuel Pearce. 2016. [Developing infrastructure to support closer collaboration of aggregators with open repositories](#). *LIBER Quarterly: The Journal of the Association of European Research Libraries*.
- Jason Priem, Heather Piwowar, and Richard Orr. 2022. [Openalex: A fully-open index of scholarly works, authors, venues, institutions, and concepts](#). *arXiv preprint arXiv:2205.01833*.
- Nils Reimers and Iryna Gurevych. 2020. [Making monolingual sentence embeddings multilingual using knowledge distillation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*, EMNLP 2020, pages 4512–4525.
- Tarek Saier and Michael Färber. 2020. [unarxive: A large scholarly data set with publications' full-text, annotated in-text citations, and links to metadata](#). *Scientometrics*, 124(3):1809–1833.
- Tarek Saier, Johan Krause, and Michael Färber. 2023. [unarxive 2022: All arxiv publications pre-processed for nlp, including structured full-text and citation network](#). In *ACM/IEEE Joint Conference on Digital Libraries*, JCDL 2023, pages 66–70.
- Tarek Saier, Mayumi Ohta, Takuto Asakura, and Michael Färber. 2024. [Hyperpie: Hyperparameter information extraction from scientific publications](#). In *Proceedings of the 46th European Conference on Information Retrieval*, ECIR'24, pages 254–269.
- Lucy Lu Wang, Kyle Lo, Yoganand Chandrasekhar, Russell Reas, Jiangjiang Yang, Douglas 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 D. Wade, Kuansan Wang, Nancy Xin Ru Wang, Chris Wilhelm, Boya Xie, Douglas Raymond, Daniel S. Weld, Oren Etzioni,

and Sebastian Kohlmeier. 2020. [Cord-19: The covid-19 open research dataset](#). *arXiv preprint*, arXiv:2004.10706.

## 9. Language Resource References

Ines Besrouer and Michael Färber. 2025a. [unarXive-2024 BM25 Index](#). Hugging Face.

Ines Besrouer and Michael Färber. 2025b. [unarXive-2024 Dataset](#). Hugging Face.

Ines Besrouer and Michael Färber. 2025c. [unarXive-2024: Dataset Construction Codebase](#). GitHub.

Ines Besrouer and Michael Färber. 2025d. [unarXive-2024 E5 Index](#). Hugging Face.