REFCAT: The Fatcat Citation Graph

Martin Czygan

Internet Archive
San Francisco, California, USA
martin@archive.org

Bryan Newbold

Internet Archive
San Francisco, California, USA
bnewbold@archive.org

Abstract

As part of its scholarly data efforts, the Internet Archive releases a first version of a citation graph dataset, named *refcat*, derived from scholarly publications and additional data sources. It is composed of data gathered by the fatcat cataloging project[1], related web-scale crawls targeting primary and secondary scholarly outputs, as well as metadata from the Open Library[2] project and Wikipedia[3]. This first version of the graph consists of over 1.3B citations. We release this dataset under a CC0 Public Domain Dedication, accessible through an archive item[4]. The source code used for the derivation process, including exact and fuzzy citation matching, is released under an MIT license[5].

Index terms— Citation Graph, Web Archiving

1 Introduction

The Internet Archive releases a first version of a citation graph dataset derived from a raw corpus of about 2.5B references gathered from metadata and data obtained by PDF extraction and annotation tools such as GROBID [18]. Additionally, we consider integration with metadata from Open Library and Wikipedia. The goal of this report is to describe briefly the current contents and the derivation of the dataset. We expect this dataset to be iterated upon, with changes both in content and processing.

According to [12] over 50M scholarly articles have been published (from 1726) up to 2009, with the rate of publications on the rise [16]. In 2014, a study based on academic search engines estimated that at least 114M English-language scholarly documents are accessible on the web [13].

Modern citation indexes can be traced back to the early computing age, when projects like the Science Citation Index (1955) [8] were first devised, living on in existing commercial knowledge bases today. Open alternatives were started such as the Open Citations Corpus (OCC) in 2010 - the first version of which contained 6,325,178 individual references [25]. Other notable projects include CiteSeer [9], CiteSeerX [29] and CitEc[6]. The last decade has seen the emergence of more openly available, large scale citation projects like Microsoft Academic [27] or the Initiative for Open Citations[7] [26]. In 2021, over one billion citations are publicly available, marking a "tipping point" for this category of data [11].

While a paper will often cite other papers, more citable entities exist such as books or web links and within links a variety of targets, such as web pages, reference entries, protocols or datasets. References can be extracted manually or through more automated methods, such as metadata access and structured data extraction from full text documents; the latter offering the benefits of scalability. The completeness of bibliographic metadata ranges from documents with one or more persistant identifiers to raw, potentially unclean strings partially describing a scholarly artifact.

2 Related Work

Two typical problems that arise in the process of compiling a citation graph dataset are related to data aquisition and citation matching. Data acquisition itself can take different forms: bibliographic metadata can contain explicit reference data as provided by publishers and aggregators; this data can be relatively consistent when looked at per source, but may vary in style and comprehensiveness when looked

---

[1] https://fatcat.wiki
[2] https://openlibrary.org
[3] https://wikipedia.org
[4] https://archive.org/details/refcat_2021-07-28
[5] https://gitlab.com/internetarchive/refcat

[6] https://citec.repec.org
[7] https://i4oc.org

at as a whole. Another way of acquiring bibliographic metadata is to analyze a source document, such as a PDF (or its text), directly. Tools in this category are often based on condital random fields [14] and have been implemented in projects such as ParsCit [3], Cermine [28], EXCITE [10] or GROBID [18].

The problem of citation matching is relatively simple when common, persistent identifiers are present in the data. Complications mount, when there is *Identity Uncertainty*, that is "objects are not labeled with unique identifiers or when those identifiers may not be perceived perfectly" [22]. Cite-Seer has been an early project concerned with citation matching [9]. A taxonomy of potential issues common in the matching process has been compiled by [21]. Additional care is required, when the citation matching process is done at scale [7]. The problem of heterogenity has been discussed in the context of datasets by [19].

Projects centered around citations or containing citation data as a core component are COCI, the "OpenCitations Index of Crossref open DOI-to-DOI citations", which was first released 2018-07-29[8] and has been regularly updated [23].

The WikiCite[9] project, "a Wikimedia initiative to develop open citations and linked bibliographic data to serve free knowledge" continously adds citations to its database[10].

Microsoft Academic Graph [27] is comprised of a number of entities[11] with *PaperReferences* being one relation among many others.

3   Dataset

We release the first version of the *refcat* dataset in an format used internally for storage and to serve queries (and which we call *biblioref* or *bref* for short). The dataset includes metadata from fatcat, the Open Library project and inbound links from the English Wikipedia. The fatcat project itself aggregates data from variety of open data sources, such as Crossref[12], PubMed[13], DataCite[14], Directory of Open Access Jourals (DOAJ)[15], dblp [17] and others, as well as metadata generated from analysis of data preserved at the Internet Archive and active crawls of publication sites on the web.

The dataset is integrated into the fatcat website and allows users to explore inbound and outbound references[16].

Set	Count
COCI (C)	1,094,394,688
refcat-doi (R)	1,303,424,212
C ∩ R	1,007,539,966
C \ R	86,854,309
R \ C	295,884,246

Table 1: Comparison between COCI and *refcat-doi*, a subset of *refcat* where entities have a known DOI. At least 50% of the 295,884,246 references only in *refcat-doi* come from links recorded within a specific dataset provider (GBIF, DOI prefix: 10.15468).

The format records source and target (fatcat release and work) identifiers, a few attributes from the metadata (such as year or release stage) as well as information about the match status and provenance.

The dataset currently contains 1,323,423,672 citations across 76,327,662 entities (55,123,635 unique source and 60,244,206 unique target work identifiers; for 1,303,424,212 - or 98.49% of all citations - we do have a DOI for both source and target). The majority of matches - 1,250,523,321 - are established through identifier based matching (DOI, PMIC, PMCID, ARXIV, ISBN). 72,900,351 citations are established through fuzzy matching techniques.

The majority of citations between *refcat* and COCI overlap, as can be seen in Table 1.

4   System Design

4.1   Constraints

The constraints for the systems design are informed by the volume and the variety of the data. The capability to run the whole graph derivation on a single machine was a minor goal as well. In total, the raw inputs amount to a few terabytes of textual content, mostly newline delimited JSON. More importantly, while the number of data fields is low, certain schemas are very partial with hundreds of different combinations of available field values found in the raw reference data. This is most likely caused by aggregators passing on reference data coming from hundreds of sources, each of which not necessarily agreeing on a common granularity for citation data and from artifacts of machine learning based structured data extraction tools.

Each combination of fields may require a slightly different processing path. For example, references with an Arxiv identifier can be processed differently from references with only a title. Over 50% of the raw reference data comes from a set of eight field set manifestations, as listed in Table 2.

Fields	Percentage
CN · CRN · P · T · U · V · Y	14%
DOI	14%
CN · CRN · IS · P · T · U · V · Y	5%
CN · CRN · **DOI** · U · V · Y	4%
PMID · U	4%
CN · CRN · **DOI** · T · V · Y	4%
CN · CRN · Y	4%
CN · CRN · **DOI** · V · Y	4%

Table 2: Top 8 combinations of available fields in raw reference data accounting for about 53% of the total data (CN = container name, CRN = contrib raw name, P = pages, T = title, U = unstructured, V = volume, IS = issue, Y = year, DOI = doi, PMID = pmid). Unstructured fields may contain any value. Identifiers emphasized.

4.2 Data Sources

Reference data comes from two main sources: explicit bibliographic metadata and PDF extraction. The bibliographic metadata is taken from fatcat, which itself harvests and imports web accessible sources such as Crossref, Pubmed, Arxiv, Datacite, DOAJ, dblp and others into its catalog (as the source permits, data is processed continously or in batches). Reference data from PDF documents has been extracted with GROBID[17], with the TEI-XML results being cached locally in a key-value store accessible with an S3 API. Archived PDF documents result from dedicated web-scale crawls of scholarly domains conducted with Heritrix[18] and a variety of seed lists targeting journal homepages, repositories, dataset providers, aggregators, web archives and other venues. A processing pipeline merges catalog data from the primary database and cached values in key-value stores and generates the set of about 2.5B references documents, which currently serve as an input for the citation graph derivation pipeline.

4.3 Methodology

Overall, a map-reduce style [5] approach is followed[19], which allows for some uniformity in the overall processing. We extract (key, document) tuples (as TSV) from the raw JSON data and sort by key. We then group documents with the same key and apply a function on each group in order to generate our target schema or perform additional

operations such as deduplication or fusion of matched and unmatched references for indexing.

The key derivation can be exact (via an identifier like DOI, PMID, etc) or based on a value normalization, like "slugifying" a title string. For identifier based matches we can generate the target schema directly. For fuzzy matching candidates, we pass possible match pairs through a verification procedure, which is implemented for *release entity*[20] pairs. This procedure is a domain dependent rule based verification, able to identify different versions of a publication, preprint-published pairs and documents, which are are similar by various metrics calculated over title and author fields. The fuzzy matching approach is applied on all reference documents without identifier (a title is currently required).

We currently implement performance sensitive parts in Go[21], with various processing stages (e.g. conversion, map, reduce, ...) represented by separate command line tools. A thin task orchestration layer using the luigi framework[22] allows for experimentation in the pipeline and for single command derivations, as data dependencies are encoded with the help of the orchestrator. Within the tasks, we also utilize classic platfrom tools such as sort [20].

With a few schema conversions, fuzzy matching can be applied to Wikipedia articles and Open Library (edition) records as well. The aspect of precision and recall are represented by the two stages: we are generous in the match candidate generation phase in order to improve recall, but we are strict during verification, in order to control precision. Quality assurance for verification is implemented through a growing list of test cases of real examples from the catalog and their expected or desired match status[23].

5 Limitations and Future Work

As other dataset in this field we expect this dataset to be iterated upon.

- The fatcat catalog updates its metadata continuously[24] and web crawls are conducted regularly. Current processing pipelines cover raw reference snapshot creation and derivation of the graph structure, which allows to rerun processing based on updated data as it becomes available.

---

[17]GROBID v0.5.5

[18]https://github.com/internetarchive/heritrix3

[19]While the operations are similar, the processing is not distributed but runs on a single machine. For space efficiency, zstd [2] is used to compress raw data and derivations.

[20]https://guide.fatcat.wiki/entity_release.html.

[21]https://golang.org/

[22]https://github.com/spotify/luigi [1], which has been used in various scientific pipeline application, like [24], [6], [15], [4] and others.

[23]The list can be found under: https://gitlab.com/internetarchive/refcat/-/blob/master/skate/testdata/verify.csv. It is helpful to keep this test suite independent of any specific programming language.

[24]A changelog can currenly be followed here: https://fatcat.wiki/changelog.

- Metadata extraction from PDFs depends on supervised machine learning models, which in turn depend on available training datasets. With additional crawls and metadata available we hope to improve models used for metadata extraction, improving yield and reducing data extraction artifacts in the process.

- As of this version, a number of raw reference docs remain unmatched, which means that neither exact nor fuzzy matching has detected a link to a known entity. On the one hand, this can hint at missing metadata. However, parts of the data will contain a reference to a catalogued entity, but in a specific, dense and harder to recover form. This also include improvements to the fuzzy matching approach.

- The reference dataset contains millions of URLs and their integration into the graph has been implemented as a prototype. A full implementation requires a few data cleanup and normalization steps.

6   Acknowledgements

7   Appendix A

A note on data quality: While we implement various data quality measures, real-world data, especially coming from many different sources will contain issues. Among other measures, we keep track of match reasons, especially for fuzzy matching to be able to zoom in on systematic errors more easily (see Table 3).

References

[1] E. Bernhardsson and E. Freider. Rouhani a. spotify/luigi-github, 2018.

[2] Y. Collet and M. Kucherawy. Zstandard compression and the application/zstd media type. *RFC 8478*, 2018.

[3] I. G. Councill, C. L. Giles, and M.-Y. Kan. Parscit: an open-source crf reference string parsing package. In *LREC*, volume 8, pages 661–667, 2008.

[4] M. Czygan. Design and implementation of a library metadata management framework and its application in fuzzy data deduplication and data reconciliation with authority data. *Informatik 2014*, 2014.

[5] J. Dean and S. Ghemawat. Mapreduce: a flexible data processing tool. *Communications of the ACM*, 53(1):72–77, 2010.

Count	Provenance	Status	Reason
934932865	crossref	exact	doi
151366108	fatcat-datacite	exact	doi
65345275	fatcat-pubmed	exact	pmid
48778607	fuzzy	strong	jaccardauthors
42465250	grobid	exact	doi
29197902	fatcat-pubmed	exact	doi
19996327	fatcat-crossref	exact	doi
11996694	fuzzy	strong	slugtitleauthormatch
9157498	fuzzy	strong	tokenizedauthors
3547594	grobid	exact	arxiv
2310025	fuzzy	exact	titleauthormatch
1496515	grobid	exact	pmid
680722	crossref	strong	jaccardauthors
476331	fuzzy	strong	versioneddoi
449271	grobid	exact	isbn
230645	fatcat-crossref	strong	jaccardauthors
190578	grobid	strong	jaccardauthors
156657	crossref	exact	isbn
123681	fatcat-pubmed	strong	jaccardauthors
79328	crossref	exact	arxiv
57414	crossref	strong	tokenizedauthors
53480	fuzzy	strong	pmiddoipair
52453	fuzzy	strong	dataciterelatedid
47119	grobid	strong	slugtitleauthormatch
36774	fuzzy	strong	arxivversion

Table 3: Table of match counts (top 25), reference provenance, match status and match reason. Provenance currently can name the raw origin (e.g. *crossref*) or the method (e.g. *fuzzy*). The match reason identifier encode a specific rule in the domain dependent verification process and are included for completeness - we do not include the details of each rule in this report.

[6] M. Erdmann, B. Fischer, R. Fischer, and M. Rieger. Design and execution of make-like, distributed analyses based on spotify's pipelining package luigi. In *Journal of Physics: Conference Series*, volume 898, page 072047. IOP Publishing, 2017.

[7] M. Fedoryszak, D. Tkaczyk, and Ł. Bolikowski. Large scale citation matching using apache hadoop. In *International Conference on Theory and Practice of Digital Libraries*, pages 362–365. Springer, 2013.

[8] E. Garfield. The evolution of the science citation index. *International microbiology*, 10(1):65, 2007.

[9] C. L. Giles, K. D. Bollacker, and S. Lawrence. Citeseer: An automatic citation indexing system. In *Proceedings of the third ACM conference on Digital libraries*, pages 89–98, 1998.

[10] A. Hosseini, B. Ghavimi, Z. Boukhers, and P. Mayr. Excite–a toolchain to extract, match and publish open literature references. In *2019 ACM/IEEE Joint Conference on Digital Libraries (JCDL)*, pages 432–433. IEEE, 2019.

[11] B. I. Hutchins. A tipping point for open citation data. *Quantitative Science Studies*, pages 1–5, 2021.

[12] Jinha. Article 50 million: an estimate of the number of scholarly articles in existence. 23, Jul 2010.

[13] Khabsa and Giles. The number of scholarly documents on the public web. May 2014.

[14] J. Lafferty, A. McCallum, and F. C. Pereira. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. 2001.

[15] S. Lampa, M. Dahlö, J. Alvarsson, and O. Spjuth. Scipipe: A workflow library for agile development of complex and dynamic bioinformatics pipelines. *GigaScience*, 8(5):giz044, 2019.

[16] Landhuis. Scientific literature: Information overload. 535(7612), Jul 2016.

[17] M. Ley. The dblp computer science bibliography: Evolution, research issues, perspectives. In *International symposium on string processing and information retrieval*, pages 1–10. Springer, 2002.

[18] P. Lopez. Grobid: Combining automatic bibliographic data recognition and term extraction for scholarship publications. In *International conference on theory and practice of digital libraries*, pages 473–474. Springer, 2009.

[19] B. Mathiak and K. Boland. Challenges in matching dataset citation strings to datasets in social science. *D-Lib Magazine*, 21(1/2):23–28, 2015.

[20] M. D. McIlroy. A research unix reader: annotated excerpts from the programmer's manual. 1971.

[21] M. Olensky, M. Schmidt, and N. J. van Eck. Evaluation of the citation matching algorithms of cwts and i fq in comparison to the w eb of science. *Journal of the Association for Information Science and Technology*, 67(10):2550–2564, 2016.

[22] H. Pasula, B. Marthi, B. Milch, S. J. Russell, and I. Shpitser. Identity uncertainty and citation matching. In *Advances in neural information processing systems*, pages 1425–1432, 2003.

[23] S. Peroni and D. Shotton. Opencitations, an infrastructure organization for open scholarship. *Quantitative Science Studies*, 1(1):428–444, 2020.

[24] W. L. Schulz, T. J. Durant, A. J. Siddon, and R. Torres. Use of application containers and workflows for genomic data analysis. *Journal of pathology informatics*, 7, 2016.

[25] D. Shotton. Publishing: open citations. *Nature News*, 502(7471):295, 2013.

[26] D. Shotton. Funders should mandate open citations. *Nature*, 553(7686):129–130, 2018.

[27] A. Sinha, Z. Shen, Y. Song, H. Ma, D. Eide, B.-J. Hsu, and K. Wang. An overview of microsoft academic service (mas) and applications. In *Proceedings of the 24th international conference on world wide web*, pages 243–246, 2015.

[28] D. Tkaczyk, P. Szostek, P. J. Dendek, M. Fedoryszak, and L. Bolikowski. Cermine–automatic extraction of metadata and references from scientific literature. In *2014 11th IAPR International Workshop on Document Analysis Systems*, pages 217–221. IEEE, 2014.

[29] J. Wu, K. Kim, and C. L. Giles. Citeseerx: 20 years of service to scholarly big data. In *Proceedings of the Conference on Artificial Intelligence for Data Discovery and Reuse*, pages 1–4, 2019.