

Preprint of the paper:

Meuschke, N. & Jagdale, A. & Spinde, T. & Mitrović, J. & Gipp, B., "A Benchmark of PDF Information Extraction Tools Using a Multi-task and Multi-domain Evaluation Framework for Academic Documents", in Information for a Better World: Normality, Virtuality, Physicality, Inclusivity, LNCS, vol. 13972, Cham: Springer Nature Switzerland, 2023, pp. 383–405, DOI: [10.1007/978-3-031-28032-0_31](https://doi.org/10.1007/978-3-031-28032-0_31).

Click to download: [BibTeX](#)

A Benchmark of PDF Information Extraction Tools using a Multi-Task and Multi-Domain Evaluation Framework for Academic Documents

Norman Meuschke¹,[\[OrcID\]](#), Apurva Jagdale², Timo Spinde¹,[\[OrcID\]](#), Jelena Mitrović^{2,3},[\[OrcID\]](#), and Bela Gipp¹,[\[OrcID\]](#)

¹ University of Göttingen, 37073 Göttingen, Germany
{meuschke, spinde, gipp}@uni-goettingen.de

² University of Passau, 94032 Passau, Germany
{apurva.jagdale, jelena.mitrovic}@uni-passau.de

³ The Institute for Artificial Intelligence R&D of Serbia, 21000 Novi Sad, Serbia

Abstract. Extracting information from academic PDF documents is crucial for numerous indexing, retrieval, and analysis use cases. Choosing the best tool to extract specific content elements is difficult because many, technically diverse tools are available, but recent performance benchmarks are rare. Moreover, such benchmarks typically cover only a few content elements like header metadata or bibliographic references and use smaller datasets from specific academic disciplines. We provide a large and diverse evaluation framework that supports more extraction tasks than most related datasets. Our framework builds upon DocBank, a multi-domain dataset of 1.5M annotated content elements extracted from 500K pages of research papers on arXiv. Using the new framework, we benchmark ten freely available tools in extracting document metadata, bibliographic references, tables, and other content elements from academic PDF documents. GROBID achieves the best metadata and reference extraction results, followed by CERMINÉ and Science Parse. For table extraction, Adobe Extract outperforms other tools, even though the performance is much lower than for other content elements. All tools struggle to extract lists, footers, and equations. We conclude that more research on improving and combining tools is necessary to achieve satisfactory extraction quality for most content elements. Evaluation datasets and frameworks like the one we present support this line of research. We make our data and code publicly available to contribute toward this goal.

Keywords: PDF · Information Extraction · Benchmark · Evaluation.

1 Introduction

The Portable Document Format (PDF) is the most prevalent encoding for academic documents. Extracting information from academic PDF documents is crucial for numerous indexing, retrieval, and analysis tasks. Document search, recommendation, summarization, classification, knowledge base construction, question answering, and bibliometric analysis are just a few examples [31].

However, the format’s technical design makes information extraction challenging. Adobe designed PDF as a platform-independent, fixed-layout format by extending the PostScript [24] page description language. PDF focuses on encoding a document’s visual layout to ensure a consistent appearance of the document across software and hardware platforms but includes little structural and semantic information on document elements.

Numerous tools for information extraction (IE) from PDF documents have been presented since the format’s inception in 1993. The development of such tools has been subject to a fast-paced technological evolution of extraction approaches from rule-based algorithms, over statistical machine learning (ML) to deep learning (DL) models (cf. Section 2). Finding the best tool to extract specific content elements from PDF documents is currently difficult because:

1. Typically, tools only support extracting a subset of the content elements in academic documents, e.g., title, authors, paragraphs, in-text citations, captions, tables, figures, equations, or references.
2. Many information extraction tools, e.g., 12 of 35 tools we considered for our study, are no longer maintained or have become obsolete.
3. Prior evaluations of information extraction tools often consider only specific content elements or use domain-specific corpora, which makes their results difficult to compare. Moreover, the most recent comprehensive benchmarks of information extraction tools were published in 2015 for metadata⁴ [55], 2017 for body text [6], and 2018 for references⁵ [54], respectively. These evaluations do not reflect the latest technological advances in the field.

To alleviate this knowledge gap and facilitate finding the best tool to extract specific elements from academic PDF documents, we comprehensively evaluate ten state-of-the-art non-commercial tools that consider eleven content elements based on a dataset of 500K pages from arXiv documents covering multiple fields.

Our code, data, and resources are publicly available at
<http://pdf-benchmark.gipplab.org>

2 Related Work

This section presents approaches for information extraction from PDF (Section 2.1), labeled datasets suitable for training and evaluating PDF information extraction approaches, and prior evaluations of IE tools (Section 2.2).

2.1 Information Extraction from PDF Documents

Table 1 summarizes publications on PDF information extraction since 1999. For each publication, the table shows the primary technological approach and the

⁴ For example author(s), title, affiliation(s), address(es), email(s)

⁵ Refers to extracting the components of bibliographic references, e.g., author(s), title, venue, editor(s), volume, issue, page range, year of publication, etc.

Table 1: Publications on information extraction from PDF documents.

Publication ¹	Year	Task ²	Method	Training Dataset ³
Palermo [44]	1999	M, ToC	Rules	100 documents
Klink [27]	2000	M	Rules	979 pages
Giuffrida [18]	2000	M	Rules	1,000 documents
Aiello [2]	2002	RO, Title	Rules	1,000 pages
Mao [37]	2004	M	OCR, Rules	309 documents
Peng [45]	2004	M, R	CRF	CORA (500 refs.)
Day [14]	2007	M, R	Template	160,000 citations
Hetzner [23]	2008	R	HMM	CORA (500 refs.)
Councill [12]	2008	R	CRF	CORA (200 refs.), CiteSeer (200 refs.)
Lopez [36]	2009	B, M, R	CRF, DL	None
Cui [13]	2010	M	HMM	400 documents
Ojokoh [42]	2010	M	HMM	CORA (500 refs.), ManCreat FLUX-CiM (300 refs.),
Kern [25]	2012	M	HMM	E-prints, Mendeley, PubMed (19K entries)
Bast [5]	2013	B, M, R	Rules	DBLP (690 docs.), PubMed (500 docs.)
Souza [53]	2014	M	CRF	100 documents
Anzaroot [3]	2014	R	CRF	UMASS (1,800 refs.)
Vilnis [57]	2015	R	CRF	UMASS (1,800 refs.)
Tkaczyk [55]	2015	B, M, R	CRF, Rules, SVM	CiteSeer (4,000 refs.), CORA (500 refs.), GROTOAP, PMC (53K docs.)
Bhardwaj [7]	2017	R	FCN	5,090 references
Rodrigues [49]	2018	R	BiLSTM	40,000 references
Prasad [46]	2018	M, R	CRF, DL	FLUX-CiM (300 refs.), CiteSeer (4,000 refs.)
Jahongir [4]	2018	M	Rules	10,000 documents
Torre [15]	2018	B, M	Rules	300 documents
Rizvi [47]	2020	R	R-CNN	40,000 references
Hashmi [22]	2020	M	Rules	45 documents
Ahmed [1]	2020	M	Rules	150 documents
Nikolaos [33]	2021	B, M, R	Attention, BiLSTM	3,000 documents

¹ Publications in chronological order; the labels indicate the first author only.² (B) Body text, (M) Metadata, (R) References, (RO) Reading order, (ToC) Table of contents³ Domain-specific datasets: Computer Science: CiteSeer [43], CORA [39], DBLP [52], FLUX-CiM [10,11], ManCreat [42]; Health Science: PubMed [40], PMC [41]

training dataset. Eighteen of 27 approaches (67%) employ machine learning or deep learning (DL) techniques, and the remainder rule-based extraction (Rules). Early tools rely on manually coded rules [44]. Second-generation tools use statistical machine learning, e.g., based on Hidden Markov Models (HMM) [8], Conditional Random Fields (CRF) [29], and maximum entropy [26]. The most recent information extraction tools employ Transformer models [56].

A preference for—in theory—more flexible and adaptive machine learning and deep learning techniques over case-specific rule-based algorithms is observable in Table 1. However, many training datasets are domain-specific, e.g., they exclusively consist of documents from Computer Science or Health Science, and comprise fewer than 500 documents. These two factors put the generalizability of the respective IE approaches into question. Notable exceptions like Ojokoh et al. [42], Kern et al. [25], and Tkaczyk et al. [55] use multiple datasets covering different domains for training and evaluation. However, these approaches address specific tasks, i.e. header metadata extraction, reference extraction, or both.

Moreover, a literature survey by Mao et al. shows that most approaches for text extraction from PDF do not specify the ground-truth data and performance metrics they use, which impedes performance comparisons [38]. A positive exception is a publication by Bast et al. [5], which presents a comprehensive evaluation framework for text extraction from PDF that includes a fine-grained specification of the performance measures used.

2.2 Labeled Datasets and Prior Benchmarks

Table 2 summarizes datasets usable for training and evaluating PDF information extraction approaches grouped by the type of ground-truth labels they offer. Most datasets exclusively offer labels for document metadata, references, or both.

Table 2: Labeled datasets for information extraction from PDF documents.

Publication ¹	Size	Ground-truth Labels
Fan [16]	147 documents	Metadata
Färber [17]	90K documents	References
Grennan [21]	1B references	References
Saier [51,50]	1M documents.	References
Ley [30,52]	6M documents	Metadata, references
Mccallum [39]	935 documents	Metadata, references
Kyle [34]	8.1M documents	Metadata, references
Ororbial [43]	6M documents.	Metadata, references
Bast [6]	12,098 documents	Body text, sections, title
Li [31]	500K pages	Captions, equations, figures, footers lists, metadata, paragraphs, references, sections, tables

¹ The labels indicate the first author only.

Only the DocBank dataset by Li et al. [31] offers annotations for 12 diverse content elements in academic documents, including, figures, equations, tables, and captions. Most of these content elements have not been used for benchmark evaluations yet. DocBank is comparably large (500K pages from research papers published on arXiv in a four-year period). A downside of the DocBank dataset is its coarse-grained labels for references, which do not annotate the fields of bibliographic entries like the author, publisher, volume, or date, as do bibliography-specific datasets like unarXive [21] or S2ORC [34].

Table 3 shows PDF information extraction benchmarks performed since 1999. Few such works exist and were rarely repeated or updated, which is sub-optimal given that many tools receive updates frequently. Other tools become technologically obsolete or unmaintained. For instance, pdf-extract⁶, lapdftext⁷, PDF-SSA4MET⁸, and PDFMeat⁹ are no longer maintained actively, while ParsCit¹⁰ has been replaced by NeuralParsCit¹¹ and SciWING¹².

Table 3: Benchmark evaluations of PDF information extraction approaches.

Publication ¹	Dataset	Metrics ²	Tools	Labels ³
Granitzer [19]	E-prints (2,452 docs.), Mendeley (20,672 docs.)	P, R	2	M
Lipinski [32]	arXiv (1,253 docs.)	Acc	7	M
Bast [6]	arXiv (12,098 docs.)	Custom	14	NL, Pa RO, W
Körner [28]	100 (German docs.)	P, R, F_1	4	Ref
Tkaczyk [54]	9,491 documents	P, R, F_1	10	Ref
Rizvi [48]	8,766 references	F_1	4	Ref

¹ The labels indicate the first author only.

² (P) Precision, (R) Recall, (F_1) F_1 -score, (Acc) Accuracy

³ (M) Metadata, (NL) New Line, (Pa) Paragraph, (Ref) Reference, (RO) Reading order, (W) Words

As Table 3 shows, the most extensive dataset used for evaluating PDF information extraction tools so far contains approx. 24,000 documents. This number is small compared to the sizes of datasets available for this task, shown in Table 2. Most studies focused on exclusively evaluating metadata and reference extraction (see also Table 3). An exception is a benchmark by Bast and Korzen

⁶ <https://github.com/CrossRef/pdfextract>

⁷ <https://github.com/BMKEG/lapdftext>

⁸ <https://github.com/eliask/pdfssa4met>

⁹ <https://github.com/dimatura/pdfmeat>

¹⁰ <https://github.com/knmnyn/ParsCit>

¹¹ <https://github.com/WING-NUS/Neural-ParsCit>

¹² <https://github.com/abhinavkashyap/sciwing>

[6], which evaluated spurious and missing words, paragraphs, and new lines for 14 tools but used a comparably small dataset of approx. 10K documents.

We conclude from our review of related work that (1) recent benchmarks of information extraction tools for PDF are rare, (2) mostly analyze metadata extraction, (3) use small, domain-specific datasets, and (4) include tools that have become obsolete or unmaintained. (5) A variety of suitably labeled datasets have not been used to evaluate information extraction tools for PDF documents yet. Therefore, we see the need for benchmarking state-of-the-art PDF information extraction tools on a large labeled dataset of academic documents covering multiple domains and containing diverse content elements.

3 Methodology

This section presents the experimental setup of our study by describing the tools we evaluate (Section 3.1), the dataset we use (Section 3.2), and the procedure we follow (Section 3.3).

3.1 Evaluated Tools

We chose ten actively maintained non-commercial open-source tools that we categorize by extraction tasks.

1. **Metadata Extraction** includes tools to extract titles, authors, abstracts, and similar document metadata.
2. **Reference Extraction** comprises tools to access and parse bibliographic reference strings into fields like author names, publication titles, and venue.
3. **Table Extraction** refers to tools that allow accessing both the structure and data of tables.
4. **General Extraction** subsumes tools to extract, e.g., paragraphs, sections, figures, captions, equations, lists, or footers.

For each of the tools we evaluate, Table 4 shows the version, supported extraction task(s), primary technological approach, and output format. Hereafter, we briefly describe each tool, focusing on its technological approach.

Adobe Extract¹³ is a cloud-based API that allows extracting tables and numerous other content elements subsumed in the *general extraction* category. The API employs the Adobe Sensei¹⁴ AI and machine learning platform to understand the structure of PDF documents. To evaluate the Adobe Extract API, we used the Adobe PDFServices Python SDK¹⁵ to access the API’s services.

Apache Tika¹⁶ allows metadata and content extraction in XML format. We used the tika-python¹⁷ client to access the Tika REST API. Unfortunately, we found that tika-python only supports content (paragraphs) extraction.

¹³ <https://www.adobe.io/apis/documentcloud/dcsdk/pdf-extract.html>

¹⁴ <https://www.adobe.com/de/sensei.html>

¹⁵ <https://github.com/adobe/pdfservices-python-sdk-samples>

¹⁶ <https://tika.apache.org/>

¹⁷ <https://github.com/chrisattmann/tika-python>

Table 4: Overview of evaluated information extraction tools.

Tool	Version	Task ¹	Technology	Output
Adobe Extract	1.0	G, T	Adobe Sensei AI Framework	JSON, XLSX
Apache Tika	2.0.0	G	Apache PDFBox	TXT
Camelot	0.10.1	T	OpenCV, PDFMiner	CSV, Dataframe
CERMINE	1.13	G, M, R	CRF, iText, Rules, SVM	JATS
GROBID	0.7.0	G, M, R, T	CRF, Deep Learning, Pdftalto	TEI XML
PdfAct	n/a	G, M, R, T	pdftotext, rules	JSON, TXT, XML
PyMuPDF	1.19.1	G	OCR, tesseract	TXT
RefExtract	0.2.5	R	pdftotext, rules	TXT
ScienceParse	1.0	G, M, R,	CRF, pdffigures2, rules	JSON
Tabula	1.2.1	T	PDFBox, rules	CSV, Dataframe

¹ (G) General, (M) Metadata, (R) References, (T) Table

Camelot¹⁸ can extract tables using either the *Stream* or *Lattice* modes. The former uses whitespace between cells and the latter table borders for table cell identification. For our experiments, we exclusively use the Stream mode, since our test documents are academic papers, in which tables typically use whitespace in favor of cell borders to delineate cells. The Stream mode internally utilizes the PDFMiner library¹⁹ to extract characters that are subsequently grouped into words and sentences using whitespace margins.

CERMINE [55] offers metadata, reference, and general extraction capabilities. The tool employs the iText PDF toolkit²⁰ for character extraction and the Docstrum²¹ image segmentation algorithm for page segmentation of document images. CERMINE uses an SVM classifier implemented using the LibSVM²² library and rule-based algorithms for metadata extraction. For reference extraction, the tool employs k -means clustering, and Conditional Random Fields implemented using the MALLETT²³ toolkit for sequence labeling. CERMINE returns a single XML file containing the annotations for an entire PDF. We employ the Beautiful Soup²⁴ library to filter CERMINE’s output files for the annotations relevant to our evaluation.

GROBID²⁵ [35] supports all four extraction tasks. The tool allows using either feature-engineered CRF (default) or a combination of CRF and DL models realized using the DeLFT²⁶ Deep Learning library, which is based on TensorFlow and Keras. GROBID uses a cascade of sequence labeling models for different components. The models in the model cascade use individual label sequencing

¹⁸ <https://github.com/camelot-dev/camelot>

¹⁹ <https://github.com/pdfminer/pdfminer.six>

²⁰ <https://github.com/itext>

²¹ <https://github.com/chulwoopack/docstrum>

²² <https://github.com/cjlin1/libsvm>

²³ <http://mallet.cs.umass.edu/sequences.php>

²⁴ <https://www.crummy.com/software/BeautifulSoup/bs4/doc/>

²⁵ <https://github.com/kermitt2/grobid>

²⁶ <https://github.com/kermitt2/delft>

algorithms and features; some models employ tokenizers. This approach offers flexibility by allowing model tuning and improves the model’s maintainability. We evaluate the default CRF model with production settings (a recommended setting to improve the performance and availability of the GROBID server, according to the tool’s documentation²⁷).

PdfAct formerly called Icecite [5] is a rule-based tool that supports all four extraction tasks, including the extraction of appendices, acknowledgments, and tables of contents. The tool uses the PDFBox²⁸ and pdftotext²⁹ PDF manipulation and content extraction libraries. We use the tool’s JAR release³⁰.

PyMuPDF³¹ extends the MuPDF³² viewer library with font and image extraction, PDF joining, and file embedding. PyMuPDF uses tesseract³³ for OCR. PyMuPDF could not process files whose names include special characters.

RefExtract³⁴ is a reference extraction tool that uses pdftotext³⁵ and regular expressions. RefExtract returns annotations for the entire bibliography of a document. The ground-truth annotations in our dataset (cf. Section 3.2), however, pertain to individual pages of documents and do not always cover the entire document. If ground-truth annotations are only available for a subset of the references in a document, we use regular expressions to filter RefExtract’s output to those references with ground-truth labels.

Science Parse³⁶ uses a CRF model trained on data from GROBID to extract the title, author, and references. It also employs a rule-based algorithm by Clark and Divvala [9] to extract sections and paragraphs in JSON format.

Tabula³⁷ is a table extraction tool. Analogous to Camelot, Tabula offers a *Stream* mode realized using PDFBox, and a *Lattice* mode realized using OpenCV for table cell recognition.

3.2 Dataset

We use the DocBank³⁸ dataset, created by Li et al. [31], for our experiments. Figure 1 visualizes the process for compiling the dataset. First, the creators gathered arXiv documents, for which both the PDF and LaTeX source code was available. Li et al. then edited the LaTeX code to enable accurate automated annotations of content elements in the PDF version of the documents. For this purpose, they inserted commands that formatted content elements in specific

²⁷ <https://GROBID.readthedocs.io/en/latest/Troubleshooting/>

²⁸ <http://pdfbox.apache.org/>

²⁹ <https://github.com/jalan/pdftotext>

³⁰ <https://github.com/ad-freiburg/pdfact>

³¹ <https://github.com/pymupdf/PyMuPDF>

³² <https://mupdf.com/>

³³ <https://github.com/tesseract-ocr/tesseract>

³⁴ <https://github.com/inspirehep/refextract>

³⁵ <https://linux.die.net/man1/pdftotext>

³⁶ <https://github.com/allenai/science-parse>

³⁷ <https://github.com/chezou/tabula-py>

³⁸ <https://github.com/doc-analysis/DocBank>




Fig. 1: Process for generating the DocBank dataset.

colors. The center part of Figure 3 shows the mapping of content elements to colors. In the last step, the dataset creators used PDFPlumber³⁹ and PDFMiner to extract and annotate relevant content elements by their color. DocBank provides the annotations as separate files for each document page in the dataset.

Table 5 shows the structure of the tab-separated ground-truth files. Each line in the file refers to one component on the page and is structured as follows. Index 0 represents the token itself, e.g., a word. Indices 1-4 denote the bounding box information of the token, where (x_0, y_0) represents the top-left and (x_1, y_1) the bottom-right corner of the token in the PDF coordinate space. Indices 5-7 reflect the token’s color in RGB notation, index 8 the token’s font, and index 9 the label for the type of the content element. Each ground-truth file adheres to the naming scheme shown in Figure 2.

Table 5: Structure of DocBank’s plaintext ground-truth files.

Index	0	1	2	3	4	5	6	7	8	9
Content	token	x0	y0	x1	y1	R	G	B	font name	label

Source: <https://doc-analysis.github.io/docbank-page/index.html>.

Prefix Identifier Postfix Page
8.tar 1501.04311.gz_pippori 27.txt

Fig. 2: Naming scheme for DocBank’s ground-truth files.

³⁹ <https://github.com/jsvine/pdfplumber>

The DocBank dataset offers ground-truth annotations for 1.5M content elements on 500K pages. Li et al. extracted the pages from arXiv papers in Physics, Mathematics, Computer Science, and numerous other fields published between 2014 and 2018. DocBank’s large size, recency, diversity of included documents, number of annotated content elements, and high annotation quality due to the weakly supervised labeling approach make it an ideal choice for our purposes.

3.3 Evaluation Procedure

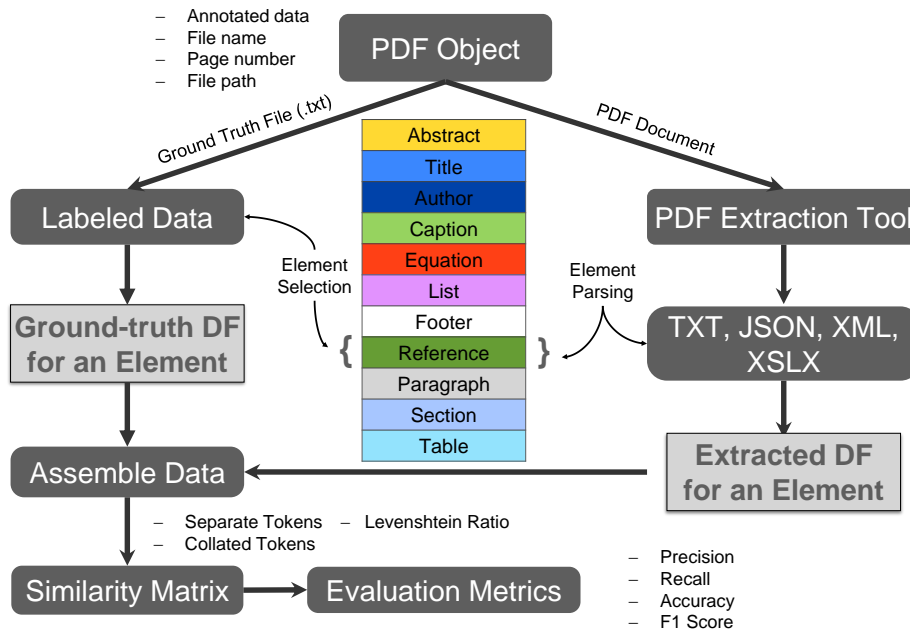


Fig. 3: Overview of the procedure for comparing content elements extracted by IE tools to the ground-truth annotations and computing evaluation metrics.

Figure 3 shows our evaluation procedure. First, we select the PDF files whose associated ground-truth files contain relevant labels. For example, we search for ground-truth files containing *reference* tokens to evaluate reference extraction tools. We include the PDF file, the ground-truth file, the document ID and page number obtainable from the file name (cf. Figure 2), and the file path in a self-defined Python object (see *PDF Object* in Figure 3).

Then, the evaluation process splits into two branches whose goal is to create two pandas data frames—one holding the relevant ground-truth data, and the other the output of an information extraction tool. For this purpose, both the ground-truth files and the output files of IE tools are parsed and filtered for

the relevant content elements. For example, to evaluate reference extraction via CERMINE, we exclusively parse reference tags from CERMINE’s XML output file into a data frame (see *Extracted DF* in Figure 3).

Finally, we convert both the *ground-truth data frame* and the *extracted data frame* into two formats for comparison and computing performance metrics. The first is the *separate tokens* format, in which every token is represented as a row in the data frame. The second is the *collated tokens* format, in which all tokens are combined into a single space-delimited row in the data frame. Separate tokens serve to compute a strict score for token-level extraction quality, whereas collated tokens yield a more lenient score intended to reflect a tool’s average extraction quality for a class of content elements. We will explain the idea of both scores and their computation hereafter.

We employ the *Levenshtein Ratio* to quantify the similarity of extracted tokens and the ground-truth data for both the separate tokens and collated tokens format. Equation (1) defines the computation of the *Levenshtein distance* of the extracted tokens t_e and the ground-truth tokens t_g .

$$lev_{t_e, t_g}(i, j) = \begin{cases} \max(i, j), & \text{if } \min(i, j) = 0, \\ \min \begin{cases} lev_{t_e, t_g}(i-1, j) + 1 \\ lev_{t_e, t_g}(i, j-1) + 1 \\ lev_{t_e, t_g}(i-1, j-1) + 1_{(t_{e_i} \neq t_{g_j})} \end{cases} & \text{otherwise.} \end{cases} \quad (1)$$

Equation (2) defines the derived Levenshtein Ratio score (γ).

$$\gamma(t_e, t_g) = 1 - \frac{lev_{t_e, t_g}(i, j)}{|t_e| + |t_g|} \quad (2)$$

Equation (3) shows the derivation of the *similarity matrix* (Δ^d) for a document (d), which contains the Levenshtein Ratio (γ) of every token in the extracted data frame with separate tokens E^s of size m and the ground-truth data frame with separate tokens G^s of size n .

$$\Delta_{m \times n}^d = \gamma [E_i^s, G_j^s]_{i, j}^{m, n} \quad (3)$$

Using the $m \times n$ similarity matrix, we compute the *Precision* P^d and *Recall* R^d scores according to Equation (4) and Equation (5), respectively. As the numerator, we use the number of extracted tokens whose Levenshtein Ratio is larger or equal to 0.7. We chose this threshold for consistency with the experiments by Granitzer et al. [19]. We then compute the F_1^d score according to Equation (6) as a token-level score for a tool’s extraction quality.

$$P^d = \frac{\#\Delta_{i, j}^d \geq 0.7}{m} \quad (4)$$

$$R^d = \frac{\#\Delta_{i, j}^d \geq 0.7}{n} \quad (5)$$

$$F_1^d = \frac{2 \times P^d \times R^d}{P^d + R^d} \quad (6)$$

Moreover, we compute the *Accuracy* score A^d reflecting a tool’s average extraction quality for a class of tokens. To obtain A^d , we compute the Levenshtein Ratio γ of the extracted tokens E^c and ground-truth tokens G^c in the collated tokens format, according to Equation (7).

$$A^d = \gamma[E^c, G^c] \quad (7)$$

Figure 4 and Figure 5 show the similarity matrices for the author names ‘Yuta,’ ‘Hamada,’ ‘Gary,’ and ‘Shiu’ using separate and collated tokens, respectively. Figure 4 additionally shows an example computation of the Levenshtein Ratio for the strings *Gary* and *Yuta*. The strings have a Levenshtein distance of six and a cumulative string length of eight, which results in a Levenshtein Ratio of 0.25 that is entered into the similarity matrix. Figure 5 analogously exemplifies computing the Accuracy score of the two strings using collated tokens.

	Yuta	Hamada	Gary	Shiu1,
Yuta	1.0	0.2	0.25	0.2
Hamada	0.2	1.0	0.2	0.0
Gary	0.25	0.2	1.0	0.0
Shiu	0.25	0.0	0.0	0.8

		Y	u	t	a
	0	1	2	3	4
G	1	2	3	4	5
a	2	3	4	5	4
r	3	4	5	6	5
y	4	3	4	5	6

Fig. 4: Left: Similarity matrix for author names using separate tokens. Right: Computation of the Levenshtein distance (6) and the optimal edit transcript (yellow highlights) for two author names using dynamic programming.

	Yuta Hamada Gary Shiu1,
Yuta Hamada Gary Shiu	0.957

Fig. 5: Similarity matrix for two sets of author names using collated tokens.

4 Results

We present the evaluation results grouped by extraction task (see Figures 6–9) and by tools (see Table 6). This two-fold breakdown of the results facilitates identifying the best-performing tool for a specific extraction task or content element and allows for gauging the strengths and weaknesses of tools more easily. Note that the task-specific result visualizations (Figures 6–9) only include tools that support the respective extraction task. See Table 4 for an overview of the evaluated tools and the extraction tasks they support.

Figure 6 shows the cumulative F_1 scores of CERMINE, GROBID, PdfAct, and Science Parse for the metadata extraction task, i.e., extracting title, abstract, and authors. Consequently, the best possible cumulative F_1 score equals three. Overall, GROBID performs best, achieving a cumulative F_1 score of 2.25 and individual F_1 scores of 0.91 for *title*, 0.82 for *abstract*, and 0.52 for *authors*. Science Parse (2.03) and CERMINE (1.97) obtain comparable cumulative F_1 scores, while PdfAct has the lowest cumulative F_1 score of 1.14. However, PdfAct performs second-best for title extraction with a F_1 score of 0.85. The performance of all tools is worse for extracting authors than for titles and abstracts. It appears that machine-learning-based approaches like those of CERMINE, GROBID, and Science Parse perform better for metadata extraction than rule-based algorithms like the one implemented in PdfAct⁴⁰.

Figure 7 shows the results for the reference extraction task. With a F_1 score of 0.79, GROBID also performs best for this task. CERMINE achieves the second rank with a F_1 score of 0.74, while Science Parse and RefExtract share the third rank with identical F_1 scores of 0.49. As for the metadata extraction task, PdfAct also achieves the lowest F_1 score of 0.15 for reference extraction. While both RefExtract and PdfAct employ pdftotext and regular expressions, GROBID performs efficient segregation of cascaded sequence labeling models⁴¹ for diverse components, which can be the reason for its superior performance [36].

Figure 8 depicts the results for the table extraction task. Adobe Extract outperforms the other tools with a F_1 score of 0.47. Camelot ($F_1 = 0.30$), Tabula ($F_1 = 0.28$), and GROBID ($F_1 = 0.23$) perform notably worse than Adobe Extract. Both Camelot and Tabula incorrectly treat two-column articles as tables and table captions as a part of the table region, which negatively affects their performance scores. The use of comparable *Stream* and *Lattice* modes in Camelot and Tabula (cf. Section 3.1) likely cause the tools’ highly similar results. PdfAct did not produce an output for any of our test documents that contain tables, although the tool supposedly supports table extraction. The performance of all tools is significantly lower for table extraction than for other content elements, which is likely caused by the need to extract additional structural information. The difficulty of table extraction is also reflected by numerous issues that users opened on the matter in the GROBID GitHub repository⁴².

⁴⁰ See Table 4 for more information on the tools’ extraction approaches.

⁴¹ <https://grobid.readthedocs.io/en/latest/Principles/>

⁴² <https://github.com/kermitt2/grobid/issues/340>




Fig. 6: Results for metadata extraction.




Fig. 7: Results for reference extraction.

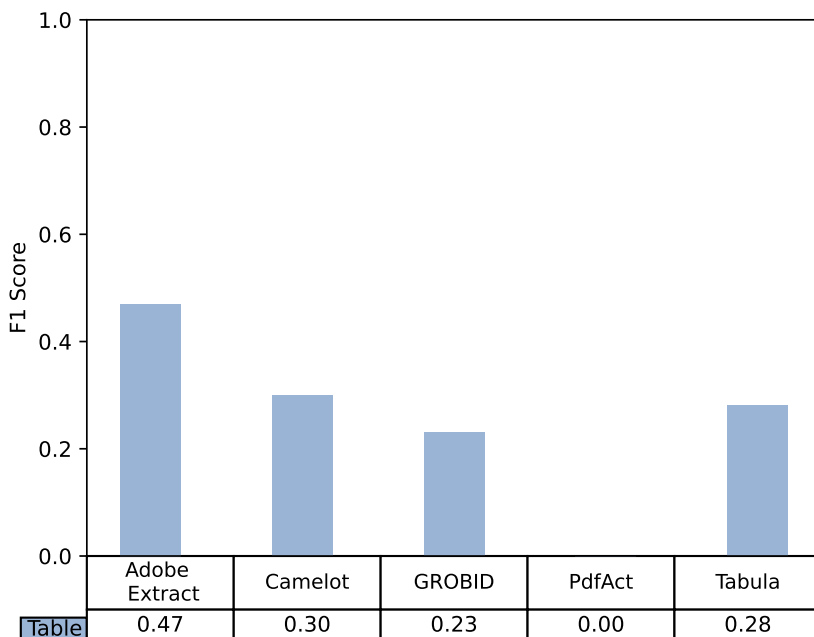


Fig. 8: Results for table extraction.

Figure 9 visualizes the results for the general extraction task. GROBID achieves the highest cumulative F_1 score of 2.38, followed by PdfAct (cumulative $F_1 = 1.66$). The cumulative F_1 scores of Science Parse (1.25), which only support paragraph and section extraction, and CERMINE (1.20) are much lower than GROBID’s score and comparable to that of PdfAct. Apache Tika, PyMuPDF, and Adobe Extract can only extract paragraphs.

For paragraph extraction, GROBID (0.9), CERMINE (0.85), and PdfAct (0.85) obtained high F_1 scores with Science Parse (0.76) and Adobe Extract (0.74) following closely. Apache Tika (0.52) and PyMuPDF (0.51) achieved notably lower scores because the tools include other elements like sections, captions, lists, footers, and equations in paragraphs.

Notably, only GROBID achieves a promising F_1 score of 0.74 for the extraction of sections. GROBID and PdfAct are the only tools that can partially extract captions. None of the tools is able to extract lists. Only PdfAct supports the extraction of footers but achieves a low F_1 score of 0.20. Only GROBID supports equation extraction but the extraction quality is comparatively low ($F_1 = 0.25$). To reduce the evaluation effort, we first tested the extraction of lists, footers, and equations on a two-months sample of the data covering January and February 2014. If a tool consistently obtained performance scores of 0, we did not continue with its evaluation. Following this procedure, we only evaluated GROBID and PdfAct on the full dataset.

For the general extraction task, GROBID outperforms other tools due to its segmentation model⁴³, which detects the main areas of documents based on layout features. Therefore, frequent content elements like paragraphs will not impact the extraction of rare elements from a non-body area by keeping the imbalanced classes in separate models. The cascading models used in GROBID also offer the flexibility to tune each model. Using layouts and structures as a basis for the process allows the association of simpler training data.




Fig. 9: Results for general data extraction.

The breakdown of results by tools shown in Table 6 underscores the main takeaway point of the results' presentation for the individual extraction tasks. The tools' results differ greatly for different content elements. Certainly, no tool performs best for all elements, rather, even tools that perform well overall can fail completely for certain extraction tasks. The large amount of content elements whose extraction is either unsupported or only possible in poor quality indicates a large potential for improvement in future work.

⁴³ <https://grobid.readthedocs.io/en/latest/Principles/>

Table 6: Results grouped by extraction tool.

Tool ¹	Label	# De- tected	# Pro- cessed ²	Acc	F ₁	P	R
Adobe Extract	Table	1,635	736	0.52	0.47	0.45	0.49
	Paragraph	3,985	3,088	0.85	0.74	0.72	0.76
Apache Tika	Paragraph	339,603	258,582	0.55	0.52	0.43	0.65
Camelot	Table	16,289	11,628	0.27	0.30	0.23	0.44
CERMINE	Title	16,196	14,501	0.84	0.81	0.81	0.81
	Author	19,788	14,797	0.43	0.44	0.44	0.46
	Abstract	19,342	16,716	0.71	0.72	0.68	0.76
	Reference	40,333	35,193	0.80	0.74	0.71	0.77
	Paragraph	361,273	348,160	0.89	0.85	0.83	0.87
	Section	163,077	139,921	0.40	0.35	0.32	0.38
GROBID	Title	16,196	16,018	0.92	0.91	0.91	0.92
	Author	19,788	19,563	0.54	0.52	0.52	0.53
	Abstract	19,342	18,714	0.82	0.82	0.81	0.83
	Reference	40,333	36,020	0.82	0.79	0.79	0.80
	Paragraph	361,273	358,730	0.90	0.90	0.89	0.91
	Section	163,077	163,037	0.77	0.74	0.73	0.76
	Caption	90,606	62,445	0.57	0.49	0.47	0.51
	Table	16,740	8,633	0.24	0.23	0.23	0.23
Equation	142,736	96,560	0.26	0.25	0.20	0.32	
PdfAct	Title	17,670	16,834	0.85	0.85	0.85	0.86
	Author	13,110	2,187	0.14	0.13	0.12	0.18
	Abstract	21,470	4,683	0.17	0.16	0.15	0.20
	Reference	30,263	12,705	0.19	0.15	0.17	0.20
	Paragraph	361,318	357,905	0.85	0.85	0.80	0.89
	Section	129,361	87,605	0.21	0.16	0.12	0.25
	Caption	83,435	53,314	0.45	0.45	0.40	0.52
Footer	32,457	26,252	0.23	0.20	0.25	0.16	
PyMuPDF	Paragraph	339,650	258,383	0.55	0.51	0.41	0.65
RefExtract	Reference	40,333	38,405	0.55	0.49	0.44	0.55
Science Parse	Title	11,696	11,687	0.79	0.70	0.70	0.70
	Author	471	471	0.54	0.52	0.52	0.53
	Abstract	14,150	14,149	0.83	0.81	0.73	0.90
	Reference	40,333	35,200	0.55	0.49	0.49	0.50
	Paragraph	361,318	355,529	0.79	0.76	0.76	0.76
Section	163,077	158,556	0.54	0.49	0.49	0.50	
Tabula	Table	10,361	9,456	0.29	0.28	0.20	0.46

¹ Boldface indicates the best value for each content element type.² The differences in the number of detected and processed items are due to PDF Read Exceptions or Warnings. We label an item as processed if it has a non-zero F₁ score.

5 Conclusion and Future Work

We present an open evaluation framework for information extraction from academic PDF documents. Our framework uses the DocBank dataset [31] offering 12 types and 1.5M annotated instances of content elements contained in 500K pages of arXiv papers from multiple disciplines. The dataset is larger, more topically diverse, and supports more extraction tasks than most related datasets.

We use the newly developed framework to benchmark the performance of ten freely available tools in extracting document metadata, bibliographic references, tables, and other content elements in academic PDF documents. GROBID, followed by CERMINE and Science Parse achieves the best results for the metadata and reference extraction tasks. For table extraction, Adobe Extract outperforms other tools, even though the performance is much lower than for other content elements. All tools struggle to extract lists, footers, and equations.

While DocBank covers more disciplines than other datasets, we see further diversification of the collection in terms of disciplines, document types, and content elements as a valuable task for future research. Table 2 shows that more datasets suitable for information extraction from PDF documents are available but unused thus far. The weakly supervised annotation approach used for creating the DocBank dataset is transferable to other LaTeX document collections.

Apart from the dataset, our framework can incorporate additional tools and allows easy replacement of tools in case of updates. We intend to update and extend our performance benchmark in the future.

The extraction of tables, equations, footers, lists, and similar content elements poses the toughest challenge for tools in our benchmark. In recent work, Grennan et al.[20] showed that the usage of synthetic datasets for model training can improve citation parsing. A similar approach could also be a promising direction for improving the access to currently hard-to-extract content elements.

Combining extraction approaches could lead to a one-fits-all extraction tool, which we consider desirable. The Sciencebeam-pipelines⁴⁴ project currently undertakes initial steps toward that goal. We hope that our evaluation framework will help to support this line of research by facilitating performance benchmarks of IE tools as part of a continuous development and integration process.

References

1. Ahmed, M.W., Afzal, M.T.: FLAG-PDFe: Features Oriented Metadata Extraction Framework for Scientific Publications. *IEEE Access* **8**, 99458–99469 (May 2020). <https://doi.org/10.1109/ACCESS.2020.2997907>
2. Aiello, M., Monz, C., Todoran, L., Worring, M.: Document Understanding for a Broad Class of Documents. *International Journal on Document Analysis and Recognition* **5(1)** (Aug 2002). <https://doi.org/10.1007/s10032-002-0080-x>
3. Anzaroot, S., Passos, A., Belanger, D., McCallum, A.: Learning Soft Linear Constraints with Application to Citation Field Extraction. In: *Proceedings of the 52nd*

⁴⁴ <https://github.com/elifesciences/sciencebeam-pipelines>

- Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers). pp. 593–602. Association for Computational Linguistics, Baltimore, Maryland (2014). <https://doi.org/10.3115/v1/P14-1056>
4. Azimjonov, J., Alikhanov, J.: Rule Based Metadata Extraction Framework from Academic Articles. arXiv CoRR **1807.09009v1** [cs.IR], 1–10 (2018). <https://doi.org/10.48550/arXiv.1807.09009>
 5. Bast, H., Korzen, C.: The Icecite Research Paper Management System. In: Web Information Systems Engineering – WISE 2013, vol. 8181, pp. 396–409. Springer Berlin, Heidelberg, Nanjing, China (2013). https://doi.org/10.1007/978-3-642-41154-0_30
 6. Bast, H., Korzen, C.: A Benchmark and Evaluation for Text Extraction from PDF. In: 2017 ACM/IEEE Joint Conference on Digital Libraries (JCDL). pp. 1–10. IEEE, Toronto, ON, Canada (2017). <https://doi.org/10.1109/JCDL.2017.7991564>
 7. Bhardwaj, A., Mercier, D., Dengel, A., Ahmed, S.: DeepBIBX: Deep Learning for Image Based Bibliographic Data Extraction. In: Proceedings of the 24th International Conference on Neural Information Processing. LNCS, vol. 10635, pp. 286–293. Springer, Guangzhou, China (2017). https://doi.org/10.1007/978-3-319-70096-0_30
 8. Borkar, V., Deshmukh, K., Sarawagi, S.: Automatic Segmentation of Text into Structured Records. SIGMOD Record **30**(2), 175–186 (Jun 2001). <https://doi.org/10.1145/376284.375682>
 9. Clark, C., Divvala, S.: PDFFigures 2.0: Mining Figures from Research Papers. In: Proceedings of the 16th ACM/IEEE-CS on Joint Conference on Digital Libraries. pp. 143–152. JCDL '16, Association for Computing Machinery, New York, NY, USA (2016). <https://doi.org/10.1145/2910896.2910904>
 10. Cortez, E., da Silva, A.S., Gonçalves, M.A., Mesquita, F., de Moura, E.S.: FLUX-CIM: Flexible Unsupervised Extraction of Citation Metadata. In: Proceedings of the 7th ACM/IEEE-CS Joint Conference on Digital Libraries. pp. 215–224. JCDL '07, Association for Computing Machinery, New York, NY, USA (2007). <https://doi.org/10.1145/1255175.1255219>
 11. Cortez, E., da Silva, A.S., Gonçalves, M.A., Mesquita, F., de Moura, E.S.: A Flexible Approach for Extracting Metadata from Bibliographic Citations. JASIST **60**(6), 1144–1158 (Jun 2009). <https://doi.org/10.1002/asi.21049>
 12. Councill, I., Giles, C.L., Kan, M.Y.: ParsCit: an Open-source CRF Reference String Parsing Package. In: Proceedings of the Sixth International Conference on Language Resources and Evaluation. European Language Resources Association, Marrakech, Morocco (2008), <https://aclanthology.org/L08-1291/>
 13. Cui, B.G., Chen, X.: An Improved Hidden Markov Model for Literature Metadata Extraction. In: Advanced Intelligent Computing Theories and Applications, vol. 6215, pp. 205–212. Springer Berlin Heidelberg, Changsha, China (2010). https://doi.org/10.1007/978-3-642-14922-1_26
 14. Day, M.Y., Tsai, R.T.H., Sung, C.L., Hsieh, C.C., Lee, C.W., Wu, S.H., Wu, K.P., Ong, C.S., Hsu, W.L.: Reference metadata extraction using a hierarchical knowledge representation framework. Decision Support Systems **43**(1), 152–167 (Feb 2007). <https://doi.org/10.1016/j.dss.2006.08.006>
 15. De La Torre, M., Aguirre, C., Anshutz, B., Hsu, W.: MATEC: Metadata-analytic text extractor and section classifier for scientific publications. In: Proceedings of the 10th International Joint Conference on Knowledge Discovery, Knowledge Engineering and Knowledge Management. vol. 1, pp. 261–267. SciTePress (2018). <https://doi.org/10.5220/0006937702610267>

16. Fan, T., Liu, J., Qiu, Y., Jiang, C., Zhang, J., Zhang, W., Wan, J.: PARDA: A Dataset for Scholarly PDF Document Metadata Extraction Evaluation. In: Collaborative Computing: Networking, Applications and Worksharing. pp. 417–431. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-12981-1_29
17. Färber, M., Thiemann, A., Jatowt, A.: A High-Quality Gold Standard for Citation-based Tasks. In: Proceedings of the Eleventh International Conference on Language Resources and Evaluation. European Language Resources Association, Miyazaki, Japan (2018), <https://aclanthology.org/L18-1296>
18. Giuffrida, G., Shek, E.C., Yang, J.: Knowledge-Based Metadata Extraction from PostScript Files. In: Proceedings of the Fifth ACM Conference on Digital Libraries. pp. 77–84. DL '00, Association for Computing Machinery, New York, NY, USA (2000). <https://doi.org/10.1145/336597.336639>
19. Granitzer, M., Hristakeva, M., Jack, K., Knight, R.: A Comparison of Metadata Extraction Techniques for Crowdsourced Bibliographic Metadata Management. In: Proceedings of the 27th Annual ACM Symposium on Applied Computing. pp. 962–964. SAC '12, Association for Computing Machinery, New York, NY, USA (2012). <https://doi.org/10.1145/2245276.2245462>
20. Grennan, M., Beel, J.: Synthetic vs. Real Reference Strings for Citation Parsing, and the Importance of Re-training and Out-Of-Sample Data for Meaningful Evaluations: Experiments with GROBID, GIANT and CORA. In: Proceedings of the 8th International Workshop on Mining Scientific Publications. pp. 27–35. Association for Computational Linguistics, Wuhan, China (2020), <https://aclanthology.org/2020.wosp-1.4>
21. Grennan, M., Schibel, M., Collins, A., Beel, J.: GIANT: The 1-Billion Annotated Synthetic Bibliographic-Reference-String Dataset for Deep Citation Parsing [Data] (2019). <https://doi.org/10.7910/DVN/LXQXAO>
22. Hashmi, A.M., Afzal, M.T., ur Rehman, S.: Rule Based Approach to Extract Metadata from Scientific PDF Documents. In: 2020 5th International Conference on Innovative Technologies in Intelligent Systems and Industrial Applications (CITISIA). pp. 1–4. IEEE, Sydney, Australia (2020). <https://doi.org/10.1109/CITISIA50690.2020.9371784>
23. Hetzner, E.: A Simple Method for Citation Metadata Extraction Using Hidden Markov Models. In: Proceedings of the 8th ACM/IEEE-CS Joint Conference on Digital Libraries. pp. 280–284. JCDL '08, Association for Computing Machinery, New York, NY, USA (2008). <https://doi.org/10.1145/1378889.1378937>
24. Kasdorf, W.E.: The Columbia Guide to Digital Publishing. Columbia University Press, USA (2003)
25. Kern, R., Jack, K., Hristakeva, M.: TeamBeam - Meta-Data Extraction from Scientific Literature. D-Lib Magazine **18**(7/8) (Jul 2012). <https://doi.org/10.1045/july2012-kern>
26. Klein, D., Manning, C.D.: Conditional Structure versus Conditional Estimation in NLP Models. In: Proceedings of the 2002 Conference on Empirical Methods in Natural Language Processing (EMNLP). pp. 9–16. Association for Computational Linguistics, Pennsylvania, Philadelphia, PA, USA (2002). <https://doi.org/10.3115/1118693.1118695>
27. Klink, S., Dengel, A., Kieninger, T.: Document Structure Analysis Based on Layout and Textual Features. In: IAPR International Workshop on Document Analysis Systems. IAPR, Rio de Janeiro, Brazil (2000)
28. Körner, M., Ghavimi, B., Mayr, P., Hartmann, H., Staab, S.: Evaluating Reference String Extraction Using Line-Based Conditional Random Fields: A Case Study

- with German Language Publications. In: *New Trends in Databases and Information Systems*. pp. 137–145. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-67162-8_15
29. Lafferty, J.D., McCallum, A., Pereira, F.C.N.: Conditional Random Fields: Probabilistic Models for Segmenting and Labeling Sequence Data. In: *Proceedings of the Eighteenth International Conference on Machine Learning*. pp. 282–289. ICML '01, Morgan Kaufmann Publishers Inc., San Francisco, CA, USA (2001), <https://dl.acm.org/doi/10.5555/645530.655813>
 30. Ley, M.: DBLP: Some Lessons Learned. *Proc. VLDB Endowment* **2**(2), 1493–1500 (Aug 2009). <https://doi.org/10.14778/1687553.1687577>
 31. Li, M., Xu, Y., Cui, L., Huang, S., Wei, F., Li, Z., Zhou, M.: DocBank: A Benchmark Dataset for Document Layout Analysis. In: *Proceedings of the 28th International Conference on Computational Linguistics*. pp. 949–960. International Committee on Computational Linguistics, Barcelona, Spain (Online) (2020). <https://doi.org/10.18653/v1/2020.coling-main.82>
 32. Lipinski, M., Yao, K., Breitingner, C., Beel, J., Gipp, B.: Evaluation of Header Metadata Extraction Approaches and Tools for Scientific PDF Documents. In: *Proceedings of the 13th ACM/IEEE-CS Joint Conference on Digital Libraries*. pp. 385–386. JCDL '13, Association for Computing Machinery, New York, NY, USA (2013). <https://doi.org/10.1145/2467696.2467753>
 33. Livathinos, N., Berrospi, C., Lysak, M., Kuropiatnyk, V., Nassar, A., Carvalho, A., Dolfi, M., Auer, C., Dinkla, K., Staar, P.: Robust PDF Document Conversion using Recurrent Neural Networks. *Proceedings of the AAAI Conference on Artificial Intelligence* **35**(17), 15137–15145 (May 2021). <https://doi.org/10.1609/aaai.v35i17.17777>
 34. Lo, K., Wang, L.L., Neumann, M., Kinney, R., Weld, D.: S2ORC: The Semantic Scholar Open Research Corpus. In: *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. pp. 4969–4983. Association for Computational Linguistics, Online (2020). <https://doi.org/10.18653/v1/2020.acl-main.447>
 35. Lopez, P.: GROBID (2008), <https://github.com/kermitt2/grobid>
 36. Lopez, P.: GROBID: Combining Automatic Bibliographic Data Recognition and Term Extraction for Scholarship Publications. In: *Research and Advanced Technology for Digital Libraries, LNCS*, vol. 5714, pp. 473–474. Springer Berlin Heidelberg (2009). https://doi.org/10.1007/978-3-642-04346-8_62
 37. Mao, S., Kim, J., Thoma, G.R.: A Dynamic Feature Generation System for Automated Metadata Extraction in Preservation of Digital Materials. In: *1st International Workshop on Document Image Analysis for Libraries*. pp. 225–232. IEEE Computer Society, Palo Alto, CA, USA (2004). <https://doi.org/10.1109/DIAL.2004.1263251>
 38. Mao, S., Rosenfeld, A., Kanungo, T.: Document structure analysis algorithms: a literature survey. In: *Proceedings Document Recognition and Retrieval X. SPIE Proceedings*, vol. 5010, pp. 197–207. SPIE, Santa Clara, California, USA (Jan 2003). <https://doi.org/10.1117/12.476326>
 39. McCallum, A.K., Nigam, K., Rennie, J., Seymore, K.: Automating the Construction of Internet Portals with Machine Learning. *Information Retrieval* **3**(2), 127–163 (Jul 2000). <https://doi.org/10.1023/A:1009953814988>
 40. National Library of Medicine: PubMed, <https://pubmed.ncbi.nlm.nih.gov/>
 41. National Library of Medicine: PubMed Central, <https://www.ncbi.nlm.nih.gov/pmc/>

42. Ojokoh, B., Zhang, M., Tang, J.: A trigram hidden Markov model for metadata extraction from heterogeneous references. *Information Sciences* **181**(9), 1538–1551 (May 2011). <https://doi.org/10.1016/j.ins.2011.01.014>
43. Ororbia, A.G., Wu, J., Khabsa, M., Williams, K., Giles, C.L.: Big Scholarly Data in CiteSeerX: Information Extraction from the Web. In: Proceedings of the 24th International Conference on World Wide Web. pp. 597–602. WWW '15 Companion, Association for Computing Machinery, New York, NY, USA (2015). <https://doi.org/10.1145/2740908.2741736>
44. Palmero, G., Dimitriadis, Y.: Structured document labeling and rule extraction using a new recurrent fuzzy-neural system. In: Proceedings of the Fifth International Conference on Document Analysis and Recognition. pp. 181–184. Springer, Bangalore, India (1999). <https://doi.org/10.1109/ICDAR.1999.791754>
45. Peng, F., McCallum, A.: Accurate Information Extraction from Research Papers using Conditional Random Fields. In: Proceedings of the Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics: HLT-NAACL. pp. 329–336. Association for Computational Linguistics, Boston, Massachusetts, USA (2004), <https://aclanthology.org/N04-1042>
46. Prasad, A., Kaur, M., Kan, M.Y.: Neural ParsCit: a deep learning-based reference string parser. *International Journal on Digital Libraries* **19**(4), 323–337 (Nov 2018). <https://doi.org/10.1007/s00799-018-0242-1>
47. Rizvi, S.T.R., Dengel, A., Ahmed, S.: A Hybrid Approach and Unified Framework for Bibliographic Reference Extraction. *IEEE Access* **8**, 217231–217245 (Dec 2020). <https://doi.org/10.1109/ACCESS.2020.3042455>
48. Rizvi, S.T.R., Lucieri, A., Dengel, A., Ahmed, S.: Benchmarking Object Detection Networks for Image Based Reference Detection in Document Images. In: 2019 Digital Image Computing: Techniques and Applications (DICTA). pp. 1–8. IEEE, Perth, WA, Australia (2019). <https://doi.org/10.1109/DICTA47822.2019.8945991>
49. Rodrigues Alves, D., Colavizza, G., Kaplan, F.: Deep Reference Mining From Scholarly Literature in the Arts and Humanities. *Frontiers in Research Metrics and Analytics* **3**, 21 (Jul 2018). <https://doi.org/10.3389/frma.2018.00021>
50. Saier, T., Färber, M.: Bibliometric-Enhanced arXiv: A Data Set for Paper-Based and Citation-Based Tasks. In: Proceedings of the 8th International Workshop on Bibliometric-enhanced Information Retrieval (BIR). CEUR Workshop Proceedings, vol. 2345, pp. 14–26. CEUR-WS.org, Cologne, Germany (2019), <http://ceur-ws.org/Vol-2345/paper2.pdf>
51. Saier, T., Färber, M.: unarXive: A Large Scholarly Data Set with Publications' Full-Text, Annotated In-Text Citations, and Links to Metadata. *Scientometrics* **125**(3), 3085–3108 (Dec 2020). <https://doi.org/10.1007/s11192-020-03382-z>
52. Schloss Dagstuhl - Leibniz Center for Informatics, University of Trier: dblp: computer science bibliography, <https://dblp.org/>
53. Souza, A., Moreira, V., Heuser, C.: ARCTIC: Metadata Extraction from Scientific Papers in Pdf Using Two-Layer CRF. In: Proceedings of the 2014 ACM Symposium on Document Engineering. pp. 121–130. DocEng '14, Association for Computing Machinery, New York, NY, USA (2014). <https://doi.org/10.1145/2644866.2644872>
54. Tkaczyk, D., Collins, A., Sheridan, P., Beel, J.: Machine Learning vs. Rules and Out-of-the-Box vs. Retrained: An Evaluation of Open-Source Bibliographic Reference and Citation Parsers. In: Proceedings of the 18th ACM/IEEE on Joint Conference on Digital Libraries. pp. 99–108. JCDL '18, Association for Computing Machinery, New York, NY, USA (2018). <https://doi.org/10.1145/3197026.3197048>

55. Tkaczyk, D., Szostek, P., Fedoryszak, M., Dendek, P.J., Bolikowski, L.: CERMINE: automatic extraction of structured metadata from scientific literature. *International Journal on Document Analysis and Recognition (IJ DAR)* **18**(4), 317–335 (Dec 2015). <https://doi.org/10.1007/s10032-015-0249-8>
56. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is All You Need. In: *Proceedings of the 31st International Conference on Neural Information Processing Systems*. pp. 6000–6010. NIPS’17, Curran Associates Inc., Red Hook, NY, USA (2017), <https://dl.acm.org/doi/10.5555/3295222.3295349>
57. Vilnis, L., Belanger, D., Sheldon, D., McCallum, A.: Bethe Projections for Non-Local Inference. In: *Proceedings of the Thirty-First Conference on Uncertainty in Artificial Intelligence*. pp. 892–901. UAI’15, AUAI Press, Arlington, Virginia, USA (2015). <https://doi.org/10.48550/arXiv.1503.01397>