

STEREO: Scientific Text Reuse in Open Access Publications

Lukas Gienapp^{1,*}, Wolfgang Kircheis¹, Bjarne Sievers¹, Benno Stein², and Martin Potthast^{1,*}

¹Leipzig University, Text Mining and Retrieval Group, Leipzig, DE-04109, Germany

²Bauhaus-Universität Weimar, Web Technology and Information Systems Group, Weimar, DE-99423, Germany

*corresponding authors: Lukas Gienapp, Martin Potthast (<first>.<last>@uni-leipzig.de)

ABSTRACT

We present the **Webis-STEREO-21** dataset, a massive collection of **Scientific Text Reuse in Open-access** publications. It contains more than 91 million cases of reused text passages found in 4.2 million unique open-access publications. Featuring a high coverage of scientific disciplines and varieties of reuse, as well as comprehensive metadata to contextualize each case, our dataset addresses the most salient shortcomings of previous ones on scientific writing. **Webis-STEREO-21** allows for tackling a wide range of research questions from different scientific backgrounds, facilitating both qualitative and quantitative analysis of the phenomenon as well as a first-time grounding on the base rate of text reuse in scientific publications.

Background & Summary

The reuse of text has a longstanding history in science. In qualitative research, besides verbatim quotations, the techniques of paraphrasing, translation, and summarization are instrumental to both teaching and learning scientific writing as well as to gaining new scientific insights¹. In quantitative research, the use of templates as an efficient way of reporting new results on otherwise standardized workflows is common². As science often progresses incrementally, authors may also reuse their texts across (different types of) subsequent publications on the same subject (also called “text recycling”)²⁻⁴. Likewise, in interdisciplinary research, reuse across publications at venues of different disciplines has been observed to promote the dissemination of new insights^{5,6}. Orthogonal to all of these manifestations of text reuse is the scientific context that stipulates its legitimacy: Plagiarism, the intentional reuse of text with lacking acknowledgment of its original source, violates codes of honor and academic integrity⁷.

Text reuse has been quantitatively studied in many scientific disciplines^{1,2,8-10}; yet few studies assess the phenomenon at scale, beyond what can be manually analyzed^{9,10}. Large-scale studies require the use of automatic text reuse detection technology. This being both algorithmically challenging and computationally expensive, lacking expertise or budget may have prevented such studies. Employing proprietary analysis software or services instead, too, is subject to budgetary limitations, in addition to their lack of methodological transparency and reproducibility.

Text reuse detection itself is still subject to ongoing research in natural language processing and information retrieval. Setting up a custom processing pipeline thus demands an evaluation against the state of the art. The challenges in constructing a competitive solution for this task arise from the aforementioned diversity of different forms of text reuse, the large solution space of detection approaches, and the need to apply heuristics that render a given solution sufficiently scalable. Preprocessing a collection of scientific publications, too, presents its own difficulties. This includes the noisy and error-prone conversion of a publication’s original PDF version to machine-readable text and the collection of reliable metadata about the publications. The available quantitative studies on scientific text reuse lack with respect to the presentation of preprocessing steps taken, the design choices of the solution to text reuse detection, and their justification in terms of rigorous evaluation. Altogether, comparable, reproducible, reliable, and accessible research on the phenomenon of scientific text reuse remains an open problem.

To provide for a solid new foundation for the investigation of scientific text reuse within and across disciplines, we compile **Webis-STEREO-21**. To overcome the aforementioned issues, we stipulate three design principles for the creation of the dataset: (1) high coverage, both in terms of the number of included publications and the variety of scientific disciplines; (2) a scalable approach to reuse detection with a focus on high precision at a competitive recall, capturing a comprehensive set of reused passages as reliable resource for research on scientific text reuse; and (3) comprehensive metadata to contextualize each case, to address a wide range of potential hypotheses, and to provide the basis for semantic post-processing, for instance, to separate benign text reuse from plagiarism.

Webis-STEREO-21 results from applying scalable text reuse detection approaches to a large collection of scientific open-access publications, exhaustively comparing all documents to extract a comprehensive dataset of reused passages between them. It contains more than 91 million cases of reused passages among 4.2 million unique publications. The cases stem from 46 scientific fields of study, grouped into 14 scientific areas in all four major scientific disciplines, and spanning over 150 years of scientific publishing between 1860 and 2018. The data is openly accessible to be useful to a wide range of researchers with different scientific backgrounds, enabling both qualitative and quantitative analysis.

Methods

To compile a dataset on text reuse in scientific writing, three things are required: a large collection of scientific publications; detailed metadata for each of these publications; and a scalable, yet accurate method for detecting reuse between the compiled texts. This section formally introduces all three of these components. Our pipeline and its individual steps are depicted in Figure 1 for reference. Note that text reuse detection is typically implemented as a bipartite process of source retrieval to identify candidate pairs of documents, followed by text alignment to identify reused portions of text between both candidate documents¹¹.

Publication Acquisition, Preprocessing and Metadata

Before the detection of scientific text reuse at scale can commence, a large amount of scientific publications and detailed metadata has to be collected. We assemble such data in the five steps document selection, plain text extraction, text preprocessing, meta data acquisition, and meta data standardization.

For raw document selection, we make use of the CORE dataset¹², one of the largest collections of open access academic publications, obtained from more than 12,000 data providers. First, we identify the 6,015,512 unique open-access DOIs in the CORE dataset dump from March 1st, 2018. Since the plain texts extracted from the publications' PDF files as supplied by the CORE data are of highly varying quality, and since no structural annotations (such as markup for citations, in-text references, section annotations) are available, we opted to obtain the original PDF files of the identified open access DOIs from various openly accessible repositories.

Plain text extraction was then repeated on the acquired PDF files using the standardized state-of-the-art toolchain GROBID¹³. The extracted texts are preprocessed as follows: For each document, abstract and main text paragraphs are extracted as indicated by the content tags provided by GROBID; in-text references, tables, and figures are omitted. We choose to further drop the bibliographical data, special characters, and numerical characters to reduce the amount of false positives for the subsequent text alignment step. Since the retrieval models underlying the text alignment detection compute the word overlap between document passages, standardized or meta text fragments such as citations, numbers, author names, affiliations, and references exponentially increase the set of passages matched between documents that otherwise may not share text. A minimum of 1,000 and a maximum of 60,000 whitespace-separated words is imposed as an effective heuristic to filter common plain text extraction errors. Overall, we obtained and extracted clean plain text for 4,267,166 documents (70% of the original CORE dataset).

Given this set, we cross-check and supplement the metadata provided by CORE with additional data from the Microsoft Open Academic Graph^{14,15}, which features fields-of-study annotations for a wide set of publications. Since the annotated disciplines do not follow a hierarchical scheme, and since they are of differing granularity per publication (e.g., "humanities" as a whole vs. "chemical solid-state research" as a subfield of chemistry), we manually map the classification found in the Microsoft Open Academic Graph to the standardized, hierarchical *DFG Classification of Scientific Disciplines, Research Areas, Review Boards and Subject Areas*^{16,1}.

Text Reuse Detection in Large Document Collections

Given the large collection of plain texts, a text reuse detection approach is applied to identify reused passages between pairs of texts. At heart, the detection of text reuse boils down to the task of *text alignment*¹⁷: Given two documents d and d' , determine the set $R_{d,d'}$ of all pairs of reused text passages (t, t') , where $t \subseteq d$ and $t' \subseteq d'$.

When given a set of documents D , the computational complexity of detecting the set R_D of all pairs of reused text passages among all pairs of documents, $\{(d, d') \mid d, d' \in D\}$, grows quadratically with the document count, $|D|$. If an exhaustive comparison exceeds the available computing capacity, a source retrieval step has to be carried out first^{18,19}: Given a document d and a document collection D , a subset $D_d \subseteq D$ of candidate documents is retrieved containing those documents $d' \in D_d$ whose likelihood of sharing a reused text passage with d is sufficiently high. The size $|D_d|$ for a given document d varies dependent on how many documents in D have a sufficiently high reuse likelihood, but is expected to be $|D_d| \ll |D|$. Note that the estimation of a reuse likelihood for all pairs $\{(d, d') \mid d, d' \in D\}$ can be operationalized with linear time complexity in $|D|$, e.g., by using locality-sensitive hashing²⁰.

¹We opt to replace the term *review board* as used in the DFG classification with the more conventional denomination *field of study*.

Source retrieval introduces an error in terms of lower recall, i.e., missed pairs of reused passages that would have been detected by an exhaustive text alignment of every document pair. The goal of source retrieval therefore is to closely approximate the recall of an exhaustive comparison while minimizing the sum of the sizes of retrieved candidate documents $\sum_{d \in D} |D_d|$, i.e., the computing capacity actually expended. Altogether, to detect reused text passages in D , we retrieve for each $d \in D$ the set D_d of candidate documents and compute for each $d' \in D_d$ the set of reused text passages $R_{d,d'}$. The union of all detected passages is used as an approximation of R_D .

Source Retrieval

The source retrieval component of our pipeline, i.e., the computation of D_d for a given d , is operationalized by treating text reuse as a locally bounded phenomenon. Candidate pairs are presumed to be identifiable by comparing text passages between two documents. If the similarity is sufficiently high for at least one combination of passages between two documents d, d' , then d' is considered a candidate for d . To achieve high scalability, we apply a locality-sensitive hash function h to each passage t in a document d , thus representing each document as a set of hash values $h(t)$. The similarity between two passages t, t' is then approximated by the extent of overlap between their hash sets, $|h(t) \cap h(t')|$. Thus, a document d' is considered a candidate source of text reuse for a document d , if at least one of their passage-level hash sets intersects²⁰: $\exists t \subseteq d \exists t' \subseteq d' : h(t) \cap h(t') \neq \emptyset$.

In our pipeline, the documents are first divided into consecutive passages of n terms, which in turn are each embedded using a bag-of-words representation, i.e., as their term occurrence vector. To find matching passages between two documents, MinHash²¹ is chosen as hashing scheme for h . For each passage vector, multiple individual hash functions are applied to each element, and the minimum hash value of each pass is saved, yielding the set minimum hashes for a passage. It can be shown that, for m individual hashes per passage, two texts with a Jaccard similarity of at least m^{-1} are guaranteed to produce a hash collision between their hash sets, rendering them a suitable approximation for the lower bound of document similarity²¹.

Hash-based source retrieval allows for a significant reduction of the required computation time, since the hash-based approximation has a linear time complexity with respect to $|D|$, as opposed to the quadratic complexity of vector comparisons²⁰. As a result, our source retrieval computation time could be fitted into the allotted budget of two months of computing time on a 130-node Apache Spark cluster, with 12 CPUs and 196 GB RAM per node.

Besides its efficiency, the outlined approach provides us with a second highly beneficial characteristic—the source retrieval step depends on only two parameters: the passage length n , which can be chosen according to computational constraints, and the number of hashes m , which is chosen according to the required minimum similarity separating two passages. Applying this scheme to all passages in all documents, we detect all document pairs that are locally similar with at least m^{-1} Jaccard similarity, i.e., which have a word overlap in their bag-of-words passage representations. By choosing this bound lower than the minimum detection threshold of the subsequent text alignment step, the search space is pruned without a loss in accuracy. Thus, our hash-based similarity search primarily impacts the precision, but not the recall of the source retrieval step with respect to the subsequent text alignment step.

Text Alignment

Given a set of candidates D_d for a document d obtained from source retrieval, the text alignment component extracts the reused passages $R_{d,d'}$ of every pair of documents d and $d' \in D_d$. Here, we follow the seed-and-extend approach to local sequence alignment²². An overview of the process is given in Figure 2: first, texts are separated into small spans (*‘chunking’*). Then, matching text pieces in the cartesian product of those spans are computed according to a similarity function ϕ . Its purpose is to identify short, matching sequences of text (*seeds*) between two documents d and d' , which are nearly guaranteed to have the same meaning or be considered instances of the same concept. Finally, matches sufficiently close to each other are joined to form larger passages (*‘extending’*). A pair of passages ($t \subseteq d, t' \subseteq d'$) across both documents then constitutes a reuse case.

With regard to the chunking and seeding steps, two methodological choices have to be made: first, how to divide a given document into small spans for comparison purposes; and second, how to compare (i.e., which similarity function to apply) pairs of text spans to identify whether they are sufficiently similar, indicating possible text reuse. For the former, a widely used approach that has been shown to produce accurate results is to divide a text into (word) n -grams, namely continuous spans of words of length n ^{9,23}. These n -grams are overlapping and generated by “sliding” a window over the text. For example, consider a 4-gram with an overlap of three: then the first chunk consists of the words 1 to 4 of a text, the second consists of the words 2 to 5, and so forth. The chunk sequences of the two documents d and d' are then exhaustively compared, i.e., every chunk of d is compared to every chunk of d' by comparing their hash values (using a string hash function, not locality-sensitive hashing). Calculating all hash collisions between the chunks of both texts can be accomplished in linear time in the length of the texts. Two parameters can be modified in this approach: the n -gram size n , and the n -gram overlap k .

The matching chunks found through hash collisions indicate “seeds” of a potentially longer case of reused text. To determine whether text has actually been reused, the *extension* step joins co-aligned matching chunks into longer passages when a sufficient number are found close to each other. In this manner, not only cases of coherent reuse, but also cases where text was copied and then paraphrased can be detected. For example, if consecutive sentences in a source text are copied into a target text, and then another (new) sentence is placed in between them, this would be still considered a case of coherent reuse. Seeding would succeed in identifying the two copied sentences on their own, yet the extension recognizes that both seeds are in close proximity in the source material, forming a single identified case of reuse. The opposite case is also possible: two chunks from different locations in a source text are placed close to each other in the target texts. Here, again, an extension is required to reconstruct the full scope of reuse. Note that the extension approach depends on a single parameter, Δ , which is the maximum distance in characters for two seeds to be considered a single reuse case.²³

We supply a massively parallelizable implementation of both the source retrieval and text alignment step, which allows for detecting text reuse in the highly scalable manner needed given the amount and length of input documents: 4.2 million publications, amounting to a total of 1.1 Terabyte of text data. Overall, the text alignment step accounted for an additional 3.5 months of computing time on the aforementioned Spark cluster.

Data Records

Two types of data records are included in the **Webis-STEREO-21** corpus: the reuse case data, which contains all identified cases of text reuse, their metadata, and the publication data for each individual document considered when computing the cases, including publication year, field of study annotations, and plain text. The records can be cross-referenced using a publications’ DOI as primary key.

Each of the 91,466,374 identified cases of potential text reuse is represented as an individual entry. A document pair can contain multiple occurrences of text reuse, each of which is treated as a unique case. For each case, three kinds of information are provided (Table 1 overviews the fields available): (1) *text*, which includes the matched text itself, and 100 characters of surrounding context on either side; (2) *locators*, which identify the matched text by its in-text location, using character offsets to mark start and end; (3) *metadata*, which provides context about the publications involved in the case both by year and by field of study. In the context of a case, we refer to the first involved publication as *a*, and to the second as *b*. This, however, does not indicate a directionality of reuse.

In the publication data, the metadata for all 4,267,166 documents considered for reuse computation is provided, including metadata that is not part of a reuse case, to help contextualize analyses derived from the case data. Table 2 provides detailed information about the fields available for each publication.

The **Webis-STEREO-21** corpus is archived at Zenodo. It is distributed as multiple partial files in the JSONL format, with one JSON-encoded case/publication per line. We provide two versions of the corpus: a freely distributed “metadata-only” version (10.5281/zenodo.5575285)²⁴, which includes locators and metadata, but not the text data, and an access-restricted complete version (10.5281/zenodo.5575320)²⁵, which includes the text-complete cases.

Technical Validation

This section motivates and details the parameter choices for the source retrieval and text alignment components of the text reuse detection pipeline. For both steps, we strive for maximum accuracy given the constraints for scalability and computational efficiency imposed by the amount of data to be processed.

Source Retrieval

Objective of the source retrieval step is to prune the search space of document pairs by reducing the number of pairs to be compared in subsequent (computationally expensive) steps. The optimization criterion is recall: Ideally, no document pair containing text reuse should be overlooked, i.e., the false negative rate should be minimized.

To allow for a fine-grained detection of local similarity, the parameters of the source retrieval step are set to $n = 50$ and $m = 10$, i.e., documents are split into chunks of 50 words, with 10 MinHash values computed per chunk. Consequently, the source retrieval is able to identify pairs of documents that share as little as a 9-word overlap between any two chunks. Note that these words do not need to be consecutive since chunks are represented in a bag-of-words model. Since the subsequent alignment step operates at the 8-gram level (= eight consecutive units, see next section), the filtering within the source retrieval step eliminates only pairs for which the subsequent alignment step is guaranteed to not find any matches. Overall, the source retrieval step yields 3.305×10^{12} total unique document pairs for further analysis. Given the initial document count of 4,656,302, this represents a pruning of the search space by over 84% compared to an exhaustive pairwise comparison, rendering the source retrieval step highly effective in improving overall efficiency.

Text Alignment

The text alignment step identifies the exact location and extent of reuse between two documents. Unlike source retrieval, text alignment is a precision-oriented task: Special emphasis is put on precision over recall for the parameter choice to minimize the number of false positives; a low noise ratio is paramount for meaningful future analyses of the dataset. To ensure the effectiveness of the text alignment, its parameters are chosen by performing a grid-search, using precision, recall, and $F_{0.5}$ as effectiveness scores. We employ the PAN-13 Text Alignment Corpus for evaluation, which has been previously published as a benchmark dataset for text alignment by the PAN Shared Task on Plagiarism Detection²⁶. It contains 10,000 pairs of texts that are subject to different so-called obfuscation strategies that simulate more difficult cases of plagiarism, where the authors tried to hide text-reuse by paraphrasing the copied text to some extent.

The applied obfuscation strategies include *random obfuscation* (shuffling, adding, deleting, and replacing words or short phrases at random), *cyclic translation obfuscation* (a text is translated into another language, and back to the original language; possibly with more languages in-between), *summary obfuscation* (human-generated summaries of source texts), non-obfuscated plagiarism, and pairs without any text reuse between source and target. Each of these strategies is equally represented in the corpus with 2,000 pairs. The grid search identifies an n -gram size of $n = 8$, an n -gram overlap of $k = 7$, and an extension range of $\Delta = 250$ as optimal. Under these conditions, our implementation of the text alignment step achieves a precision of 0.93 at a recall of 0.46, and thus an $F_{0.5}$ score of 0.77. This renders our approach as highly competitive when compared to other approaches evaluated as part of the PAN Shared Tasks, placing it among the best for precision and $F_{0.5}$, while being the only approach adhering to the scalability requirements imposed by the scale of our analysis. A full overview of the attained scores for comparison with competing systems is given in Table 3.

When evaluating separately per obfuscation strategy (Table 4), precision is very high throughout, exceeding 0.88 in all cases. This fact places our system within 0.08 for the best-performing, yet computationally complex approach for each obfuscation strategy. Yet, a drop in recall can be observed for heavily obfuscated text, down to 0.1 for summary obfuscation. This effect is expected and also noticeable among the other approaches presented at PAN. Moreover, our focus is not on plagiarism but on establishing a general baseline for text reuse in science, where more literal reuse (e.g., citations, idiomatic reuse, template writing, etc.) is expected to be the norm. The heavily obfuscated test sets studied at PAN were dedicated to study extreme cases of plagiarism, where an author expends much effort to hide the fact via severe forms of paraphrasing. The investment of such an effort, however, goes against the time savings that can be expected from reusing text, so that it can be presumed that the vast amount of paraphrased text reuse will only make smaller changes to a text instead of changing every last n -gram. This observation is partially corroborated when reviewing the many cases of academic plagiarism found in dissertation theses throughout the recent years²⁷, where heavy paraphrasing is hardly ever observed. We therefore deemed the investment of an exponentially higher computation cost into retrieving such cases to be uneconomical.

Nevertheless, two measures to tackle this recall issue have been proposed by PAN participants: (1) customized approaches for each of the different obfuscation strategies, employing heuristics to detect which kind a document pair is exhibiting, and (2) ensemble methods encompassing different seeding and extension strategies, combined into a single result. However, the first is not applicable to our situation, as there is no ground truth data available to fine-tune such a classification to scientific writing. The second comes at a very high runtime and algorithmic complexity. This is reflected in Table 3 as well: for each approach, the (asymptotic) algorithmic complexity is noted, for two given sequences of length n and m . Only six of the approaches presented at PAN perform in sub-quadratic time, a necessary requirement for large-scale detection. Furthermore, the data specificity for each approach is listed: it denotes how much fine-tuning to the PAN data has taken place, for example by crafting specialized corpus-dependent features, using ensemble methods, trained classifiers, or specific approaches for each of the different obfuscation strategies, all of which reduce the generalizability and transferability of the approaches to other data.

Out of the four other approaches to combine sub-quadratic runtime complexity with low data specificity, ours performs best with regards to recall and $F_{0.5}$ score, and is second in precision by a very close margin. Against this background, our achieved detection performance is comparably outstanding, given the specialized requirements in terms of scalability as well as the focus on a high-precision identification.

Usage Notes

As outlined in the ‘Data Records’ section, the dataset is distributed in two different versions, where the full corpus is access-restricted, while a smaller metadata-only version is readily available. The rationale for this split is twofold: First, the full dataset is significantly larger in size, which may exceed the processing capabilities available to some researchers interested in working with the data. Moreover, a wide range of research questions can be addressed

based on the metadata alone; providing a focused version of the data lowers the barrier of access without impeding utilizability. Second, the full version contains texts which (while being open-access) are the intellectual property of their respective authors. Therefore, the full version is access-restricted at Zenodo and available on request only. Requests can be issued following Zenodo's standardized procedures and will be granted to third parties for research and other non-commercial purposes.

Alongside the dataset the text preprocessing component of our processing pipeline is included as a standalone Python script for use by researchers who obtain the metadata-only version of the corpus. This allows for transforming GROBID extraction output into a compatible text format, such that the locators included in the dataset can be applied to reconstruct the text portions of individual reuse cases. It alleviates the need to obtain the full corpus, as text data for small-scale case studies can be retroactively obtained from secondary sources. Since the publications in the corpus is based upon are open access, the original PDF files can be easily retrieved by their DOI, further lowering the barrier of access. We also provide the code used to calculate the corpus statistics given throughout this article as a usage example to others interested in working with the data.

Code Availability

The complete source code used for candidate retrieval and text alignment is openly accessible and permanently available on GitHub². The data processing pipeline is written in Python 3.7, utilizing the `pyspark` framework. The compute cluster on which we carried out the data processing and our experiments run Spark version 2.4.8. The text alignment component is written in Go 1.16 and can be used as a standalone application. Detailed documentation about each pipeline component, recommendations for compute resources, and suggestions for parameter choices are distributed alongside the code to facilitate code reuse.

References

1. Sun, Y.-C. & Yang, F.-Y. Uncovering published authors' text-borrowing practices: Paraphrasing strategies, sources, and self-plagiarism. *J. Engl. for Acad. Purp.* **20**, 224–236, 10.1016/j.jeap.2015.05.003 (2015).
2. Anson, I. G. & Moskovitz, C. Text recycling in stem: A text-analytic study of recently published research articles. *Accountability Res.* **28**, 349 – 371 (2020).
3. Moskovitz, C. Self-Plagiarism, Text Recycling, and Science Education. *BioScience* **66**, 5–6, 10.1093/biosci/biv160 (2015). <https://academic.oup.com/bioscience/article-pdf/66/1/5/7459167/biv160.pdf>.
4. Hall, S., Moskovitz, C. & Pemberton, M. A. Attitudes toward text recycling in academic writing across disciplines. *Accountability Res.* **25**, 142–169, 10.1080/08989621.2018.1434622 (2018). <https://doi.org/10.1080/08989621.2018.1434622>.
5. Bird, S. J. Self-plagiarism and dual and redundant publications: What is the problem? *Sci. engineering ethics* **8**, 543–544 (2002).
6. Wen, Q. & Gao, Y. Dual publication and academic inequality. *Int. J. Appl. Linguist.* **17**, 221–225 (2007).
7. Eberle, M. E. Paraphrasing, plagiarism, and misrepresentation in scientific writing. *Transactions Kansas Acad. Sci. (1903-)* **116**, 157–167 (2013).
8. Ganascia, J.-G., Glaudes, P. & Del Lungo, A. Automatic detection of reuses and citations in literary texts. *Lit. Linguist. Comput.* **29**, 412–421, 10.1093/llc/fqu020 (2014).
9. Citron, D. T. & Ginsparg, P. Patterns of text reuse in a scientific corpus. *Proc. Natl. Acad. Sci.* **112**, 25–30, 10.1073/pnas.1415135111 (2015). <https://www.pnas.org/content/112/1/25.full.pdf>.
10. Horbach, S. S. & Halffman, W. W. The extent and causes of academic text recycling or 'self-plagiarism'. *Res. Policy* **48**, 492–502 (2019).
11. Stein, B., Meyer zu Eißén, S. & Potthast, M. Strategies for Retrieving Plagiarized Documents. In Clarke, C., Fuhr, N., Kando, N., Kraaij, W. & de Vries, A. (eds.) *30th International ACM Conference on Research and Development in Information Retrieval (SIGIR 2007)*, 825–826, 10.1145/1277741.1277928 (ACM, New York, 2007).
12. Knoth, P. & Zdráhal, Z. CORE: three access levels to underpin open access. *D Lib Mag.* **18**, 10.1045/november2012-knoth (2012).

²<https://github.com/webis-de/arxiv21-stereo-scientific-text-reuse>

13. Lopez, P. & Romary, L. GROBID - information extraction from scientific publications. *ERCIM News* **2015** (2015).
14. Tang, J. *et al.* Arnetminer: extraction and mining of academic social networks. In Li, Y., Liu, B. & Sarawagi, S. (eds.) *Proceedings of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Las Vegas, Nevada, USA, August 24-27, 2008*, 990–998, 10.1145/1401890.1402008 (ACM, 2008).
15. Sinha, A. *et al.* An overview of microsoft academic service (MAS) and applications. In Gangemi, A., Leonardi, S. & Panconesi, A. (eds.) *Proceedings of the 24th International Conference on World Wide Web Companion, WWW 2015, Florence, Italy, May 18-22, 2015 - Companion Volume*, 243–246, 10.1145/2740908.2742839 (ACM, 2015).
16. Deutsche Forschungsgemeinschaft. DFG classification of scientific disciplines, research areas, review boards and subject areas. https://web.archive.org/web/20201126170513/https://www.dfg.de/download/pdf/dfg_im_profil/gremien/fachkollegien/amtsperiode_2016_2019/fachsystematik_2016-2019_en_grafik.pdf (2016). Accessed on 2021-05-27.
17. Potthast, M. *et al.* Overview of the 5th International Competition on Plagiarism Detection. In Forner, P., Navigli, R. & Tufis, D. (eds.) *Working Notes Papers of the CLEF 2013 Evaluation Labs*, vol. 1179 of *Lecture Notes in Computer Science* (2013).
18. Hagen, M., Potthast, M. & Stein, B. Source Retrieval for Plagiarism Detection from Large Web Corpora: Recent Approaches. In Cappellato, L., Ferro, N., Jones, G. & San Juan, E. (eds.) *Working Notes Papers of the CLEF 2015 Evaluation Labs*, vol. 1391 of *Lecture Notes in Computer Science* (2015).
19. Hagen, M., Potthast, M., Adineh, P., Fatehifar, E. & Stein, B. Source Retrieval for Web-Scale Text Reuse Detection. In Lim, E.-P. *et al.* (eds.) *26th ACM International Conference on Information and Knowledge Management (CIKM 2017)*, 2091–2094, 10.1145/3132847.3133097 (ACM, 2017).
20. Alshomary, M. *et al.* Wikipedia Text Reuse: Within and Without. In Azzopardi, L. *et al.* (eds.) *Advances in Information Retrieval. 41st European Conference on IR Research (ECIR 2019)*, vol. 11437 of *Lecture Notes in Computer Science*, 747–754, 10.1007/978-3-030-15712-8_49 (Springer, Berlin Heidelberg New York, 2019).
21. Broder, A. Z. On the resemblance and containment of documents. In Carpentieri, B., Santis, A. D., Vaccaro, U. & Storer, J. A. (eds.) *Compression and Complexity of SEQUENCES 1997, Positano, Amalfitan Coast, Salerno, Italy, June 11-13, 1997, Proceedings*, 21–29, 10.1109/SEQUEN.1997.666900 (IEEE, 1997).
22. Potthast, M. *et al.* Overview of the 4th International Competition on Plagiarism Detection. In Forner, P., Karlgren, J. & Womser-Hacker, C. (eds.) *Working Notes Papers of the CLEF 2012 Evaluation Labs* (2012).
23. Stammatos, E. Plagiarism detection using stopword n -grams. *J. Assoc. Inf. Sci. Technol.* **62**, 2512–2527, 10.1002/asi.21630 (2011).
24. Gienapp, L., Kircheis, W., Bjarne, S., Stein, B. & Potthast, M. Webis-STEREO-21 Corpus (Metadata Only Version). Zenodo, 10.5281/zenodo.5575285 (2021).
25. Gienapp, L., Kircheis, W., Bjarne, S., Stein, B. & Potthast, M. Webis-STEREO-21 Corpus (Full Version). Zenodo, 10.5281/zenodo.5575320 (2021).
26. Potthast, M. *et al.* Overview of the 5th International Competition on Plagiarism Detection. In Forner, P., Navigli, R. & Tufis, D. (eds.) *Working Notes Papers of the CLEF 2013 Evaluation Labs* (2013).
27. Various Authors. Vroniplag Wiki. <https://vroniplag.fandom.com/> (2021). Accessed: 2021-12-14.
28. Potthast, M. *et al.* Overview of the 6th International Competition on Plagiarism Detection. In Cappellato, L., Ferro, N., Halvey, M. & Kraaij, W. (eds.) *Working Notes Papers of the CLEF 2014 Evaluation Labs*, vol. 1180 of *Lecture Notes in Computer Science* (2014).
29. Forner, P., Karlgren, J. & Womser-Hacker, C. (eds.). *CLEF 2012 Evaluation Labs and Workshop – Working Notes Papers, 17-20 September, Rome, Italy* (CEUR-WS.org, 2012).
30. Forner, P., Navigli, R. & Tufis, D. (eds.). *CLEF 2013 Evaluation Labs and Workshop – Working Notes Papers, 23-26 September, Valencia, Spain* (CEUR-WS.org, 2013).
31. Cappellato, L., Ferro, N., Halvey, M. & Kraaij, W. (eds.). *Working Notes Papers of the CLEF 2014 Evaluation Labs*, CEUR Workshop Proceedings (CEUR-WS.org, 2014).

Author Contributions Statement

L.G., M.P., W.K. and B.St. conceived the experiment. L.G. conducted the experiment. W.K. and B.Si. assembled the metadata and carried out the plain text extraction. All authors wrote and reviewed the manuscript.

Competing Interests

The authors declare no competing interests.

Figures & Tables

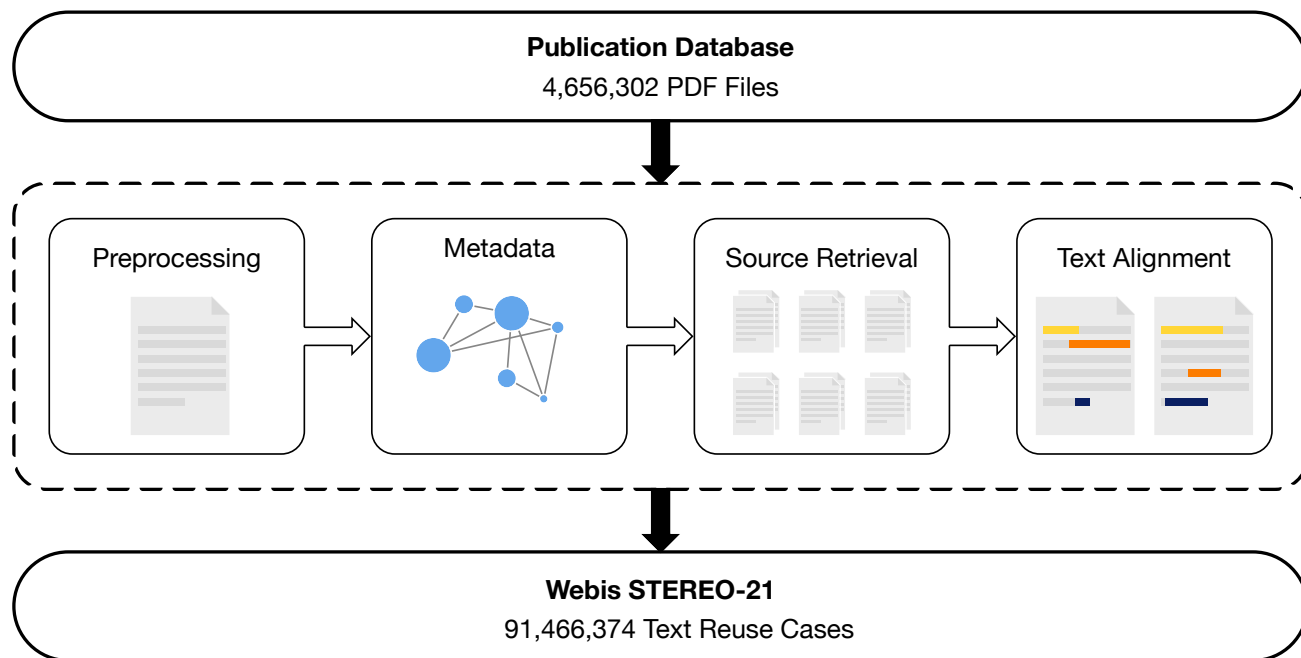


Figure 1. Schematic overview of the text reuse detection pipeline. Each document is pre-processed and supplied with metadata. Source retrieval identifies document pairs with local similarities, and alignment is applied to identify reuse cases between those.

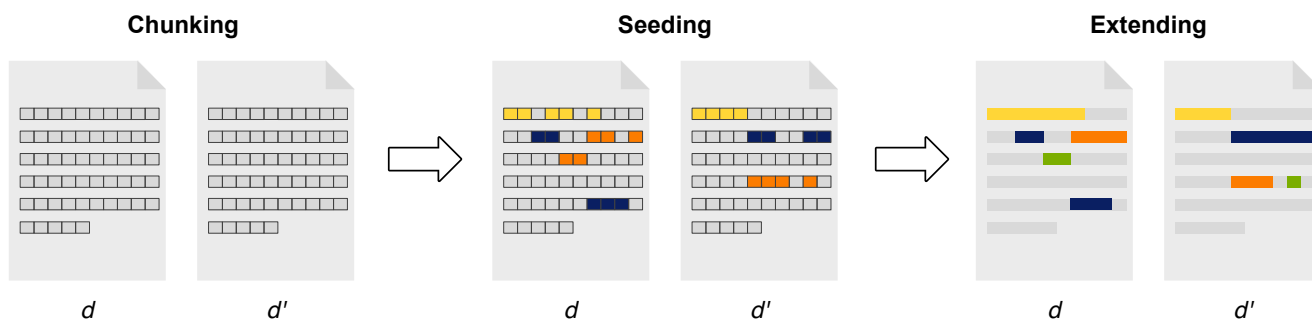


Figure 2. Schematic overview of the text alignment process. Each document is divided into chunks, matching chunks are identified between documents, and matches are extended to whole reuse cases.

Table 1. Data record overview for case data. *Key* denotes the top-level JSON key by which fields are identified, *Description* provides an explanation of the contained data, and *Type* denotes the data type; “optional” indicates fields that can be empty. The text fields are included only in the full version of the corpus.

	Key	Description	Type	
	id	Unique identifier for this case, in UUID format.	String	
Publication A	Text	text_a	Matched text in publication A.	String (full only)
		before_a	Text context before the matched text, 100 characters long.	String (full only)
		after_a	Text context after the matched text, 100 characters long.	String (full only)
	Locator	begin_a	Start location of matched text, measured as character offset.	Integer
		end_a	End location of matched text, measured as character offset.	Integer
		doc_length_a	Total length of publication A in characters.	Integer
	Metadata	doi_a	DOI identifier for publication A.	String
		year_a	Publication year for publication A.	Integer (optional)
		field_a	Field(s) of study for publication A.	String Array (optional)
		area_a	Scientific area(s) for publication A.	String Array (optional)
discipline_a		Scientific discipline(s) for publication A.	String Array (optional)	
Publication B	Text	text_b	Matched text in publication B.	String (full only)
		before_b	Text context before the matched text, 100 characters long.	String (full only)
		after_b	Text context after the matched text, 100 characters long.	String (full only)
	Locator	begin_b	Start location of matched text, measured as character offset.	Integer
		end_b	End location of matched text, measured as character offset.	Integer
		doc_length_b	Total length of publication B in characters.	Integer
	Metadata	doi_b	DOI identifier for publication B.	String
		year_b	Publication year for publication B.	Integer (optional)
		field_b	Field(s) of study for publication B.	String Array (optional)
		area_b	Scientific area(s) for publication B.	String Array (optional)
discipline_b		Scientific discipline(s) for publication B.	String Array (optional)	

Table 2. Data record overview for publication data. *Key* denotes the top-level JSON key fields are identified by, *Description* provides an explanation of the contained data, and *Type* denotes the data type; “optional” indicates fields that can be empty.

	Key	Description	Type
Metadata	doi	DOI identifier of the publication.	String
	doc_length	Total length of the publication in characters.	Integer
	year	Publication year of the publication.	String Array (optional)
	field	Field of study of the publication.	String Array (optional)
	area	Scientific area of the publication.	String Array (optional)
	discipline	Scientific discipline of the publication.	String Array (optional)

Table 3. Precision, Recall, and $F_{0.5}$ score of competing alignment approaches by teams participating at the PAN Shared Tasks, taken from Potthast et al. (2013)²⁶ and Potthast et al. (2014)²⁸ and sorted descending by precision. The listed approaches are described in detail in the PAN Workshop proceedings²⁹⁻³¹. Complexity denotes the runtime complexity of the approach; Data Specificity denotes the degree of PAN-specific optimizations reducing transferability to other data domains, i.e. ensemble methods, trained approaches, feature extraction, or special handling of corpus characteristics.

PAN13 Evaluation Corpus					
Team	Precision↓	Recall	$F_{0.5}$	Complexity	Data Specificity
Glinos (2014)	0.96	0.79	0.92	$O(nm)$	Medium
Jayapal (2012)	0.95	0.22	0.57	$O(n+m)$	Low
Nourian (2013)	0.95	0.43	0.76	—	—
Gross (2014)	0.93	0.77	0.89	$O(nm)$	Low
Alvi (2014)	0.93	0.55	0.82	$O(n+m)$	Medium
Ours	0.93	0.46	0.77	$O(n+m)$	Low
Baseline (2014)	0.93	0.34	0.69	$O(n+m)$	Low
Palkovskii (2014)	0.92	0.83	0.90	$O(nm)$	High
Torrejon (2014)	0.90	0.77	0.87	$O(nm)$	High
Gillam (2014)	0.89	0.17	0.48	$O(nm)$	High
Oberreuter (2014)	0.89	0.86	0.88	—	—
Oberreuter (2012)	0.89	0.77	0.86	—	—
Gillam (2012)	0.89	0.27	0.61	$O(nm)$	High
Torrejon (2013)	0.89	0.76	0.86	$O(nm)$	High
Gillam (2013)	0.88	0.26	0.60	$O(nm)$	High
Jayapal (2013)	0.88	0.38	0.70	$O(n+m)$	Low
Sanchez-Perez (2014)	0.88	0.88	0.88	$O(nm)$	High
Kueppers (2012)	0.87	0.51	0.76	$O(nm)$	Low
Shrestha (2013)	0.87	0.74	0.84	$O(nm)$	Medium
Saremi (2013)	0.87	0.77	0.85	—	—
Shrestha (2014)	0.86	0.84	0.86	$O(nm)$	Medium
Kong (2012)	0.85	0.82	0.84	$O(nm)$	Low
Suchomel (2012)	0.84	0.65	0.79	$O(n+m)$	Medium
Kong (2014)	0.84	0.81	0.83	$O(nm)$	Low
Kong (2013)	0.83	0.81	0.83	$O(nm)$	Low
Torrejon (2012)	0.83	0.75	0.81	$O(nm)$	High
Palkovskii (2013)	0.82	0.54	0.74	$O(nm)$	High
Palkovskii (2012)	0.82	0.76	0.81	$O(nm)$	High
Abnar (2014)	0.77	0.61	0.73	$O(nm)$	Medium
Suchomel (2013)	0.73	0.77	0.74	$O(n+m)$	Medium
Sanchez-Vega (2012)	0.40	0.56	0.42	$O(nm)$	Medium

Table 4. Precision, Recall, and $F_{0.5}$ of the text alignment component per obfuscation strategy and on the complete evaluation corpus.

PAN13 Evaluation Corpus			
	Precision	Recall	$F_{0.5}$
No Obfuscation	0.88	0.90	0.88
No Plagiarism	1.00	1.00	1.00
Random Obfuscation	0.90	0.11	0.37
Summary Obfuscation	0.99	0.10	0.36
Translation Obfuscation	0.88	0.16	0.46
Entire Corpus	0.93	0.46	0.77