

Improving Academic Plagiarism Detection for STEM Documents by Analyzing Mathematical Content and Citations

Norman Meuschke^{1,2}, Vincent Stange², Moritz Schubotz¹, Michael Kramer², Bela Gipp^{1,2}

¹University of Wuppertal, Germany (last@uni-wuppertal.de)

²University of Konstanz, Germany (first.last@uni-konstanz.de)

ABSTRACT

Identifying academic plagiarism is a pressing task for educational and research institutions, publishers, and funding agencies. Current plagiarism detection systems reliably find instances of copied and moderately reworded text. However, reliably detecting concealed plagiarism, such as strong paraphrases, translations, and the reuse of nontextual content and ideas is an open research problem. In this paper, we extend our prior research on analyzing mathematical content and academic citations. Both are promising approaches for improving the detection of concealed academic plagiarism primarily in Science, Technology, Engineering and Mathematics (STEM). We make the following contributions: i) We present a two-stage detection process that combines similarity assessments of mathematical content, academic citations, and text. ii) We introduce new similarity measures that consider the order of mathematical features and outperform the measures in our prior research. iii) We compare the effectiveness of the math-based, citation-based, and text-based detection approaches using confirmed cases of academic plagiarism. iv) We demonstrate that the combined analysis of math-based and citation-based content features allows identifying potentially suspicious cases in a collection of 102K STEM documents. Overall, we show that analyzing the similarity of mathematical content and academic citations is a striking supplement for conventional text-based detection approaches for academic literature in the STEM disciplines. The data and code of our study are openly available at <https://pur1.org/hybridPD>

1 INTRODUCTION

Academic plagiarism (AP) is *'the use of ideas, concepts, words, or structures without appropriately acknowledging the source to benefit in a setting where originality is expected'* [8, 9]. Forms of AP range from copying content (*copy&paste*) to reusing slightly modified content, e.g., interweaving text from multiple sources, to heavily concealing content reuse, e.g., by paraphrasing or translating text, and reusing data or ideas without proper attribution [42].

The easily recognizable copy&paste-type AP is more prevalent among students [22], while concealed AP is more characteristic of researchers, who have strong incentives to avoid detection [2]. Plagiarized research publications can have a severe negative impact by distorting the mechanisms for tracing and correcting research results and causing inefficient allocations of research funds [9]. Therefore, detecting concealed AP in research publications is a pressing problem affecting many stakeholders, including research institutions, academic publishers, digital library providers, funding agencies, and of course other researchers [28].

As we present in Section 2, many plagiarism detection (PD) approaches have been proposed that employ lexical, semantic, syntactical, or cross-lingual text analysis. These approaches reliably

detect copied or moderately altered forms of AP; some approaches can also identify paraphrased and translated text. However, current approaches are computationally expensive and thus require a computationally less demanding selection of candidate documents before their application. The performance of methods that analyze textual features to retrieve candidate documents has reached a plateau. Therefore, the candidate retrieval step currently limits the effectiveness of PD approaches to detect concealed forms of AP.

Prior research (cf. Section 2) showed that approaches that analyze nontextual content features, such as academic citations, images, and mathematical content, complement the many text analysis approaches to improve the identification of concealed forms of AP. Nontextual content features in academic documents are a valuable source of semantic information that is largely independent of natural language text. Considering these sources of semantic information for similarity analysis raises the effort plagiarists must invest for obfuscating reused content [9, 24].

We extend the research on analyzing nontextual content features for PD by devising a novel PD approach that combines the analysis of mathematical content with the analysis of academic citations. We structure the presentation of our contributions as follows. Section 2 presents an overview of prior research on PD to show the benefit of analyzing nontextual content features for this purpose. Section 3 explains the conceptual design and technical realization of the math-based, citation-based and text-based PD approaches we investigate. Section 4 presents our test collection and describes the methodology of our evaluation. Section 5 compares the performance of the three PD approaches using confirmed cases of AP. Section 6 presents the results of an exploratory study that investigates the ability of the novel PD approach to discover so far unknown cases of AP. Section 7 concludes the paper and presents future work.

2 RELATED WORK

External plagiarism detection is an information retrieval task with the objective of comparing an input document to a large collection and retrieving all documents exhibiting similarities above a threshold [39]. External PD approaches typically employ a two-stage process consisting of candidate retrieval and detailed analysis [23, 39]. In the candidate retrieval stage, the approaches employ computationally efficient retrieval methods to limit the collection to a set of documents that may have been the source for the content in the input document. In the detailed analysis stage, the systems perform computationally more demanding analysis steps to substantiate the suspicion and to align components in the input document and potential source documents that are similar [2, 23].

Text retrieval research has yielded mature systems that reliably detect copied or moderately altered text in an input document and retrieve its source if the source is part of the system's reference

collection. For the candidate retrieval stage, such systems typically employ character-gram or word-gram fingerprinting [30, 41] or term-based vector space models [19]. For the detailed analysis stage, such systems often perform exhaustive string comparisons [41] or computationally more efficient text alignment [30]. Text alignment approaches typically use matching strings as seeds, which the procedures extend and then filter using heuristics [36].

To detect monolingual paraphrases, researchers have proposed approaches that analyze semantic and syntactic features, mostly during the detailed analysis stage of the PD process. Several researchers adapted semantic text analysis methods, such as Singular Value Decomposition [5], Latent Semantic Analysis [38] and Explicit Semantic Analysis [27] for the PD use case. Other PD approaches employ linguistic resources, such as WordNet¹, to analyze exactly matching and semantically related words [16]. Some works combine the analysis of word-based semantic similarity with an analysis of similarity in semantic arguments derived using Semantic Role Labeling [31]. Other approaches employ part-of-speech tagging to also compare the syntactic structure of documents [16].

To detect cross-lingual (CL) plagiarism, researchers have proposed approaches that leverage lexical similarities of languages, e.g., CL character n -gram matching, employ thesauri, parallel corpora, and machine translation followed by a mono-lingual analysis [3].

The extensive research on semantic and syntactic methods for monolingual and cross-lingual PD has yielded approaches that are highly effective for the detailed analysis stage. To our knowledge, Gupta et al. reported the highest retrieval effectiveness for the detailed analysis of realistically obfuscated plagiarism [16]. The authors analyzed artificially created plagiarism in the Webis Text Reuse Corpus 2012 (Webis-TRC-2012) [33], which is a standard corpus to evaluate PD systems as part of the PAN Workshop series². For manually paraphrased instances of simulated plagiarism, Gupta et al. reported the precision $P = 0.81$, recall $R = 0.80$, and $F_1 = 0.80$.

While these results for the detailed analysis stage are promising, significant potential for improvement remains regarding the complete retrieval process. Semantic and syntactic analysis approaches like the one of Gupta et al. are computationally expensive and thus require a prior candidate retrieval step. Kong et al. achieved the best recall ($R = 0.65$) for the candidate retrieval task in all four PAN workshops (2012-2015) that evaluated research contributions for this task [20]. The organizers of the PAN workshop series noted that the workshops had reached a stable production phase, in which the submitted approaches no longer exhibited "[...] *real innovations with respect to recall-oriented source retrieval.*" [17].

For cross-lingual PD, the candidate retrieval stage likewise seems to present an upper bound for the otherwise higher effectiveness of the analysis methods in the detailed analysis stage. Ehsan et al. reported an approach that achieved an F_1 -score of 0.87 ($P = 0.93$, $R = 0.82$) for the detailed analysis of cross-lingual plagiarism in the Webis-TRC-2012 [7]. For candidate retrieval, Ehsan and Shakeri reported a maximum recall of $R = 0.75$ using the same corpus [6]. These results suggest that the candidate retrieval step, which is necessary to enable applying semantic, syntactic and cross-lingual PD approaches, currently limits the effectiveness of the PD approaches.

PD approaches that analyze nontextual content features in academic documents are a promising complement to the variety of text analysis approaches both for the candidate retrieval and the detailed analysis stages of the PD process. In prior research, we showed that an analysis of in-text citation patterns in academic documents, i.e., identical citations occurring in proximity or in a similar order within two documents, can identify concealed forms of AP in real-world, large-scale collections [9–12]. This approach is computationally efficient enough to be applied in the candidate retrieval stage [9, 12, 24]. Pertile et al. confirmed the positive effect of combining citation and text analysis and devised a hybrid approach using machine learning [32]. The benefits of citation-based PD are twofold. First, citations encode semantic information that cannot easily be substituted, since leaving out citations to relevant prior work or citing sources that are not relevant to the topic would likely raise the suspicion of expert peer reviewers. Second, while citations are independent of natural language text, analyzing in-text citation patterns can indicate shared structural and semantic similarity among texts. Assessing this semantic and structural similarity using citation patterns requires significantly less computational effort than approaches for semantic and syntactic text analysis.

We also showed that analyzing image similarity in academic documents, e.g., the similarity of figures and plots, improves the detection capabilities for concealed forms of AP [25].

In a recent short paper [26], we extended the idea of citation-based PD. We proposed that mathematical expressions share many characteristics of academic citations and hence are promising nontextual content features to be considered when searching for concealed forms of AP. Similar to academic citations, mathematical expressions are essential components of academic documents in the Science, Technology, Engineering and Mathematics fields. Furthermore, mathematical expressions are independent of natural language text and contain rich semantic information. Additionally, some STEM disciplines, such as mathematics and physics, are known for their comparably sparse use of academic citations [29]. A citation-based analysis alone is, therefore, less likely to reveal potentially suspicious content similarity for these disciplines.

Our piloting study [26] investigated measures to quantify the similarity of mathematical content features in the detailed analysis stage of the PD process. Given the infancy of the research area, we evaluated the suitability of comparing basic presentational features of mathematical expressions, i.e., elements of mathematical notation, such as identifiers, numbers, operators, and special symbols. Our goal was to identify the type of similar mathematics that we had observed in confirmed cases of AP, which we collected, e.g., by reviewing journal retractions. We embedded the retracted test documents together with their sources in the dataset of the NTCIR-11 Math task (105,120 arXiv documents) [1]. We performed pairwise comparisons of all documents in the dataset (detailed analysis approach) and evaluated similarity measures that consider identifiers, numbers, operators, and combinations thereof as the features. The best performing approach, a set-based comparison of the frequency of mathematical identifiers, retrieved eight of ten test cases at the top rank and achieved a mean reciprocal rank of 0.86.

¹<https://wordnet.princeton.edu/>

²<http://pan.webis.de/>

The present paper extends our pilot study [26] by making four contributions: i) devising a candidate retrieval stage that analyzes mathematical expressions, citations, and textual features; ii) proposing new math-based similarity measures for the detailed analysis stage that consider the order of mathematical content features; iii) comparing the effectiveness of the math-based, citation-based and text-based PD approaches using the confirmed cases of AP gathered for our pilot study; iv) analyzing math-based and citation-based content features to discover potentially suspicious cases of document similarity in a large-scale dataset of STEM documents.

3 DETECTION APPROACH

Figure 1 gives an overview of our system HyPlag [28] that implements the analysis of math-based, citation-based, and text-based content features. HyPlag allows the combination of the analysis approaches as part of a hybrid detection process that consists of five stages: preprocessing, indexing, candidate retrieval, detailed analysis, and human inspection. We describe the stages hereafter.

3.1 Preprocessing

HyPlag preprocesses input documents in two steps. In the first step, the system converts the documents to a unified XML-based document format used for the second preprocessing step. Our unified document format uses a subset of the TEI standard³ defined by the information extraction tool GROBID⁴ to represent in-text citations and bibliographic references. Additionally, the unified document format employs a subset of the Mathematical Markup Language (MathML)⁵ to represent mathematical formulae.

For this study, we processed documents in two formats: PDF (confirmed cases of AP) and LaTeX source code (NTCIR-11 MathIR Task dataset). We used GROBID to obtain bibliographic references from documents in both formats because the tool achieved excellent results for extracting header metadata, citations, and references [4, 40]. Since GROBID cannot recognize mathematical formulae, we semi-automatically invoked InftyReader⁶ to convert the PDFs for confirmed cases of AP to an intermediate LaTeX format. InftyReader is currently the most commonly used OCR-based recognition system for mathematical content [18]. While the tool typically achieves a recall of at least 0.90, the precision can be as low as approx. 0.15 in some scenarios [18]. To prevent bias from recognition errors, we manually checked and corrected the LaTeX output of InftyReader.

We employed LaTeXML⁷ to convert LaTeX documents, i.e., the documents in the NTCIR-11 dataset and the PDF that InftyReader converted to LaTeX, to the unified document format. The LaTeXML library offers mathematical content conversion from LaTeX source code to a MathML representation. To enable conversion to our unified document format, we contributed an XSL style sheet that transforms LaTeXML's native output to TEI. The new conversion option has been included in the LaTeXML distribution⁸.

³<http://tei-c.org/>

⁴<https://github.com/kermitt2/grobid>

⁵<https://w3.org/Math/>

⁶<http://www.inftyreader.org/>

⁷<https://dlmf.nist.gov/LaTeXML>

⁸<https://github.com/bruce miller/LaTeXML/blob/master/lib/LaTeXML/resources/XSLT/LaTeXML-tei.xsl>

To recognize in-text citations, LaTeXML requires the use of LaTeX tags, such as `\cite{}`. Many documents in the dataset do not contain such markup but state in-text citations as plain-text. In such cases, our preprocessing pipeline did not recognize the in-text citations. Additionally, many documents do not use in-text citations at all but only reference items in the bibliography. Due to both errors, the number of unique in-text citations for 68,743 documents (67% of the dataset) is smaller than the number of references.

In the second preprocessing step, HyPlag splits the unified document format into separate data structures holding plain text, mathematical formulae, in-text citations, and bibliographic references. To extract plain text, the system removes all XML structures, images, formulae, and formatting instructions. Formulae in Content MathML are extracted as they are. In-text citations are linked to the corresponding reference entries in the bibliography; reference entries are split into author, title, and venue fields.

3.2 Indexing

In the indexing stage, HyPlag stores into an Elasticsearch⁹ index the following data extracted from the preprocessed documents: *Document metadata*: title, authors, publication date and filename.

Mathematical features: the sequence of all mathematical identifiers in the order of their occurrence in the document, and the unordered histogram of the occurrence frequencies of identifiers, i.e., how often an identifier occurs in the document. We focused on analyzing identifiers since they achieved the best retrieval effectiveness in our pilot study [26]. We used the MathML `<formula>` element to extract formulae. Additionally, we used `<ci>` elements in MathML formulae to extract identifiers in a formula.

Citation features: bibliographic references and in-text citations. The system consolidates the data about referenced documents by comparing the title and author names extracted from reference strings to the data of previously indexed documents while accounting for minor spelling variations utilizing the Levenshtein distance.

Textual features: full text and text fingerprints formed by chunking the document into word 3-grams and applying probabilistic chunk selection (average chunk retention rate 1/16). To realize the text fingerprinting approach, we adapted the Sherlock¹⁰ tool.

The indexing process is identical for all documents, i.e., documents that ought to be analyzed need to be indexed first.

3.3 Candidate Retrieval

In the candidate retrieval stage, the system queries the index using mathematical identifiers, in-text citations, and text fingerprints extracted from an input document to retrieve a set of candidate documents for the subsequent detailed analysis.

To retrieve candidate documents, we employed "Lucene's practical scoring function" implemented in the Elasticsearch server as a computationally efficient, well-established heuristic. The scoring function combines a tf/idf weighted vector space model with a Boolean retrieval approach¹¹.

⁹<https://www.elastic.co/products/elasticsearch>

¹⁰The tool's website went offline recently. The source code and documentation are still available via the Web archive: <https://web.archive.org/web/20180219024142/http://web.it.usyd.edu.au/~scilect/sherlock/>

¹¹documentation of the scoring function: <https://www.elastic.co/guide/en/elasticsearch/guide/1.x/practical-scoring-function.html>

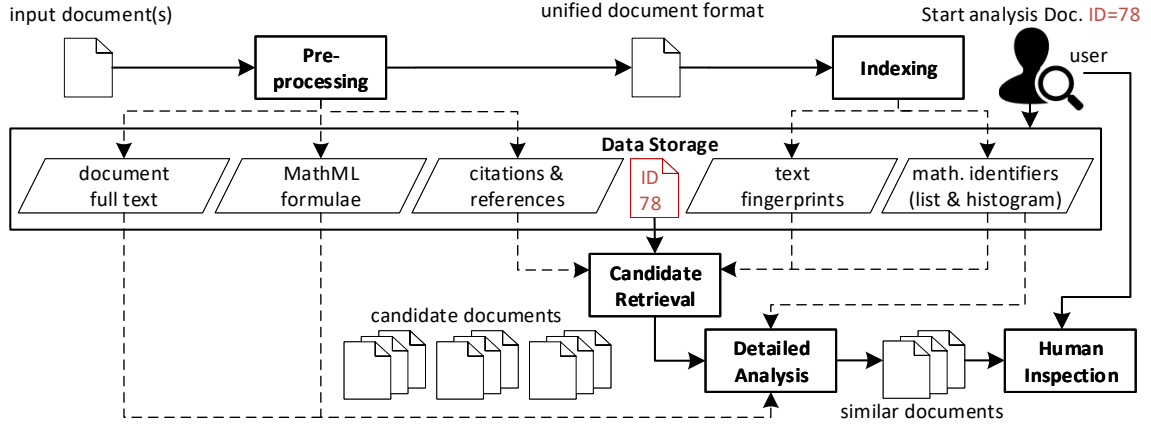


Figure 1: Overview of the hybrid plagiarism detection approach.

We performed three queries, each retrieving the 100 documents with the highest relevance scores. For the citation-based and text-based retrieval of candidate documents, in-text citations and, respectively, text fingerprints of the input document represented the terms of the query. Analogously, indexed documents were represented by their sets of in-text citations and text fingerprints. We used the default parameters of Lucene’s scoring function.

For the math-based retrieval of candidate documents, the set of mathematical identifiers occurring in a document was the query. The indexed documents were represented by the sequence of mathematical identifiers in a document, i.e., identifiers can occur more than once. However, using Lucene’s default parameters for the relevance scoring yielded unsatisfactory results in the case of mathematical features. This finding is in line with research by Sojka and Liška [37]. Similar to Sojka and Liška, we found that query terms, i.e., mathematical identifiers, should be given additional weight for multiple occurrences. Therefore, we set the boost value $\text{boost}(t)$ for the term t in the query, i.e., individual identifiers, to the number of occurrences of the term (identifier) in the query document.

Since we sought to investigate the effectiveness of the math-based, the citation-based, and the text-based PD approach independently of each other, we did not consolidate the three sets of 100 candidate documents each retrieved by the three queries.

3.4 Detailed Analysis

In the detailed analysis stage, HyPlag compares the input document(s) to all documents in each of the three sets of 100 candidate documents retrieved in the previous stage. For each document comparison, HyPlag computes the following similarity measures.

3.4.1 Math-based similarity measures. We only computed math-based similarity scores for document pairs that share 20 or more identifiers to prevent high similarity scores resulting from a few shared identifiers, such as the occurrence of x and y . For documents that meet this threshold, we computed three similarity measures.

First, we computed the similarity of *frequency histograms of mathematical identifiers (Histo)*, which performed best in our pilot study [26]. The measure quantifies the similarity of two documents d and d' as the difference in the relative occurrence frequencies of

identifiers f_i in d and d' according to Equation (1).

$$s(d, d') = 1 - \frac{\sum_{i \in I} |f_{i,d} - f_{i,d'}|}{\sum_{i \in I} \max(f_{i,d}, f_{i,d'})}. \quad (1)$$

The Histo score reflects the global overlap of identifiers in two documents. The measure is most suitable for documents with comparable numbers of identifiers. Typically, this requirement is not met if the two documents strongly differ in length.

In addition to the set-based, order-agnostic Histo measure proposed in our pilot study [26], we devised two new similarity measures that consider the order of mathematical identifiers. We did not evaluate the influence of the sequential similarity of features in our pilot study. The two new measures consider the *Longest Common Subsequence of Identifiers (LCIS)* and the set of *Greedy Identifier Tiles (GIT)* for score computation. The longest common subsequence and greedy tiling algorithms are well-established approaches to identify sequential patterns and have been applied successfully for the text-based [21] and citation-based [10] detection of academic plagiarism, as well as for source code plagiarism detection [35].

The longest common subsequence of features, e.g., characters, citations, or mathematical identifiers, is the maximum number of features that match in both documents in the same order but not necessarily in a contiguous block. Like Histo, the LCIS measure quantifies the global similarity of documents. We compute the similarity score $s_{\text{LCIS}}(d, d') = |L(d, d')| I_d^{-1}$ that represents the number of identifiers in the query document I_d that are part of the longest common identifier sequence whose length is given by L .

Greedy tiles are the set of all individually longest blocks of shared features in identical order that cannot be extended to the left or right without encountering a non-matching feature [43]. Greedy tiles are well-suited to identify confined regions with high similarity [10]. We computed the similarity of two documents using the GIT approach as $s_{\text{GIT}}(d, d') = |T_l| I_d^{-1}$, where T_l is the set of tiles with a length greater or equal to 5 matching identifiers and I_d is the total number of identifiers in the query document. In other terms, the score quantifies the number of identifiers in the query document that are part of identifier tiles with a minimum length of five.

3.4.2 Text-based similarity assessment. For the detailed text-based analysis, we used the *Encoplot algorithm (Enco)* developed by Grozea et al. [14]. Encoplot is an efficient character 16-gram comparison that achieves $O(n)$ time-complexity by ignoring repeated matches. The similarity score is the ratio of shared character 16-grams to all 16-grams of the shorter document.

3.4.3 Citation-based similarity assessment. We used three approaches that proved effective in our prior research[9–12, 24].

Bibliographic Coupling (BC), quantifies the fraction of shared bibliographic references. The similarity score is calculated as $s(d, d') = |R_d \cap R_{d'}| (R_d \cup R_{d'})^{-1}$, where R_d and $R_{d'}$ are the sets of references in the query and the comparison document. Like the Histo measure for mathematical features, BC is an order-agnostic measure that quantifies the global citation-based similarity of documents. It achieves high scores if documents with similar numbers of references share a significant fraction of those references.

The *Longest Common Citation Sequence (LCCS)* and *Greedy Citation Tiling (GCT)* measures follow the same idea as LCIS and GIT but consider in-text citations instead of mathematical identifiers. The similarity scores for the LCCS and GCT approaches are calculated analogously to the scores for LCIS and GIT, with the exception that the minimum length for citation tiles is two matching citations opposed to five matching identifiers for GIT. We only computed the three citation-based similarity measures if both documents we compared contained at least three bibliographic references.

3.5 Human Inspection

We used HyPlag’s web-based frontend to visualize content similarity for inspection. For details on the visualizations, see [28].

4 EVALUATION

To ensure the reproducibility of our research, the data and code of our study are available at <https://purl.org/hybridPD>.

4.1 Dataset

To achieve comparability to our prior research, we reused the dataset of our previous experiments [26]. The dataset consists of ten cases that we selected after manually reviewing 44 research publications in STEM disciplines that have been officially retracted for plagiarism and involve mathematical content. We restricted the dataset to ten cases for three reasons. First, we chose cases from research fields within our area of expertise to enable us to assess the severity of identified similarities. Second, we selected cases that are most representative of the types of mathematical similarity we observed. Third, the effort required for checking the output of InftyReader and correcting incorrectly recognized mathematical expressions prevented us from converting more cases.

We chose using real cases of AP over creating artificial test cases, although gathering and converting real cases is time-consuming and thus resulted in a smaller dataset. The reason is that we see the ability to identify real cases of AP committed by researchers who are experts in their fields and have a strong incentive to avoid detection as the ultimate performance test for any PD approach. Therefore, we see real cases of AP as best suited to devise and evaluate the novel hybrid detection approach.

Table 1: Overview of content features in our dataset.

Features	Total	Avg. per doc.
references	2,201,094	21
unique references	1,445,059	–
citations	3,068,865	30
text fingerprints	26,539,276	256
formulae	52,271,908	504
math. identifiers	156,706,600	1513

We embedded the ten retracted documents and the ten source documents for our test cases in the topically related NTCIR-11 MathIR Task dataset. The NTCIR dataset, which is available for research purposes [1], consists of 105,120 scientific papers in LaTeX format from computer science, mathematics, physics, and statistics that were published in the arXiv preprint repository¹².

Using our preprocessing pipeline (cf. section 3.1), we converted all LaTeX source files of the NTCIR dataset and of our test cases to HyPlag’s unified document format. We excluded 2,616 documents, for which LaTeXML or our TEI parser encountered critical processing errors. Approximately one-third of the remaining documents did not contain markup for authors and title. To achieve the best possible data quality, we used the arXiv API¹³ to obtain author and title information for all documents instead of extracting the information from the LaTeX source files. For 6,770 documents we were unable to extract bibliographic references due to missing markup. Since the arXiv API does not offer the data of bibliographic reference, we indexed these documents without reference data.

Table 1 shows the number of content features we obtained for the final dataset of 102,524 documents. The numbers confirm that this collection of STEM documents contains a significantly higher number of mathematical formulae (52M) than academic citations (3M). Therefore, analyzing both mathematical formulae and citations is more promising in these disciplines than analyzing citations alone. The formulae contain more than 156M identifiers, which the system grouped into 4,063,354 identifier histogram entries. On average, documents contained 70 different mathematical identifiers.

4.2 Investigations

To evaluate the effectiveness of the PD approaches, we performed two conceptually different investigations. The first investigation reflects the typical scenario in external PD, i.e., checking an input document for similarity to documents in a collection. We submitted the retracted paper for each test case to our system HyPlag. For each query document, the system used the math-based, citation-based, and text-based retrieval heuristics (cf. Section 3.3) to retrieve three sets of 100 candidate documents each. In the subsequent detailed analysis stage, each query document was compared to all the candidate documents in the three sets without consolidating the sets. Section 5 presents the results of this investigation.

The second investigation assesses the effectiveness of combining the math-based and citation-based similarity measures to discover so far unknown cases of potentially suspicious document similarity.

¹²<http://www.arxiv.org>

¹³<https://arxiv.org/help/api/>

We submitted each of the $N = 102,524$ documents in our dataset to HyPlag. We retrieved the three sets of candidate documents by applying the math-based, citation-based, and text-based retrieval heuristics for all N documents. Opposed to the evaluation of confirmed plagiarism cases, we formed the union of the sets to enable the exploration of approaches that combine measures. In the detailed analysis stage, we compared each of the N documents in the dataset to its consolidated set of candidate documents C . We manually examined the retrieved documents with the highest similarity scores. Section 6 presents the results of this investigation.

5 RESULTS CONFIRMED CASES OF AP

To not establish a link between a paper on academic plagiarism detection and legitimate research papers, i.e., the source documents of our test cases and unsuspecting documents retrieved in our experiments, we do not cite the documents we discuss hereafter. However, all documents are accessible via our data and code repository.

5.1 Candidate Retrieval

Table 2 shows the effectiveness of the candidate retrieval approaches. Plus signs (+) in the table indicate that HyPlag retrieved the source document among the 100 candidate documents when the retracted document for each of the ten test cases (C1 ... C10) was the query. Minus signs (-) indicate that an analysis approach did not retrieve the source document among the candidate documents. The right-most column in Table 2 shows the recall of the approaches.

Table 2: Recall for candidate retrieval stage.

	C1	C2	C3	C4	C5	C6	C7	C8	C9	C10	R
Math	+	+	+	-	-	-	+	+	+	+	0.7
Cit.	+	+	-	+	+	+	+	+	+	+	0.9
Text	+	+	+	+	+	+	-	+	+	+	0.9

Both the citation-based and the text-based approaches achieved a recall of 0.9; the math-based approach achieved a recall of 0.7. Notably, the three approaches failed to retrieve the source document among the candidates for distinct sets of test cases. Combining the three sets of candidate documents would result in a perfect recall.

5.2 Detailed Analysis

To quantify the effectiveness of the similarity measures employed in the detailed analysis stage, we performed a score-based assessment and a rank-based assessment.

5.2.1 Score-based assessment. This assessment determines which scores are significant, i.e., potentially suspicious, for our similarity measures and dataset. To our knowledge, no study (including our pilot study [26]) has quantified the mathematical similarity that can be expected by chance to derive a significance threshold.

To establish significance thresholds for the scores of all similarity measures, we analyzed a random sample of 1 million document pairs as follows. We randomly picked two documents from the dataset. If the chosen documents had (a) common author(s) or if one of the documents cited the other, we discarded the pair. We continued the process until reaching the number of 1M document

Table 3: Significance thresholds for similarity measures.

	Histo	LCIS	GIT	BC	LCCS	GCT	Enco
s	$\geq .56$	$\geq .76$	$\geq .15$	$\geq .13$	$\geq .22$	$\geq .10$	$\geq .06$

pairs. The selection criteria ought to eliminate document pairs that exhibit high content similarity for likely legitimate reasons, i.e., reusing own work and referring to the work of others with due attribution. Our goal was to estimate an upper bound for similarity scores that likely result from random feature matches. To do so, we manually assessed the topical relatedness of the top-ranked document pairs within the random sample of 1M documents for each similarity measure. We picked as the significance threshold for a similarity measure the rank of the first document pair for which we could not identify a topical relatedness. Table 3 shows the significance scores we derived using this procedure.

Figure 2 shows the distribution of the similarity scores s (vertical axis) computed using each similarity measure for the random sample of 1M documents. Large horizontal bars shaded in blue indicate the median score; small horizontal bars shaded in grey mark the minimum and maximum scores; small horizontal bars shaded in green indicate the significance thresholds for each measure (cf. Table 3). The grey shapes in the chart show the smoothed probability density functions of the score frequencies, which were generated by applying a kernel-based density estimation. Red dots in the plot indicate the similarity scores of test cases for which the respective measure was applied, i.e., if the document pairs contained enough features to compute a score (cf. Section 3.4).

As shown in Figure 2, the probability density function (PrDF) of Histo is symmetrical while the PrDF for any other measure is negatively skewed, i.e., exhibits the highest frequencies at lower scores. The stronger the PrDF of scores is negatively skewed, the more selective the measure is. For the math-based similarity measures (Histo, LCIS, GIT), considering the order of identifiers strongly increases the selectivity of the measures. The PrDF for the order-agnostic Histo measure is symmetrical. The PrDF of scores for the LCIS measure, which leniently considers the order of identifiers in the entire document, is slightly skewed towards lower values, while the PrDF for the GIT measure, which focuses on considering identifier order, is strongly skewed towards lower values.




Figure 2: Similarity scores in 1M random document pairs.

Table 4: Retrieval effectiveness of detection approaches for confirmed cases of plagiarism.

Case	Math						Citation						Text				
	Histo		LCIS		GIT		BC			LCCS			GCT			Enco	
	r	s	r	s	r	s	r	s	s^*	r	s	s^*	r	s	s^*	r	s
C1	1	<u>.68</u>	1	.40	1	<u>.21</u>	1	.06	<u>.15</u>	1	.06	.10	-	-	.04	1	<u>.13</u>
C2	1	<u>.60</u>	1	.39	1	.12	10'	.05	<u>.28</u>	1	<u>.33</u>	<u>.42</u>	-	-	-	1	<u>.16</u>
C3	3	.29	1	<u>.88</u>	1	<u>.78</u>	-	-	-	-	-	-	-	-	-	1	<u>.36</u>
C4	(1)	(.36)	(99)	(.37)	(3)	(.03)	-	-	<u>.35</u>	-	-	<u>.44</u>	-	-	<u>.25</u>	1	<u>.15</u>
C5	(1)	(.57)	(86)	(.30)	(1)	(.23)	5	.02	<u>.18</u>	7'	.02	<u>.23</u>	-	-	.05	1	<u>.45</u>
C6	(19)	(.14)	(98)	(.40)	(1)	(.15)	2	.04	<u>.32</u>	1	.11	<u>.44</u>	-	-	<u>.22</u>	1	<u>.27</u>
C7	2	.52	98	.25	1	.09	-	-	.04	-	-	.05	-	-	-	(4)	(.02)
C8	1	<u>.76</u>	1	.65	1	<u>.37</u>	1	.11	<u>.37</u>	-	-	<u>.25</u>	-	-	-	1	<u>.32</u>
C9	1	<u>.69</u>	1	.51	1	<u>.27</u>	1	.03	<u>.26</u>	1	.08	<u>.39</u>	-	-	-	1	<u>.68</u>
C10	1	<u>.85</u>	1	<u>.81</u>	1	<u>.63</u>	1	.03	.03	1	.04	.04	-	-	-	1	<u>.51</u>
MRR	.58 (.79)		.60 (.60)		.79 (.93)		.48 (.48)			.60 (.60)			.00 (.00)			.90 (.93)	

Given our prior research on citation-based similarity [9, 10, 12, 13], we expected similar characteristics for the citation-based measures. However, as shown in Figure 2, the order-agnostic BC measure is more selective than the order-considering LCCS measure in this case. The reason is the errors in citation extraction (cf. Section 3.1). The mismatch of references and in-text citations causes that the LCCS and GCT measures can only consider a fraction of the citations in the dataset. This fraction is smaller than the fraction of extracted references, which the BC measure uses. Therefore, the BC measure is more selective than the LCCS measure for this dataset, since overlaps of the comparably sparse in-text citations increased the LCIS score more than overlaps in references increased the BC score. Unrecognized in-text citations also cause the GCT measure to be overly selective for this dataset. Due to a shortage of data points, the PrDF for scores of the GCT measure shows interpolation artifacts, i.e., the PrDF is not monotonically decreasing for larger scores. HyPlag could identify in-text citation tiles above the exclusion thresholds (cf. Section 3.4) for only 41 document pairs in our sample of 1M documents and for none of our test cases.

The PrDF of the Encoplot scores shows that the text-based measure is highly selective. Nine of the test cases have scores above the significance threshold, i.e., most verified cases of AP have a significant textual overlap with the respective source document. This characteristic is common for confirmed cases of AP [34]. Identifying literal text overlap is easier for reviewers and better supported by productive PD systems than identifying concealed content similarity. Therefore, documents with (near) copied text are more likely to be discovered and hence likely overrepresented in our dataset.

5.2.2 Combined rank-based & score-based assessment. In addition to assessing the significance of the similarity scores, we also examined the ranks r at which the similarity measures retrieved the source document for each of the test cases. To indicate the average ranking performance of the measures, we computed the *Mean Reciprocal Rank (MRR)*. In the case of tied ranks, we considered the mean rank, i.e., the pessimistically rounded average of the number of document pairs that share the same rank. The best possible score

of 1 is assigned if a similarity measure exclusively retrieves the source document at rank 1 for each test case.

Table 4 shows the results of both the rank-based and the score-based assessment. For each of the test cases (C1...C10) the table lists the rank r at which HyPlag retrieved the source document and which score s the similarity measure assigned. We mark the mean rank, which we list in the case of tied ranks, with an apostrophe, e.g., 7'. Scores above the significance threshold of a measure (see Table 3) are underlined. To gauge the performance of the similarity measures specifically for the detailed analysis stage, we also state the ranks and similarity scores for the cases not retrieved in the candidate retrieval stage. We mark such entries with parentheses, e.g., (0.15). To compute the ranks and scores for these documents, we performed a comparison of the query document to all documents in the dataset. Minus characters (-) indicate that HyPlag computed no similarity score due to the exclusion criteria of the measure. Because of the incomplete and error-prone extraction of bibliographic data, we state a separate score s^* for the citation-based measures. The score indicates the true citation-based similarity of the test cases. To compute s^* , we manually corrected erroneous data for in-text citations and references before applying the similarity measures.

The text-based approach consisting of word 3-gram fingerprinting (Sherlock) for the candidate retrieval stage and efficient string matching (Encoplot) for the detailed analysis stage achieved the best individual result. The approach retrieved nine of the ten test cases at the top rank. Only test case C7 exhibits a textual similarity that is too low to retrieve the source document in the candidate retrieval stage and mark the document as suspicious in the detailed analysis stage. The Encoplot scores for 6 of the 10 test cases exceed 0.25, hence are clearly suspicious. For the cases C1, C2, and C4, the Encoplot scores exceed our significance threshold of 0.06, yet are lower than 0.20. Anecdotal evidence¹⁴ suggests that 10% - 20% of text overlap is not immediately suspicious but often tolerated by journal reviewers and editors. The practices regarding acceptable

¹⁴www.researchgate.net/post/What_percentage_of_plagiarism_is_generally_treated_as_acceptable

text overlap vary between research fields and even between venues. Whether a productive text-based PD system would flag C1, C2, and C4 as suspicious is thus unclear. The retraction note of C1 names the unattributed reuse of a mathematical analysis, not the textual overlap with the source, as the reason for the retraction. The scores for Histo (0.68) and Git (0.21), which both exceed the significance thresholds, reflect this similarity in mathematical content.

The math-based similarity measures achieved the second-best result when considering both the candidate retrieval and detailed analysis stages. GIT, which we devised as a new similarity measure for this study, performed particularly well, retrieving seven cases at the top rank. When only considering the detailed analysis stage, GIT achieved the same effectiveness as the text-based analysis (9 test cases retrieved at rank one, MRR=0.93). To enable this result for the detailed analysis stage, the candidate retrieval procedure could simply combine the results of the math-based, citation-based and text-based approaches as discussed in Section 3.3.

GIT outperformed the Histo measure, which achieved the best results in our pilot study [26]. In this prior study, Histo achieved an MRR score of 0.86. Our current implementation exhibits a slightly lower MRR of 0.79. We attribute the difference to using a different conversion and data extraction process. The good performance of GIT suggests that the pattern of reusing (nearly) identical content in confined parts of a document known as "shake&paste" or "patchwriting" [42] also applies to mathematical content.

For our test cases, LCIS achieved no significant improvement over the set-based Histo measure. Both LCIS and Histo achieved good results for test cases that share a large fraction of their mathematical content. For such documents, the amount of shared math sufficed to retrieve the documents using the Histo approach. That the large overlap in mathematical content also yielded long identifier subsequences did not significantly improve the similarity score.

The citation-based measures achieved the lowest overall performance, largely due to the deficiencies of the extracted data. Despite the sub-optimal data, the LCCS measure retrieved 5 cases at rank one achieving an MRR score of 0.60. The similarity scores s^* , which assume the bibliographic data in the documents would have been extracted and matched correctly, give a better indication of the potential effectiveness of the citation-based measures. Notably, LCCS would yield scores of approx. two times the significance threshold of 0.22 and hence strongly suspicious for C2, C4, C6, and C9. Given that C2 and C4 exhibit a textual overlap that is significant but not strongly suspicious (0.16 and 0.15), the high LCCS score could provide an indicator for suspicious similarity.

For all cases except C7, which none of the measures flags as suspicious, at least one math-based or citation-based measure yields a similarity score above the individual significance thresholds. For Case C7, the Histo score has the smallest difference to the measure's significance threshold making Histo the most likely measure to retrieve the case despite the comparably low score.

In summary, the evaluation using confirmed cases of AP showed that the combined analysis of math-based and citation-based similarity identified all cases that also a text-based analysis flagged as strongly suspicious. Moreover, the two nontextual detection approaches provide valuable indicators for suspicious document similarity for cases with a comparably low textual similarity.

Table 5: Top-ranked documents in exploratory study.

Rank	1	2	3	4	5	6	7	8	9	10
Case	C3	C11	C12	C13	C10	C14	C15	C16	C17	C18
Rating	Plag.	Susp.	CR	FP	Plag.	FP	CR	CR	CR	CR

6 EXPLORATORY STUDY RESULTS

In this section, we describe our findings from manually investigating the top-ranked documents that HyPlag retrieved when applying math-based and citation-based content features to compare each document of the dataset to its individual set of candidate documents.

Given the size of the result set (approx. 6M document pairs) and our primary goal of searching for undiscovered cases of plagiarism, we employed several filters to reduce the result set and to focus our manual investigations on the most critical similarities. To eliminate cases, in which authors likely reused own content, we excluded document pairs that shared at least one author. This exclusion prevents the identification of potential self-plagiarism, i.e., the undue reuse of own previously published work. Similarly, we pruned document pairs, for which the older document cites the newer document, to reduce results in which authors reproduced previous work with due attribution. We make these restrictions for two reasons. First, the definition of what constitutes self-plagiarism varies greatly in different research fields and even for different venues. The vagueness of the problem definition prevents a well-founded assessment of the retrieved documents. Second, because we analyze all documents in the dataset, the number of results is much larger than in the typical PD scenario, i.e., analyzing a single input document for similarities to documents in a reference collection.

Since we are particularly interested in the benefit that a math-based similarity assessment can add to a combined approach, we excluded documents with a Histo score below 0.25. i.e., with little math-based similarity, and sorted the remaining results according to the GIT score in descending order. To not exclude cases, in which documents contained unequal amounts of identifiers, e.g., because one document is significantly shorter (cf. Section 3.4), we did not require a Histo score above the significance threshold of 0.56 but only a score that is greater or equal to 0.25.

Table 5 shows the ten top-ranked document pairs and our rating of the observed similarities. We use the abbreviated ratings 'Plag.' for confirmed cases of plagiarism, 'Susp.' for suspicious content similarity, 'CR' for notable but legitimate content reuse and 'FP' for false positives, i.e., documents with insignificant content overlap.

The highest ranked document pair is the confirmed case of plagiarism C3. The author of the retracted paper copied three geometric proofs with few changes from a significantly longer paper, thus resulting in a high GIT (0.78) but low Histo score (0.29). Another confirmed case of AP (C10) was retrieved at rank 5. The main contribution of the retracted paper in C10, a model in Nuclear Physics, was taken from the source paper while partially renaming identifiers. Almost the entire mathematical content of the retracted paper overlaps with the source document, resulting in the highest Histo score (0.85) in our exploratory study. The differences of identifiers in the source document and the retracted document result in a lower but still suspiciously high GIT score (0.63).

The later document in C11 (rank 2) is a mixture of idea reuse and content reuse. The author of the later paper reused the argumentative structure, sequence of formulae, several of the cited sources, many descriptions of formulae, and non-trivial remarks about the implications of the research from the earlier paper. By doing so, the author of the later paper derived a minor generalization of an entropy model for a specific type of black holes introduced in the earlier paper. The later paper cites other papers by the author of the earlier paper but not the earlier paper itself. We contacted the author of the earlier paper about our findings. In his view, the later paper "*certainly constitutes a case of plagiarism*". In coordination with the author of the earlier paper, we contacted the journal that published the later paper. The journal's editorial board currently examines the case. Since the journal has not published an official determination about the legitimacy of the paper, we classify the document as suspicious. This case exemplifies the benefits of a combined math-based, citation-based, and text-based similarity analysis. Only a combined analysis reveals the full extent of content similarity that encompasses approx. 80% of the paper's content.

The five cases of legitimate content reuse (C12, C15, C16, C17, and C18) exhibit similar characteristics. In all five cases, the authors of the later papers reproduce and properly cite extensive mathematical models proposed in the earlier papers. HyPlag failed to recognize the citations and to exclude the document pairs due to two challenges. First, the use of severely abridged citation styles, e.g., only stating the author name(s) and the arXiv identifier of a paper. Second, some authors cite the arXiv preprint of a paper, whereas other authors cite the journal version. The journal versions regularly exhibit differences in the order of authors and the title compared to the respective arXiv preprints. Both cases were not handled correctly by our preprocessing pipeline (cf. Section 3.1). Clearly, we need to improve our procedures for extracting and disambiguating such challenging references in STEM documents.

However, retrieving these five cases at top ranks is justified given the overlap in mathematical content (typically multiple pages). We expect that reviewers would like to be made aware of such content overlap, e.g., to verify the correct citation of the previous work.

The two false positives, C13 (rank 4) and C14 (rank 6) that HyPlag retrieved reveal potential improvements for the math-based similarity measures. C13 comprises of two papers in Combinatorics that contain long lists of all possible combinations of the identifiers a , b , and c according to a set of production rules. Similarly, C14 comprises two documents that analyze partition functions and contain long matches entirely made up of the identifiers p and q that occur in large quantities within unrelated formulae.

To increase the effectiveness of the math-based similarity measures and prevent such false positives, we plan to devise measures that are confined to individual formulae. Likewise, we plan to research how an assessment of structural and semantic similarity of formulae can be adapted for the plagiarism detection use case. Research on formula search and other mathematical information retrieval tasks has provided approaches that could prove valuable for the PD scenario [15, 26]. The math-based approach to PD is at an early stage of development. Like the early approaches for text reuse detection [2, 23], we investigated basic, computationally efficient feature analysis methods to identify the reuse of identical

and slightly different mathematical content. The results of our investigations show that the math-based analysis approaches increase the detection capabilities for STEM documents, particularly when being combined with other similarity assessments.

7 CONCLUSION & FUTURE WORK

By reviewing prior research, we showed that semantic, syntactic and cross-lingual PD approaches achieve high detection effectiveness, even for concealed forms of AP. However, these approaches require a high computational effort. The recall level of efficient text-based candidate retrieval methods stagnates. Approaches that analyze nontextual content features in academic documents, such as citations and mathematical content, show promise for being employed as computationally modest methods to retrieve candidate documents and as more elaborate detailed analysis methods.

The paper at hand extends a pilot study [26], in which we explored the potential of analyzing the similarity of mathematical content in the detailed analysis stage of the external PD process. In the current paper, we additionally devised a computationally efficient candidate retrieval stage that analyzes mathematical content features, academic citations, and textual features using production-ready information retrieval technology. Moreover, we created the GIT and LCIS measures, which consider the order of mathematical identifiers, for the detailed analysis stage. We implemented the newly developed math-based measures, as well as established citation-based and text-based measures in HyPlag - a working prototype of a hybrid plagiarism detection system.

Using HyPlag, we compared the effectiveness of the math-based, citation-based, and text-based PD approaches using confirmed cases of AP. We showed that a simple unification of the modestly sized sets of candidate documents retrieved by each retrieval heuristic achieved perfect recall for the candidate retrieval stage. For the detailed analysis stage, the newly developed GIT measure exceeded the effectiveness of the best performing approach (Histo) in our prior study and achieved the same effectiveness as the text-based similarity measure in our current study.

Errors in the acquisition of in-text citations and bibliographic references decreased the effectiveness of the citation-based similarity measures in our experiments. Despite the limitations, citation-based measures added a significant benefit to the hybrid approach, particularly for the candidate retrieval stage. LCCS also performed decently for the detailed analysis stage (MRR=0.60 for our test cases). The error-corrected similarity scores showed that the true effectiveness of the citation-based measures is much higher.

Overall, the combined analysis of math-based and a citation-based similarity identified all cases that a text-based analysis flagged as strongly suspicious. Moreover, the two nontextual detection approaches provided valuable indicators for suspicious document similarity for cases with a comparably low textual similarity. This result indicates that the best detection effectiveness can be achieved by combining the heterogeneous similarity assessments.

In an exploratory study, we showed the effectiveness of analyzing math-based and citation-based similarity for discovering unknown cases of potential AP. We used the GIT and Histo measures in combination with the citation relations of documents to reduce a result set of approx. 6M document pairs to 10 document pairs

that we investigated manually. The highest ranked document pair was a confirmed case of AP. The document retrieved at the second rank was rated as an undiscovered case of AP by the author of the apparent source document. The remaining 8 cases include one confirmed case of AP, 5 documents with high but legitimate overlap in mathematical content and 2 false positives. The citation-based filter criteria would have eliminated 5 cases of legitimate content reuse if the bibliographic data had been extracted correctly.

Our results show the large potential of analyzing mathematical content and academic citations as a complement to text-based PD approaches. However, our future work must improve the extraction of citation data for STEM fields to leverage the full potential of citation-based similarity measures. Additionally, we will increase the number of confirmed cases of AP and their degree of obfuscation to further support our results. For this purpose, we are collaborating with a major mathematical publishing service.

We will also research improvements to the math-based similarity measures. We expect that the performance of math-based heuristics for candidate retrieval can be improved by incorporating positional information about mathematical features. To improve the math-based measures employed for the detailed analysis stage, we will investigate approaches that consider the structural and semantic similarity of formulae. Another promising direction for future research that we plan to pursue is the application of machine learning to balance the weights of the similarity measures. Such an approach could train a detection system tailored to the domain-specific characteristics of a collection.

In summary, we see the integrated analysis of textual and non-textual features as the most promising approach to deter and to detect academic plagiarism in STEM research publications.

REFERENCES

- [1] Akiko Aizawa, Michael Kohlhase, Iadh Ounis, and Moritz Schubotz. 2014. NTCIR-11 Math-2 Task Overview. In *Proc. NTCIR*.
- [2] Salha M. Alzahrani, Naomie Salim, and Ajith Abraham. 2012. Understanding Plagiarism Linguistic Patterns, Textual Features, and Detection Methods. In *IEEE Trans. Syst., Man, Cybern. C, Appl. Rev.*, Vol. 42. 133–149.
- [3] Alberto Barrón-Cedeño, Parth Gupta, and Paolo Rosso. 2013. Methods for Cross-language Plagiarism Detection. *Know-Based Syst.* 50 (2013), 211–217.
- [4] Hannah Bast and Claudius Korzen. 2017. A Benchmark and Evaluation for Text Extraction from PDF. In *Proc. JCDL*.
- [5] Zdenek Ceska. 2008. Plagiarism Detection Based on Singular Value Decomposition. In *Advances in Natural Language Processing*. LNCS, Vol. 5221. Springer.
- [6] Nava Ehsan and Azadeh Shakery. 2016. Candidate Document Retrieval for Cross-lingual Plagiarism Detection Using Two-level Proximity Information. *Inf. Process. Manage.* 52, 6 (2016), 1004–1017.
- [7] Nava Ehsan, Frank Wm. Tompa, and Azadeh Shakery. 2016. Using a Dictionary and N-gram Alignment to Improve Fine-grained Cross-Language Plagiarism Detection. In *Proc. DocEng*.
- [8] Teddy Fishman. 2009. "We know it when we see it"? is not good enough: toward a standard definition of plagiarism that transcends theft, fraud, and copyright. In *Proc. Asia Pacific Conf. on Educational Integrity*.
- [9] Bela Gipp. 2014. *Citation-based Plagiarism Detection - Detecting Disguised and Cross-language Plagiarism using Citation Pattern Analysis*. Springer.
- [10] Bela Gipp and Norman Meuschke. 2011. Citation Pattern Matching Algorithms for Citation-based Plagiarism Detection: Greedy Citation Tiling, Citation Chunking and Longest Common Citation Sequence. In *Proc. DocEng*.
- [11] Bela Gipp, Norman Meuschke, and Joeran Beel. 2011. Comparative Evaluation of Text- and Citation-based Plagiarism Detection Approaches using GuttenPlug. In *Proc. JCDL*.
- [12] Bela Gipp, Norman Meuschke, Corinna Breitingner, Jim Pitman, and Andreas Nuernberger. 2014. Web-based Demonstration of Semantic Similarity Detection using Citation Pattern Visualization for a Cross Language Plagiarism Case. In *Proc. Int. Conf. on Enterprise Inform. Sys.*
- [13] Bela Gipp, Norman Meuschke, and Mario Lipinski. 2015. CITREC: An Evaluation Framework for Citation-Based Similarity Measures based on TREC Genomics and PubMed Central. In *Proc. iConference*.
- [14] Christian Grozea, Christian Gehl, and Marius Popescu. 2009. ENCOPLLOT: Pair-wise Sequence Matching in Linear Time Applied to Plagiarism Detection. In *Proc. PAN WS*.
- [15] Ferruccio Guidi and Claudio Sacerdoti Coen. 2016. A Survey on Retrieval of Mathematical Knowledge. *Mathem. in Computer Science* 10, 4 (2016), 409–427.
- [16] Deepa Gupta, Vani K, and Charan Kamal Singh. 2014. Using Natural Language Processing techniques and fuzzy-semantic similarity for automatic external plagiarism detection. In *Proc. Int. Conf. on Advances in Computing, Communications and Informatics*.
- [17] Matthias Hagen, Martin Potthast, and Benno Stein. 2015. Source Retrieval for Plagiarism Detection from Large Web Corpora. In *Proc. PAN WS*.
- [18] Kenichi Iwatsuki, Takeshi Sagara, Tadayoshi Hara, and Akiko Aizawa. 2017. Detecting In-line Mathematical Expressions in Scientific Documents. In *Proc. DocEng*.
- [19] Vani K and Deepa Gupta. 2015. Investigating the impact of combined similarity metrics and POS tagging in extrinsic text plagiarism detection system. In *Proc. Int. Conf. on Advances in Computing, Communications and Informatics*.
- [20] Leilei Kong, Haoliang Qi, Cuixia Du, Mingxing Wang, and Zhongyuan Han. 2013. Approaches for Source Retrieval and Text Alignment of Plagiarism Detection. In *Proc. PAN WS*.
- [21] Arun kumar Jayapal. 2012. Similarity Overlap Metric and Greedy String Tiling at PAN 2012. In *Proc. PAN WS*.
- [22] Donald L. McCabe. 2005. Cheating among College and University Students: A North American Perspective. *Int. J. for Academic Integrity* 1, 1 (2005), 1–11.
- [23] Norman Meuschke and Bela Gipp. 2013. State-of-the-art in detecting academic plagiarism. *Int. J. for Educational Integrity* (2013).
- [24] Norman Meuschke and Bela Gipp. 2014. Reducing Computational Effort for Plagiarism Detection by using Citation Characteristics to Limit Retrieval Space. In *Proc. JCDL*.
- [25] Norman Meuschke, Christopher Gondek, Daniel Seebacher, Corinna Breitingner, Daniel A. Keim, and Bela Gipp. 2018. An Adaptive Image-based Plagiarism Detection Approach. In *Proc. JCDL*.
- [26] Norman Meuschke, Moritz Schubotz, Felix Hamborg, Tomas Skopal, and Bela Gipp. 2017. Analyzing Mathematical Content to Detect Academic Plagiarism. In *Proc. CIKM*.
- [27] Norman Meuschke, Nicolas Siebeck, Moritz Schubotz, and Bela Gipp. 2017. Analyzing Semantic Concept Patterns to Detect Academic Plagiarism. In *Proc. Int. WS on Mining Scientific Publ. (WOSP) at JCDL*.
- [28] Norman Meuschke, Vincent Stange, Moritz Schubotz, and Bela Gipp. 2018. Hy-Plug: A Hybrid Approach to Academic Plagiarism Detection. In *Proc. SIGIR*.
- [29] H.F. Moed, W.J.M. Burger, J.G. Frankfort, and A.F.J. Van Raan. 1985. The application of bibliometric indicators: Important field- and time-dependent factors to be considered. 8, 3-4 (1985), 177–203.
- [30] Gabriel Oberreuter, Gaston L'Huilier, Sebastián Ríos, and Juan. Velásquez. 2011. Approaches for Intrinsic and External Plagiarism Detection. In *Proc. PAN WS*.
- [31] Merin Paul and Sangeetha Jamal. 2015. An improved SRL based plagiarism detection technique using sentence ranking. *Proc. CS* 46 (2015), 223–230.
- [32] Solange de L. Pertile, Viviane P. Moreira, and Paolo Rosso. 2016. Comparing and combining Content- and Citation-based approaches for plagiarism detection. *JASIST* 67, 10 (2016), 2511–2526.
- [33] Martin Potthast, Tim Gollub, Matthias Hagen, Jan Graßegger, Johannes Kiesel, Maximilian Michel, Arnd Oberländer, Martin Tippmann, Alberto Barrón-Cedeño, Parth Gupta, Paolo Rosso, and Benno Stein. 2012. Overview of the 4th International Competition on Plagiarism Detection. In *Proc. PAN WS*.
- [34] Martin Potthast, Benno Stein, Alberto Barrón Cedeño, and Paolo Rosso. 2010. An Evaluation Framework for Plagiarism Detection. In *Proc. ACL*.
- [35] Lutz Prechelt, Guido Malpohl, and Michael Philippsen. 2002. Finding plagiarisms among a set of programs with JPlag. *J. of Univ. CS* 8, 11 (2002), 1016.
- [36] Miguel A. Sanchez-Perez, Alexander Gelbukh, and Grigori Sidorov. 2015. Adaptive Algorithm for Plagiarism Detection: The Best-Performing Approach at PAN 2014 Text Alignment Competition. In *Proc. CLEF (LNCS)*, Vol. 9283.
- [37] Petr Sojka and Martin Liška. 2011. Indexing and Searching Mathematics in Digital Libraries – Architecture, Design and Scalability Issues. In *Proc. Int. Conf. on Intelligent Computer Mathematics (LNCS)*, Vol. 6824.
- [38] S. Soleman and A. Purwarianti. 2014. Experiments on the Indonesian plagiarism detection using latent semantic analysis. In *Int. Conf. on ICT*.
- [39] Benno Stein, Sven Meyer zu Eissen, and Martin Potthast. 2007. Strategies for Retrieving Plagiarized Documents. In *Proc. SIGIR*.
- [40] Dominika Tkaczyk, Paweł Szostek, Mateusz Fedoryszak, Piotr Jan Dendek, and Lukasz Bolikowski. 2015. CERMINE: Automatic Extraction of Structured Metadata from Scientific Literature. *Int. J. Doc. Anal. Recognit.* 18, 4 (2015), 317–335.
- [41] Juan D. Velásquez, Yerko Covacevich, Francisco Molina, Edison Marrese-Taylor, Cristián Rodríguez, and Felipe Bravo-Marquez. 2016. DOCODE 3.0 (DOcument COpy DeteCtor). *Information Fusion* 27 (2016).
- [42] Dora Weber-Wulff. 2014. *False Feathers: A Perspective on Academic Plagiarism*.
- [43] Michael J. Wise. 1993. String Similarity via Greedy String Tiling and Running Karp-Rabin Matching. TR (Univ. of Sydney. Basser Dept. of CS) 463.

Citation for this Paper

Citation Example:

N. Meuschke, V. Stange, M. Schubotz, M. Kramer, and B. Gipp. Improving Academic Plagiarism Detection for STEM Documents by Analyzing Mathematical Content and Citations. In *Proceedings ACM/IEEE Joint Conference on Digital Libraries (JCDL)*, 2019.

Bibliographic Data:

RIS Format	BibTeX Format
TY - CONF AU - Meuschke, Norman AU - Stange, Vincent AU - Schubotz, Moritz AU - Kramer, Michael AU - Gipp, Bela T1 - Improving Academic Plagiarism Detection for STEM Documents by Analyzing Mathematical Content and Citations T2 - Proceedings ACM/IEEE Joint Conference on Digital Libraries (JCDL) Y1 - 2019	@InProceedings{Meuschke2019, author = {Meuschke, Norman and Stange, Vincent and Schubotz, Moritz and Kramer, Michael and Gipp, Bela}, title = {Improving Academic Plagiarism Detection for STEM Documents by Analyzing Mathematical Content and Citations}, booktitle = {Proceedings ACM/IEEE Joint Conference on Digital Libraries (JCDL)}, year = {2019} }