

Forms of Plagiarism in Digital Mathematical Libraries

Moritz Schubotz^{1,2}, Olaf Teschke², Vincent Stange¹, Norman Meuschke¹, and Bela Gipp¹

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

² FIZ Karlsruhe / zbMATH, Berlin, Germany
olaf.teschke@fiz-karlsruhe.de

Abstract. We report on an exploratory analysis of the forms of plagiarism observable in mathematical publications, which we identified by investigating editorial notes from zbMATH. While most cases we encountered were simple copies of earlier work, we also identified several forms of disguised plagiarism. We investigated 11 cases in detail and evaluate how current plagiarism detection systems perform in identifying these cases. Moreover, we describe the steps required to discover these and potentially undiscovered cases in the future.

1 Introduction

Plagiarism is ‘the use of ideas, concepts, words, or structures without appropriately acknowledging the source to benefit in a setting where originality is expected’ [5, 7]. Plagiarism represents severe research misconduct and has strongly negative impacts on academia and the public. Plagiarized research papers compromise the scientific process and the mechanisms for tracing and correcting results. If researchers expand or revise earlier findings in subsequent research, papers that plagiarized the original paper remain unaffected. Wrong findings can spread and affect later research or practical applications.

Furthermore, academic plagiarism causes a significant waste of resources [6]. Reviewing plagiarized research papers and grant applications causes unnecessary work. For example, Wager [29] quotes a journal editor stating that 10% of the papers submitted to the respective journal suffered from plagiarism of an unacceptable extent. If plagiarism remains undiscovered, funding agencies may even award grants for plagiarized ideas or accept plagiarized research papers as the outcomes of research projects. Studies showed that some plagiarized papers are cited at least as often as the original [14]. This is problematic, since publication and citation counts are widely used as indicators of research performance, e.g., for funding or hiring decision. Moreover, universities and research institutions invest considerable resources to investigate and sanction plagiarism.

The waste of resources and the deterioration of academic quality due to plagiarism is also a pressing concern for zbMATH. The zbMATH abstracting

and reviewing service organizes and reviews the world’s literature in mathematics and related areas since 1931. zbMATH includes its predecessor *Jahrbuch für die Fortschritte der Mathematik*, which goes back to 1868. As of today, zbMATH comprises about 4 million publications, reviewed with the help of more than 7’000 international domain experts. These expert reviewers do not repeat the foregone peer-review process of the venue publishing the article in question. Rather the zbMATH reviewers provide an unbiased view on the originality and innovative potential of articles for their subject areas, i.e., across venues. Ensuring the quality of articles is an integral part of zbMATH’s mission statement and central to the usefulness of the service.

However, so far zbMATH’s reviewing process for articles is conducted manually without the help of automated originality checks. Quality and originality insurance relies entirely on the knowledge of the expert reviewers who have to identify potential content overlap with other papers. The ability to spot such similarities requires reviewers to be closely familiar with similar articles. If an article exhibits significant overlap with prior works without discussing this fact, the reviewers will assign an editorial remark to the article.

The growth in publications indexed by zbMATH has raised concerns regarding the viability of the entirely manual reviewing process. Thus, zbMATH is investigating the use of automated approaches to support their reviewers. To guide and support this process, we performed an exploratory analysis of cases for which zbMATH reviewers questioned the originality of articles or identified plagiarism with certainty. Our goal is to develop an understanding of the typical characteristics of mathematical articles with questionable originality to derive the requirements for an automated system that can detect such cases.

We structure our report on this investigation as follows. Section 2 discusses related work before Section 3 presents the methodology of our investigation. Section 4 describes our findings and discusses their implications. Section 5 summarizes our results and gives an outlook on the next steps of the project.

2 Background and Related Work

The problem of academic plagiarism has been present for centuries [30]. However, the advancement of information technology has made plagiarizing easier than ever [6]. Forms of academic plagiarism range from copying content (*copy&paste*) over “patch-writing”, i.e., interweaving text from multiple sources with moderate adjustments, to heavily concealing content reuse, e.g., by paraphrasing or translating text, and reusing data or ideas without proper attribution [30]. The easily recognizable copy&paste-type plagiarism is more prevalent among students whose main motivation for plagiarizing is typically to save time [15]. Concealed forms of plagiarism are more characteristic of researchers, who have strong incentives to avoid detection [2].

While making plagiarizing easier, information technology also facilitated the detection of plagiarism. Researchers proposed many plagiarism detection approaches that employ lexical, semantic, syntactical, or cross-lingual text analysis

[17, 4, 28]. Such approaches typically employ a two-stage process consisting of candidate retrieval and detailed analysis [25, 17]. 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 [17, 2, 6].

Current plagiarism detection approaches reliably detect copied or moderately altered text; some approaches are also effective for finding paraphrased and translated text. Most plagiarism detection systems available for productive use focus on reliably and efficiently identifying plagiarism forms with little to no obfuscation, which are characteristic of students. We conjecture that student plagiarism is a more profitable market segment for commercial providers of plagiarism detection services. Arguments in favor of this hypothesis are the higher number of students compared to researchers, the higher frequency of plagiarism among students than among researchers [27, 15] and the availability of well-established, efficient methods to find literal text reuse [17, 4, 28].

The market leader for plagiarism detection services, *iParadigms LLC*³, offers its products under several brand names. The system *turnitin*⁴ is tailored to providing originality checks and academic writing training for students. The *iThenticate*⁵ service is mainly offered to academic publishers and conference organizers for checking research publications. The iThenticate service is also licensed to other academic service providers, such as *Crossref* who offers the service to its members as *Similarity Check*⁶.

The detection methods employed by commercial systems, such as *turnitin* and *iThenticate*, are trade secrets. However, the performance of the systems in benchmark evaluations [31] suggests that they mainly use efficient text retrieval methods, such as word-based fingerprinting and vector space models. Text fingerprinting approaches first split a document into (possibly overlapping) word or character n -grams, which are used to create a representation of the document or passage (the ‘fingerprint’) [17]. To enable efficient retrieval, most approaches select a subset of the fingerprints, which they store in an index. To speed up the comparison of fingerprints, some approaches hash or compress the fingerprints, which reduces the lengths of the strings to compare and allows for computationally more efficient numerical comparisons [6].

To complement the many text analysis approaches and to improve the detection capabilities of concealed forms of academic plagiarism, researchers proposed approaches that analyze nontextual content features, such as academic citations [9, 8, 7, 11, 10, 16, 23] and images [18]. Nontextual content features in academic documents are a valuable source of semantic information that are largely

³ <https://www.crunchbase.com/organization/iparadigms-inc>

⁴ <https://www.turnitin.com>

⁵ <https://www.turnitin.com/products/ithenticate>

⁶ <https://www.crossref.org/services/similarity-check/>

independent of natural language text. Considering these sources of semantic information for similarity analysis raises the effort plagiarists must invest for obfuscating reused content [16, 22].

Nontextual feature analysis appears to be a promising approach to plagiarism detection for mathematics and related fields. In these fields, much of the semantic content of publications is expressed in terms of mathematical notation. Research showed that classical text retrieval methods, which are similar to the methods employed by commercial plagiarism detection services, are less effective for documents in mathematics, physics and other domains that routinely interweave natural language and mathematical notation [32].

However, only a few studies have addressed the detection of plagiarism in digital mathematical libraries [19, 21, 22] regardless of the detection approach. We briefly describe the main findings of these studies hereafter.

Following up on discussions at the doctoral consortium of the SIGIR conference 2015, Meuschke et al. (2017) described mathematics-based plagiarism detection (MathPD) as a discrete sub-problem within mathematical information retrieval. The authors argued that the different approaches to query formulation and query processing distinguish MathPD from the mathematical document retrieval problem as defined by Guidi and Sacerdoti Coen [12].

To test whether an exclusive analysis of mathematical similarity holds promise for plagiarism detection, Meuschke et al. (2017) gathered documents that have been retracted for plagiarism and contain significant amounts of mathematical content from three sources. The first source was an earlier study on retracted publications by Halevi and Bar-Ilan [13]. This study had queried Elsevier’s full text database *Science Direct* for the term "RETRACTED" in October 2014. The search yielded 988 retracted articles, of which 276 had been retracted for plagiarism [13]. Meuschke et al. limited these 276 publications to a set of 39 publications that contain significant amounts of mathematical content. The second source was the blog *Retraction Watch*⁷ from which Meuschke et al. obtained two confirmed cases of plagiarism. The third source was the crowd-sourced project *VroniPlag*⁸, from which Meuschke et al. obtained three additional cases. The authors then limited the 44 cases they had gathered to cases that matched their area of expertise, i.e., mathematics, physics, and computer science, which resulted in 19 cases the authors then investigated manually. They categorized the types of shared mathematical content they observed in the analysis of these 19 cases into six broad categories [19]:

Identical: an exact copy of math in the source document.

Equivalent: equivalent forms, e.g., due to commutativity or distributivity.

Order changes: order of expressions within document differs.

Different presentation: structurally and semantically identical.

Splits or merges: of expressions that are semantically identical.

Different concepts: different, yet semantically (nearly) identical, concepts, e.g., summation over vector components instead of matrix multiplication.

⁷ <http://www.retractionwatch.com>

⁸ <http://www.vroniplag.wikia.com>

Of the 19 cases reviewed manually, Meuschke et al. selected 10 cases that were most representative of the types of content similarity they observed. The authors converted the ten source documents and ten retracted documents of those cases from PDF to LaTeX using InftyReader [26] and subsequently from LaTeX to XHTML using LaTeXML⁹. They embedded the converted documents in the dataset of the NTCIR-11 Math task [1] (105'120 arXiv documents). Using pairwise document comparisons, Meuschke et al. evaluated the retrieval effectiveness of similarity measures that consider basic representational math features, i.e., identifiers, numbers, operators, and combinations thereof. 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.

In a follow-up study, Meuschke et al. (2019) introduced similarity measures that consider the order of mathematical identifiers and presented a two-stage retrieval process consisting of a candidate retrieval and a detailed analysis stage that replaced the exclusive use of pairwise document comparisons [22]. They implemented the process in the HyPlag prototype that also offers a user interface to investigate the identified similarities [21]. The candidate retrieval stage employs efficient index-based retrieval methods based on mathematical features. The candidate documents retrieved in the first stage then undergo pairwise comparisons in the detailed analysis stage of the process. The authors compared the effectiveness of their math-based analysis to citation-based and text-based approaches. They found that the order-observing similarity measures for mathematical identifiers achieved better results than the order-agnostic measures in their previous study [19]. Most of their ten test cases also exhibit a high textual similarity. A combined analysis of math-based and citation-based similarity performed equally well as the text-based analysis for the ten test cases. In an analysis of all 105K documents in their dataset, Meuschke et al. demonstrated that the combined analysis of math-based and citation-based similarity has advantages over a text-based analysis by identifying interesting instances of content reuse that a text-based analysis could not detect.

In summary, Meuschke et al. (2019) demonstrated that the combined analysis of text-based and nontextual similarity, e.g., the similarity of mathematical content and citations, achieves promising results for retrieving confirmed cases of plagiarism that involve mathematical content. However, the set of documents they analyzed is small and likely does not reflect the full spectrum of possible content reuse in digital mathematical libraries. To aid in the advancement of plagiarism detection methods for mathematics and related disciplines, we analyze a larger set of cases obtained from the zbMATH collection. An interesting property of this collection is that it exclusively includes published work since 2007, i.e., from a period in which larger publishers already employed services like CrossRef's Similarity Check. Also digitization has reached maturity during this period, so we can expect that these examples were not detected by standard tools.

⁹ <http://dlmf.nist.gov/LaTeXML/>

3 Method

Subsection 3.1 describes our approach to identify cases of *noticeable content reuse* (NCR). Subsection 3.2 presents the challenges we faced in regard to processing the cases using the HyPlag detection system. Lastly, Subsection 3.3 defines the properties we evaluate for a small number of example cases.

3.1 Identification of noticeable content reuse

The zbMATH collection currently contains 3’981’836 publications. Since 2007, zbMATH follows a intra-organizational procedure for marking publications that exhibit noticeable content reuse. The managing editor and the deputy editor-in-chief decide on a case-by-case basis on the actions to be taken if illegitimate content reuse is observed. From 2007 to 2018, 1’226’203 new publications have been added to the zbMATH database. Although all publications underwent peer-review by the publishing venue before being submitted to zbMATH, 446 cases of questionable content reuse have been reported to the managing editor of zbMATH. After careful investigation of these reports, 144 cases received an editorial note about NCR. This list includes cases of content reuse by the same authors in different papers as well as content reuse by different authors. Moreover, in contrast to Meuschke et al. (2019), this dataset includes both true positives (144), i.e., suspicious documents, and false positives (302).

The list of the 144 suspicious cases is openly available from the zbMATH website¹⁰ using the query shown in Listing 1.1. The logic of the query is as follows: Line 1 filters for publications that appeared between 2007 and 2018, which matches the period investigated in this paper. ‘py’ is the query term for publication year. Line 2 filters the abstracts (ab) for the keywords ‘editorial remark’ or ‘editorial note.’ Line 3 filters for remarks that indicate plagiarism. Examples for excluded editorial notes are incorrect results, conceptual flaws, or organizational comments. Note that we applied this filter to all fields, not just the abstract, since sometimes the indicator for an editorial note is included in other fields, e.g., the keywords. For example the, document 1191.35223¹¹ is in the category ‘suspected plagiarism,’ although this is not indicated in the abstract of the document. Line 4 corrects the search result and excludes false positives.¹²

3.2 The challenge of full text availability

After having identified the set of cases to investigate, we had to import the documents to our plagiarism detection system HyPlag (cf. Section 2), which

¹⁰ <https://zbmath.org>

¹¹ We use zbMATH identifiers for referring to cases throughout the paper. The identifiers resolve via, e.g., <https://zbmath.org/1191.35223> to documents accessible without subscription.

¹² See <https://zbmath.org/general-help/> for the details of the search syntax.

Listing 1.1. zbMATH search query to retrieve papers with a noticeable amount of content reuse. Click here to execute the query. Subscription required for full display!

```

1 py:2007–2018 &
2 ( ab:"editorial remark" | ab:"editorial note" ) &
3 ( "very similar" | "high similarity" | overlap | plagiari* |
   identical | substantial* | essentially )
4 !( so:ieee | se:00000250 | se:00001661 | pu:AIP |
   an:0584.10010 | an:0712.35001 | an:0597.14041 |
   an:1375.14126 | an:0156.05104 | an:1345.15011 |
   an:1262.11083)

```

requires text encoding initiative (TEI) format as input. Unfortunately, the zbMATH dataset does not include full texts. The full texts might be obtained by following the DOI, an arXiv link, or a link to another digital repository. Yet, in most cases (except for arXiv) the LaTeX sources of the documents are unavailable. While HyPlag includes a conversion procedure that generates the TEI input format from PDF files via the PDF processor GROBID¹³, this procedure cannot process mathematical formulae. As a result, TEI documents generated from PDF documents miss mathematical formulae.

In a previous study [19], we employed the image to formula conversion tool InftyReader (cf. Section 2). However, the results were unsatisfactory. While InftyReader extracted parts of the formulae correctly, its extraction accuracy regarding the structure of the document was significantly worse than that of GROBID. We also evaluated several alternative tools, including the combination of maxtract [3] and pdfminer¹⁴. We also discovered and tested new machine learning based approaches¹⁵. However, these approaches only worked in a few exceptional cases.

Eventually, we were unable to identify a tool capable of converting the zbMATH documents including all essential features, i.e, text, formulae and figures. Consequently, we modified the setup of our study and decided to investigate only a small sample of the collected documents. To still analyze a diverse and representative test dataset, we decided to use a semi-random method. The managing editor of zbMATH manually selected interesting cases that had undergone extended internal discussions, which we see as an indicator that the potential overlap was not initially obvious. Due to the diversity of the cases and the different decisions zbMATH reviewers made regarding the legitimacy of the observed similarities, the documents can form an interesting test collection for PD systems in mathematics. The small test collection consists of the 11 items listed in Table 1.

¹³ <https://grobid.readthedocs.io>

¹⁴ <https://pypi.org/project/pdfminer/>

¹⁵ for example: <http://cs231n.stanford.edu/reports/2017/pdfs/815.pdf>

3.3 Properties to investigate

After choosing the test cases, we specified the following properties to investigate:

1. Which type of plagiarism is observable in the article?
2. Is the article digitally available? Has it been retracted?
3. How important were the text, the figures, the references, and the formulae during the discussion about issuing an editorial remark?
4. When were the articles published? What are the languages of the documents?
5. What is the impact of the case? Has the article that received an editorial remark been cited. Does other statistical data on the use of this article exist?
6. Which requirements on a plagiarism detection system derive from this case?

4 Results

In this section, we first present descriptive statistics for the cases exhibiting noticeable content reuse. Second, we describe the 11 selected test cases in detail. Third, we derive requirements on a plagiarism detection system from the exemplary results of our analysis.

4.1 Distribution of the source documents

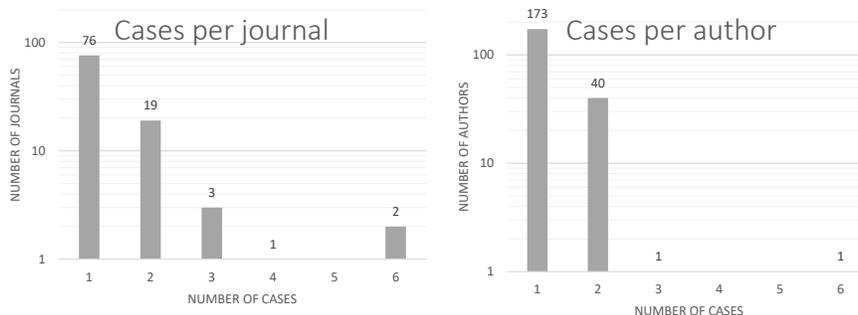


Fig. 1. Distribution of journal (left) and author (right) frequencies on a logarithmic scale.

The 144 publications that received an editorial remark for NCR consist of 139 journal articles and 5 books. The 139 articles originate from 101 journals. Two journals published 6 articles with editorial remarks each (cf. Figure 1), while 76 journals published only one such article. The 144 documents have 215 authors. One person was an author of 6 publications that received an editorial remark, while 173 authors were involved in only one case (cf. Figure 1). For 78 of the 144 documents, the full texts are available from at least one source (DOI (76), arXiv (3), eudml (4), or emis (3)). See Appendix A for the complete list of cases.

Table 1. Overview of manually investigated cases. The similarity scores were computed using HyPlag. l_i denotes the later document of case i and $e_{i,j}$ the j -th earlier document of case i .

#	l_i	$e_{i,j}$	retr.	avail. via	text sim	cit sim	spread
1	1349.46021 (2015)	06696052 (2015) 1353.46015 (2014)	-	doi, eudml, ams	.23	-	40 reads on researchgate
2	1381.51005 (2008)	1162.51304 (2007)	yes	ams	.03	-	-
3	1119.11307 (2001)	1062.11019 (2000)	yes	doi, ams	.33	-	40 reads on researchgate, 21 downloads on springer link
4	1112.35034 (2005)	0632.65108 (1987)	-	doi, ams	.33	-	6 reads on researchgate
5	1183.05037 (2008)	0247.05143 (1972)	-	doi, ams	.0	.75	5 reads on researchgate
6	1121.35118 (2005)	1062.81046 (2004)	-	doi, eudml, arXiv	.06	-	5 reads & 3 cits. on researchgate, 61 downloads on springer link
7	1176.08001 (2009)	1036.08001 (2003)	-	-	-	-	-
8	1219.30004 (2011)	1040.30002 (2004)	-	doi, eudml	.14	.19	21 reads & 4 cits. on researchgate
9	0946.35085 (1999)	0816.47056 (1994)	-	doi	.09	-	4 reads & 1 cit. on researchgate, 27 downloads on springer link
10	1155.35429 (2006)	1142.35593 (2004)	-	doi	.16	-	112 reads & 41 cit. on researchgate, 38 cits. on ScienceDirect
11	1360.81259 (2017)	1185.81005 (2009) - (-)	-	doi	0.1	-	-

4.2 Manual investigation

Hereafter, we present our findings from manually analyzing the 11 test cases:

Case 1 (legitimate content reuse) consists of the inspected document l_1 and the earlier works $e_{1,1}$ and $e_{1,2}$.

While the zbMATH editorial remark just reads: “Almost the same results were previously obtained in ($e_{1,1}$)”, more information is available through the reviews at <https://mathscinet.ams.org/mathscinet-getitem?mr=3390281>:

- Theorem 2.3(2) is Theorem 2.4(7)
- Corollary 2.4 is Corollary 2.6
- Corollary 2.5(2),(4) and (5) is Corollary 2.5(5),(3) and (7)
- Theorem 2.7 and its Corollary 2.8 are respectively Theorem 2.11 and Corollary 2.12 of the older document (by noting from Theorem 2.11 of the older document that the order bounded operators $T : E \rightarrow E$ in Corollary 2.12 may be replaced by positive ones).

Moreover, the reviewer states for the second prior document $e_{1,2}$

- Note also that Theorem 2.3(5) and Corollary 2.5(8) are respectively Theorem 2.7(8) and Corollary 2.10(5) of the second prior document. The proofs of the results that are republished in the paper under review are similar to those of the other publication.

Despite the overlap in content, the editor decided to publish the article. The rationale of this decision is that the similar structure seems to be the obvious solution to the given problem. However, this explanation does not justify the high textual similarity (0.23) that HyPlag computed for l_1 and $e_{1,1}$. Despite this moderately high value¹⁶ for the overall textual similarity, the HyPlag visualization¹⁷ (cf. Figure 2) shows that the text similarity in the section described by the reviewers is particularly high. To our knowledge, neither the reviewer, nor the editor had access to a text similarity visualization such as the one of HyPlag when they decided not to retract the article.

Note that 1394.47040 also used theorems from $e_{1,1}$ in combination with results from 1336.46019. The authors published an erratum 06644537 and acknowledged that ‘some results’ have been proven by $e_{1,1}$ and 1336.46019. They ‘feel sorry’ that they have ignored the original sources. In contrast to the other sources 1394.47040 draws upon, We did not identify a textual similarity of 1394.47040 with $e_{1,1}$ and 1336.46019.

Case 2 (retracted) consists of the retracted paper l_2 which overlaps with e_2 . The retraction note states ‘The author and the editor regret an oversight that this paper has significant overlap with the prior publication.’

¹⁶ A value above 0.2 is considered as suspicious.

¹⁷ <https://www.hyplag.org/> user cicm@hyplag.org pw: cicm2019

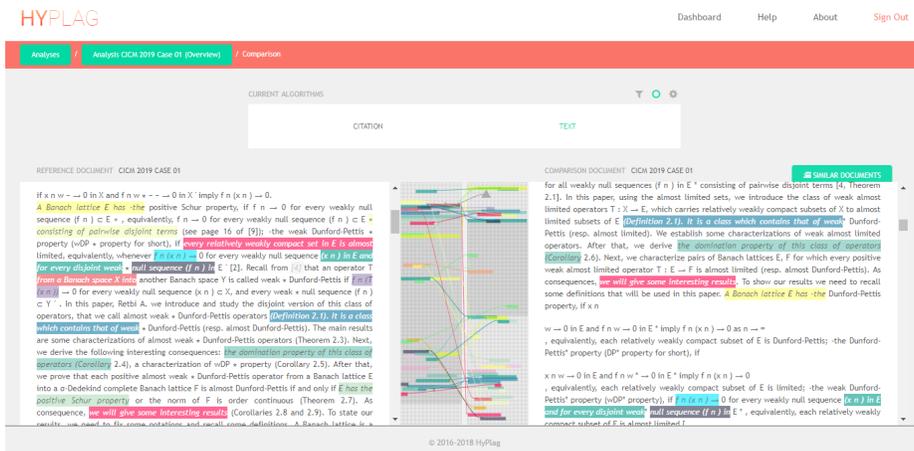


Fig. 2. Visualization of case 1 in HyPlag prototype.

From a plagiarism detection perspective, this case is particularly interesting since the later paper does not reuse text from the earlier paper (text similarity score 0.03). Moreover, the documents do not have common citations. However, many formulae and figures in the documents can be used to identify the overlap.

Unfortunately, the extraction of the figures from the PDF was challenging. State of the art image extraction technology was unable to identify the images. The internal representation of the figures in the PDF object stream prevents the identification of the basic geometric objects as a meaningful figure.

Case 3 (retracted) consists of the paper l_3 which is similar to e_3 without citing the earlier work. The later paper reuses several identical formulae and has a significant text overlap (text similarity score 0.33). The paper was published in 2001 and retracted in 2007. In the meantime, it received 2 citations in zbMATH - interestingly, by two papers that also cited e_3 .

For this case, the LaTeX sources were available to us. Thus, we can visualize the formula similarity using HyPlag. Despite the high text similarity, the identical formulae are apparent within the visualization (cf. Figure 3).

Case 4 (plagiarism) In this case, the authors of 1253.35026, which they also published as l_4 , copied results from e_4 published in 1987. The papers have a high text similarity of 0.33, even though we extracted the text from the earlier paper using OCR. However, one may speculate that the limited availability of the original in digital form (only available as a scan) contributed to the late discovery of the case. The republished version of the article l_4 is unavailable in digital form. 1253.35026 has been cited once. The case illustrates the limitations resulting from insufficient digitisation.

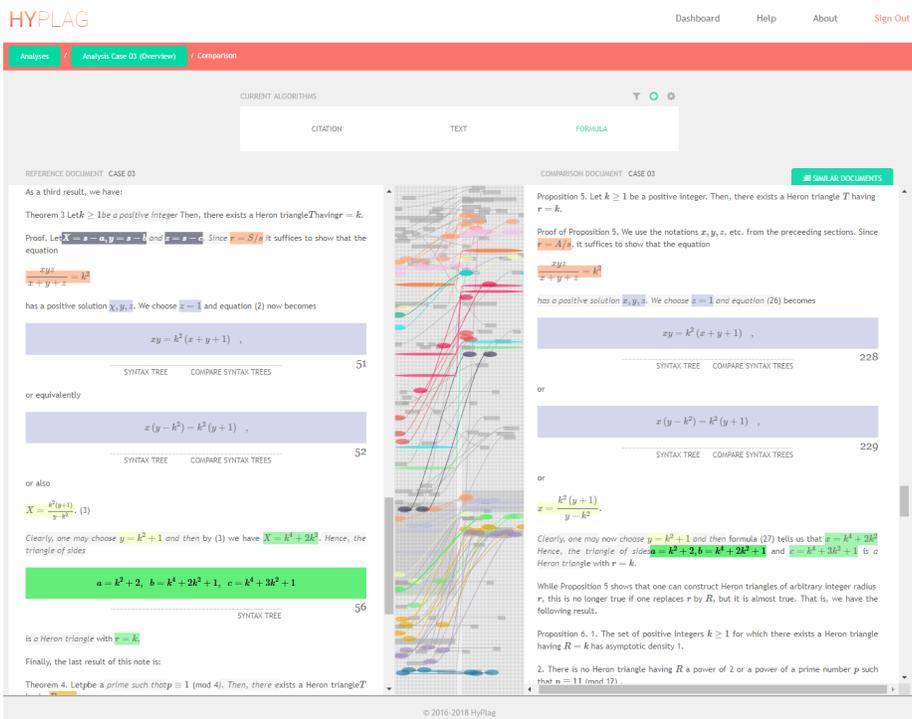


Fig. 3. Visualization of case 3 in HyPlag prototype.

Case 5 (translation) is a prototypical example for translation plagiarism. The original work e_5 written in French and published 1972 was translated to l_5 published in 2008 without acknowledging the original source.

This case likewise illustrates the limitations induced by sub-optimal digitisation. While both papers are available via DOI, the scan of the original only allows for an OCR quality that is insufficient for matching the citations in both documents. For instance, using Adobe Acrobat version 11.0.0, one receives the following representation of the references:

- [1] A. RÉNYI, On connected graphs, Magyar Tud. Akad. Mat. Kutató Int. Közl., 4, 1959, p. 385–388.
- [2] G. FORD, G. UHLENBECK, Combinatorial problems in the theory of graphs, Proceedings of the National Academy of Sciences of the U. S. A., 42, 1956, p. 122–128.
- [3] J. VAN MOON, Enumerating Labeled Trees, dans: Graph Theory and Theoretical Physics, editor F. Harary, Academic Press, New York–London, 1967, p. 261–271.

The quality of the extracted references was too low to match references and in-text citations in the documents. To demonstrate the potential of the approach, we manually corrected the OCR text. Figure 4 shows the result: the first three citations appear in identical order in both documents resulting in a high citation-based similarity score. e_5 has been neither retracted nor cited so far.



Fig. 4. Visualization of case 5 in HyPlag prototype

Case 6 (topical relatedness) consists of l_6 . The main result of this paper had been proven by an earlier article e_6 using similar methods (or, as <https://mathscinet.ams.org/mathscinet-getitem?mr=2177689> states: “The same theory can be found in an essentially identical presentation in [the] earlier paper [...] not cited by the present author.”). A subsequent discussion by expert reviewers remained inconclusive; thus, the paper has not been retracted. Technical PD tools fail to detect any significant similarity of the articles. Perhaps due to the awareness created by the reviewers’ discussion, l_6 received only two (self-)citations, while e_6 is cited frequently.

Case 7 (distribution stopped) represents the translation of a book. The German original e_7 from 2003 was translated to English and published in 2009 as l_7 . The publisher stopped the distribution of the English version¹⁸, which therefore was digitally unavailable to us. Despite its short period of availability, l_7 has been cited at least three times, while only one citation of e_7 is known so far.

Case 8 (identical) is an example for a work that reuses an earlier work, changes the mathematical notation, and presents a weaker result. The later paper l_8 which reused content from e_8 is still online without referencing the original work. The later paper received three citations according to Google scholar, was downloaded 4013 times, and viewed 7615 times according to the publisher¹⁹. Despite the significant overlap, the publisher did not issue a retraction but only a statement of priority for the original result. The statement appeared inconspicuously in a later issue and is not linked from the article itself. The similarity of the text and the references, although slightly below a critical threshold, should have issued at least a warning by a PD system.

¹⁸ <https://www.emis.de/misc/articles/ext05526289.html>

¹⁹ <https://www.scirp.org/journal/PaperInformation.aspx?PaperID=3820>

Case 9 (unclear) comprises l_9 which adopted content from the earlier work e_9 . The notation was changed and the text differs (text similarity 0.09). However, expert reviewers qualified the later article as: “derivative work”. So far, no retraction has been issued. According to the publisher²⁰, the derived work was downloaded 27 times and received one citation.

Case 10 (unclear) consists of the paper l_{10} , parts of which reuse material from e_{10} literally identical. However, the overall text similarity is below the critical threshold (score 0.16). Another noticeable difference is that one paper uses the computer algebra system Maple while the other paper uses Mathematica. No retraction has been issued. According to a later comment²¹, the later article derives incorrect formulae. Nevertheless, both articles achieved citation counts that are well above the average for math articles. The citations seem to originate from a rather peculiar community. A PD system allowing for a uniform representation of Maple/Mathematica content would have facilitated a clear detection of the similarities. However, such a system is a distant prospect.

Case 11 (compilation of text elements) consists of paper l_{11} which combines material from $e_{11,1}$ and a paper²² not part of the zbMATH corpus. A reviewer who had critically noted the authors’ way of working before indicated the case to zbMATH. So far, the journal has not reacted on the comments in any way. Due to the low visibility of the journal, the compilation seem to have had little impact so far. The comparably low text similarity of the article derives intrinsically from being a compilation of two sources. For such cases, an adapted measure resulting in an adequate warning would be desirable. Humans spot the respective adaptations quickly.

4.3 Requirements for a plagiarism detection system at zbMATH

The investigation of the selected cases indicates that the application of a purely text-based system, such as the commercial service of iThenticate, appears insufficient for analyzing content overlap in mathematical publications. Many publishers already use iThenticate as part of their submission pipelines.

The major obstacle for using the open source solution HyPlag is the availability of high quality sources. While PDF files were often available, the mathematical formulae could not be extracted from these PDFs. Therefore, the math-based similarity detection of HyPlag [20] could not be evaluated in this paper. Moreover, some PDFs are of a low quality and resulted in OCR text that is too erroneous for citation matching. Another problem with the PDF sources was that figures could not be identified.

²⁰ <https://link.springer.com/article/10.1007%2FBF02463791>

²¹ <https://doi.org/10.1016/j.camwa.2011.01.043>

²² https://math.berkeley.edu/~kwey/papers/string_theory.pdf

5 Conclusion and Outlook

We created an openly accessible dataset of 144 papers with noticeable content reuse in zbMATH. The dataset can serve as a training set or test set, e.g., for plagiarism detection competitions, such as PAN. In a second step, we extended the 144 confirmed cases of NCR with cases for which the content similarity was eventually rated as legitimate. To not discredit authors who were incorrectly accused of wrongdoing, we refrain from publishing the complete dataset of 446 cases. Instead, we composed a list of 11 typical cases that illustrate the spectrum of reported content reuse. Moreover, we will continue our analysis with more cases to derive general patterns.

Using the 11 cases, we investigated how the plagiarism detection system HyPlag would perform in identifying the documents as suspicious. In a recent study [22], we applied the system to a large test collection producing only a small number of false positives. However, nine of ten test cases in [22] could have been discovered using traditional text-based detection methods. In contrast, for the zbMATH collection, this number is only three of eleven. For the text-based and the math-based detection methods, we needed to transcribe the test data manually since the sources were unavailable and the quality of formulae, citations, and figures extracted from PDF was insufficient for reliably matching these features. However, we demonstrated that our detection system would have discovered the similarities in content if the data would be available in LaTeX or XHTML format.

The dataset of 446 cases supports zbMATH's work towards the goal of installing a system that supports the editor in identifying potentially suspicious documents, even if the final decision is not to issue a public note on content reuse. In other words, the notification threshold of the system needs to be lower than for most plagiarism detection systems. Furthermore, the system must enable the zbMATH editor to easily understand why a document has been retrieved as potentially suspicious. This requirement is even more important than automatically performing a highly accurate binary classification of documents as suspicious or unsuspecting. To achieve this goal, visualizing the topical similarity is a key feature required of the future system.

The next steps for realizing such a system are to establish an automated workflow for receiving the full-texts of the papers submitted to zbMATH. Moreover, we need to obtain mathematical formulae in a machine-readable format for at least a fraction of the zbMATH collection. We will continue our efforts to extract LaTeX formulae from PDF documents and are looking forward to the results of this year's CHROME competition²³. Especially the results for machine-readable formulae will be the foundation to conceive more sophisticated math-based detection methods. In the long run, we plan to lower the notification threshold for content reuse, which will undoubtedly require more sophisticated detection methods for formula similarity. The idea is to also identify papers that did not plagiarize but have limited novelty.

²³ <https://www.cs.rit.edu/~crohme2019/index.html>

Acknowledgements This work was supported by the German Research Foundation (DFG grant GI-1259-1).

References

- [1] A. Aizawa et al. “NTCIR-11 Math-2 Task Overview”. In: *Proc. NTCIR Conf. on Evaluation of Information Access Technologies*. 2014.
- [2] S. M. Alzahrani, N. Salim, and A. Abraham. “Understanding Plagiarism Linguistic Patterns, Textual Features, and Detection Methods”. In: *IEEE Trans. Syst., Man, Cybern. C, Appl. Rev.* Vol. 42. 2. 2012.
- [3] J. B. Baker, A. P. Sexton, and V. Sorge. “MaxTract: converting PDF to LaTeX, MathML and text.” In: *Proc. Int. Conf. on Intelligent Computer Mathematics*. Vol. 7362. 2012. DOI: 10.1007/978-3-642-31374-5_29.
- [4] T. A. E. Eisa, N. Salim, and S. M. Alzahrani. “Existing plagiarism detection techniques: A systematic mapping of the scholarly literature”. In: *Online Information Review* 39.3 (2015).
- [5] T. Fishman. “‘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*. 2009.
- [6] T. Foltyniek, N. Meuschke, and B. Gipp. “Academic Plagiarism Detection: A Systematic Literature Review”. Journal article in review. 2019.
- [7] B. Gipp. *Citation-based Plagiarism Detection - Detecting Disguised and Cross-language Plagiarism using Citation Pattern Analysis*. Springer, 2014. DOI: 10.1007/978-3-658-06394-8.
- [8] B. Gipp and N. Meuschke. “Citation Pattern Matching Algorithms for Citation-based Plagiarism Detection: Greedy Citation Tiling, Citation Chunking and Longest Common Citation Sequence”. In: *Proc. ACM Symposium on Document Engineering (DocEng)*. DOI: 10.1145/2034691.2034741.
- [9] B. Gipp, N. Meuschke, and J. Beel. “Comparative Evaluation of Text- and Citation-based Plagiarism Detection Approaches using GuttenPlag”. In: *Proc. ACM/IEEE-CS Joint Conf. on Digital Libraries (JCDL)*. 2011. DOI: 10.1145/1998076.1998124.
- [10] B. Gipp, N. Meuschke, and C. Breitingner. “Citation-based Plagiarism Detection: Practicability on a Large-scale Scientific Corpus”. In: *JASIST* 65.2 (2014). DOI: 10.1002/asi.23228.
- [11] B. Gipp et al. “Web-based Demonstration of Semantic Similarity Detection using Citation Pattern Visualization for a Cross Language Plagiarism Case”. In: *Proc. Int. Conf. on Enterprise Information Systems*. 2014. DOI: 10.5220/0004985406770683.
- [12] F. Guidi and C. Sacerdoti Coen. “A Survey on Retrieval of Mathematical Knowledge”. In: *Mathem. in Computer Science* 10.4 (2016). DOI: 10.1007/s11786-016-0274-0.
- [13] G. Halevi and J. Bar-Ilan. “Post Retraction Citations in Context”. In: *Proc. BIRNDL Workshop at JCDL*. 2016. DOI: 10.1007/s11192-017-2242-0.

- [14] T. C. Long et al. “Responding to Possible Plagiarism”. In: *Science* 323.5919 (2009). DOI: 10.1126/science.1167408.
- [15] D. L. McCabe. “Cheating among College and University Students: A North American Perspective”. In: *Int. J. for Educational Integrity* 1.1 (2005). DOI: 10.21913/IJEI.v1i1.14.
- [16] N. Meuschke and B. Gipp. “Reducing Computational Effort for Plagiarism Detection by using Citation Characteristics to Limit Retrieval Space”. In: *Proc. IEEE/ACM Joint Conf. on Digital Libraries (JCDL)*. 2014. DOI: 10.1109/JCDL.2014.6970168.
- [17] N. Meuschke and B. Gipp. “State of the Art in Detecting Academic Plagiarism”. In: *Int. J. for Educational Integrity* 9.1 (2013). DOI: 10.21913/IJEI.v9i1.847.
- [18] N. Meuschke et al. “An Adaptive Image-based Plagiarism Detection Approach”. In: *Proc. ACM/IEEE-CS Joint Conf. on Digital Libraries (JCDL)*. 2018. DOI: 10.1145/3197026.3197042.
- [19] N. Meuschke et al. “Analyzing Mathematical Content to Detect Academic Plagiarism”. In: *Proc. ACM Conf. on Information and Knowledge Management, (CIKM)*. 2017. DOI: 10.1145/3132847.3133144.
- [20] N. Meuschke et al. “Analyzing Semantic Concept Patterns to Detect Academic Plagiarism”. In: *Proc. WOSP Workshop held at ACM/IEEE-CS Joint Conf. on Digital Libraries (JCDL)*. 2017. DOI: 10.1145/3127526.3127535.
- [21] N. Meuschke et al. “HyPlag: A Hybrid Approach to Academic Plagiarism Detection”. In: *Proc. Int. ACM SIGIR Conf. on Research and Development in Information Retrieval (SIGIR)*. 2018. DOI: 10.1145/3209978.3210177.
- [22] N. Meuschke et al. “Improving Academic Plagiarism Detection for STEM Documents by Analyzing Mathematical Content and Citations”. In: *Proc. ACM/IEEE-CS Joint Conf. on Digital Libraries (JCDL)*. 2019.
- [23] S. d. L. Pertile, V. P. Moreira, and P. Rosso. “Comparing and combining Content- and Citation-based approaches for plagiarism detection”. In: *JASIST* 67.10 (2016).
- [25] B. Stein, S. Meyer zu Eissen, and M. Potthast. “Strategies for Retrieving Plagiarized Documents”. In: *Proc. Int. ACM SIGIR Conf. on Research and Development in Information Retrieval (SIGIR)*. 2007. DOI: 10.1145/1277741.1277928.
- [26] M. Suzuki et al. “An Integrated OCR Software for Mathematical Documents and Its Output with Accessibility”. In: *In Proc. Int. Conf. Computers for Handicapped Persons*. 2004. DOI: 10.1007/978-3-540-27817-7_97.
- [27] J. P. Swazey, M. S. Anderson, and K. S. Louis. “Ethical Problems in Academic Research”. In: *American Scientist* 81.6 (1993).
- [28] K. Vani and D. Gupta. “Study on Extrinsic Text Plagiarism Detection Techniques and Tools”. In: *J. of Engineering Science & Technology Review* 9.4 (2016).

- [29] E. Wager. “Defining and responding to plagiarism”. In: *Learned Publishing* 27.1 (2014). DOI: 10.1087/20140105.
- [30] D. Weber-Wulff. *False Feathers: A Perspective on Academic Plagiarism*. Springer, 2014. DOI: 10.1007/978-3-642-39961-9.
- [31] D. Weber-Wulff. *Portal Plagiat - Tests of Plagiarism Software*. Online Source. Retrieved March 12, 2019 from: <http://plagiat.htw-berlin.de/software-en/>. 2019.
- [32] M. Wolska. “A Language Engineering Architecture for Processing Informal Mathematical Discourse”. In: *Proc. DML WS Towards Digital Mathematics Library*. 2008.

A List of documents with noticeable content reuse

1360.53021, 1357.30013, 1353.39029, 1353.30019, 1345.15011, 1359.62073, 1356.01026, 1337.16003, 1354.47018, 1340.90030, 1360.47003, 1345.92082, 1318.46035, 1400.34041, 1388.42037, 1388.42036, 1343.65150, 1330.35490, 1322.93076, 1321.81036, 1307.65177, 1308.81133, 1358.47017, 1309.65163, 1304.57008, 1325.47059, 1301.16002, 1293.65167, 1359.62055, 1291.30077, 1323.65125, 1328.47074, 1328.47073, 1299.65168, 1295.35151, 1294.35189, 1290.26023, 1282.91334, 1279.91096, 1281.35058, 1306.90186, 1311.90164, 1273.91086, 1287.81012, 1386.18011, 1301.45006, 1290.18001, 1278.68235, 1278.53033, 1271.54029, 1271.39024, 1266.65214, 1266.33002, 1264.34048, 1342.34118, 1266.30001, 1265.39016, 1264.81239, 1257.11089, 1250.78038, 1246.90035, 1246.90034, 1250.78039, 1252.68177, 1234.34034, 1364.47004, 1399.35153, 1274.76184, 1252.83109, 1288.49015, 1249.60023, 1250.47059, 1231.83033, 1227.34015, 1219.30004, 1236.58009, 1230.46033, 1213.60020, 1211.34093, 1211.34092, 1295.91090, 1242.49079, 1234.60021, 1262.11083, 1221.81113, 1234.60020, 1211.46021, 1217.34137, 1211.11127, 1203.06007, 1203.06006, 1212.49026, 1193.35074, 1191.35223, 1253.60034, 1235.37020, 1186.54007, 1189.35123, 1188.16002, 1183.37156, 1371.91006, 1371.91005, 1258.74210, 1257.78018, 1192.34093, 1195.55004, 1212.60016, 1201.60017, 1184.20030, 1176.91147, 1173.90327, 1206.34097, 1177.35217, 1170.34353, 1173.34354, 1279.90096, 1153.91544, 1189.35124, 1177.35218, 1175.86006, 1162.30319, 1170.42304, 1165.35336, 1162.30309, 1153.86318, 1154.94319, 1250.49003, 1166.47308, 1153.91523, 1155.26016, 1157.05036, 1162.83357, 1139.81335, 1213.35364, 1169.46304, 1169.42310, 1144.81475, 1141.90010, 1231.93121, 1132.14304, 1144.46044, 1134.60382, 1129.83326, 06921286.

Listing 1.2. Use the following BibTeX code to cite this article

```
@inproceedings{Schubotz2019,  
  author      = {Schubotz, Moritz and Teschke, Olaf and Stange,  
                Vincent and Meuschke, Norman and Gipp, Bela},  
  location    = {Czech Republic},  
  booktitle   = {Intelligent Computer Mathematics - 12th  
                International Conference, {CICM} 2019, Prague, Czech  
                Republic, July 8-12, 2019, Proceedings},  
  date        = {2019-07},  
  keywords    = {preprint, SchubotzCV},  
  preprint    = {https://ag-gipp.github.io/bib/preprints/  
                schubotz2019.pdf},  
  title       = {Forms of Plagiarism in Digital Mathematical  
                Libraries},  
}
```