

Preprint of the paper:

Satpute, A. & Greiner-Petter, A. & Gießing, N. & Beckenbach, I. & Schubotz, M. & Teschke, O. & Aizawa, A. & Gipp, B., "Taxonomy of Mathematical Plagiarism", in 46th European Conference on Information Retrieval (ECIR), Cham: Springer Nature Switzerland, 2024.

Click to download: BibTeX

Taxonomy of Mathematical Plagiarism

Ankit Satpute^{1,2}[0000-0003-3219-026X], André Greiner-Petter²[0000-0002-5828-5497], Noah Gießing¹[0009-0006-5268-2519], Isabel Beckenbach¹[0000-0001-6691-7362], Moritz Schubotz¹[0000-0001-7141-4997], Olaf Teschke¹[0009-0003-4089-9647], Akiko Aizawa³[0000-0001-6544-5076], and Bela Gipp²[0000-0001-6522-3019]

¹ FIZ Karlsruhe Leibniz Institute for Information Infrastructure, Berlin, Germany

`FirstName.LastName@fiz-karlsruhe.de`

² Georg August University of Göttingen, Göttingen, Germany

³ National Institute of Informatics, Tokyo, Japan

Abstract. Plagiarism is a pressing concern, even more so with the availability of large language models. Existing plagiarism detection systems reliably find copied and moderately reworded text but fail for idea plagiarism, especially in mathematical science, which heavily uses formal mathematical notation. We make two contributions. First, we establish a taxonomy of mathematical content reuse by annotating potentially plagiarised 122 scientific document pairs. Second, we analyze the best-performing approaches to detect plagiarism and mathematical content similarity on the newly established taxonomy. We found that the best-performing methods for plagiarism and math content similarity achieve an overall detection score (PlagDet) of 0.06 and 0.16, respectively. The best-performing methods failed to detect most cases from all seven newly established math similarity types. Outlined contributions will benefit research in plagiarism detection systems, recommender systems, question-answering systems, and search engines. We make our experiment's code and annotated dataset available to the community: <https://github.com/gipplab/Taxonomy-of-Mathematical-Plagiarism>.

Keywords: Math Reuse · Plagiarism · Math Similarity Taxonomy.

1 Introduction

Plagiarism is “the use of ideas, concepts, words, or structures without appropriately acknowledging the source in settings expecting originality” [7, p. 5]. Plagiarism is a pressing concern as it wastes peer reviewers’ time, funding an unoriginal work and depriving original authors of the benefits [13,12]. Large Language Models (LLMs) make plagiarism easier due to produced unreferenced content [19]. Most plagiarism detection systems (PDS) have addressed identifying unoriginal text [2,15] and, to a lesser extent, unoriginal non-textual content such as formulae or images [8]. Research on plagiarized mathematical content is incipient, but prior studies have highlighted publications with questionable reuse of math content, leading to retractions [18,23].

A common approach [11] to find similar mathematical content is determining whether or not two mathematical formulae are identical. Due to the complex nature of math and possible underlying assumptions not stated explicitly, it is difficult to judge semantic similarities between expressions since similarity in presentation does not imply a semantic similarity. The unavailability of an annotated dataset of similar mathematical content in a machine-processable format like LaTeX or MathML currently hinders progress in detecting human-modified similar mathematical content. For tasks like paraphrase identification (PI), established corpora like ETPC [14] provide a taxonomy of paraphrasing types. Establishing these types has been greatly helpful in shaping advanced PI methods [28,3] to detect various ways similar text can be represented. Thus far, a taxonomy of mathematical content similarity in math plagiarised instances hasn't been produced due to the subject's complicated nature [18], impeding methods to detect human-modified similar math. Establishing annotated corpora would benefit analysis of the performance of existing math similarity detection systems and the development of new systems.

In this work, we analyze and annotate 122 scientific document pairs from zbMATH Open that experts judged to be potential cases of plagiarism. We do not claim that the cases are plagiarized. The final decision of whether they are plagiarized is up to the field experts. We establish a taxonomy of mathematical content similarity by identifying a type of reuse and classifying it into a set of rules that describes the reuse. It has great potential to accelerate LLMs, enabling them to identify similar structural and semantic math contents.

2 Related work

The datasets of the workshop series *Plagiarism Analysis, Authorship Identification, and Near-Duplicate Detection* (PAN) [24] are frequently utilized to develop and evaluate PDS. The PAN datasets consist of artificially created plagiarism, e.g., by randomly removing, inserting, or replacing words or phrases. The representativeness of simulated plagiarism in PAN datasets to real plagiarism is unclear [9], limiting the generalizability of evaluation results from these datasets.

For mathematical content, resources of a similar scale to text reuse are missing. Only two works analyzed mathematical content in scientific documents to identify plagiarism [18,17]. Both studied basic mathematical symbol occurrences and used a small evaluation dataset of 10 document pairs. Currently, No PDS considers semantic textual and non-textual content similarity [15,8]. While identifying mathematical plagiarism has barely been studied, most mathematics-related tasks focused on search engines and similarity analysis, such as NTCIR [26] and ARQMath [16]. In NTCIR-12, given a formulae query, relevant formulae were retrieved from the arXiv collection and Wikipedia articles. The ArqMath-3 [16] competition has released formula pairs from MathStackExchange question-answer pairs. Formulae from the questions and their answers are considered candidates, and the results submitted by participants are ranked on a scale of 0 (Irrelevant) to 3 (Highly relevant) by student annotators. Even though both

NTCIR and ARQMath contain relevant formulae pairs, they were annotated by human post-detection, and none of the pairs have any information about how the content is modified.

Meuschke et al. [18] have formed categories of similar mathematical content by analyzing 10 document pairs retracted for plagiarism collected from computer science, biology, etc., research areas. However, most categories focus on presentation rather than semantic manipulation. The most comprehensive dataset of real-world plagiarised cases in mathematical academia was extracted from zbMATH Open [23]. They analyzed 10 document pairs out of the suspected 446, only visually pointing to similar math in a document under inspection and a potential source document.

3 Dataset

This work uses real-world plagiarised cases in mathematical documents from zbMATH Open [23]. zbMATH Open is the most comprehensive and longest-running (1868 - present) abstracting and reviewing service in mathematics. The majority of 4.5 million entries in zbMATH Open are in English, and even if the underlying article is not in English, the review, summary, or abstract is typically written in English. We use 446 entries from zbMATH Open identified for noticeable content reuse [23] and analyze each entry to annotate which exact parts of the documents are similar, i.e., reuse cases. Further, we categorize each reuse case into a type, eventually forming a taxonomy of reuse. We call these reuse types *obfuscation Operators* since they present obstacles to automatic reuse detection. *obfuscation Operators* does not imply malicious motivation to any author.

Annotation procedure: zbMATH Open does not contain full-text documents. We manually collected the full texts by accessing the publisher or repository hosting them. Unfortunately, we could only obtain full texts for 122 pairs (inspected document - potential source document) out of 446 due to inactive full-text hyperlinks present at zbMATH Open, removal of full texts by publishers due to retraction notices, etc. In most cases, we could find the document in PDF format. Hence, we used MathPiX [1] to convert PDF to LaTeX, eventually having access to all math formulae in machine-processable format LaTeX.

A group of 7 field experts with a postgraduate degree in Mathematics and 6 having Ph.D. annotated similar mathematical contents. They used the publicly available annotation tool TEIMMA [22] for annotating reuse. They regarded content from an inspected document as an instance of content reuse if it overlaps significantly with the content from the source document. They assigned *obfuscation operators* to each case, depicting the possible way in which content from the source document was modified. In most annotations, multiple obfuscation operators were found simultaneously. In the following, we describe each obfuscation operator formed by annotators.

Obfuscation Operator	Abbreviation	Short Explanation	Cases count	
			Comb.	Uniq.
Paraphrasing	P	Text or Math rewording	1354	133
Insertions and Deletions	ID	Insert or delete text/math expressions	766	21
Substitutions	S	Substitute a term/expression or replace by cross reference	94	2
Text to Math/ Math to Text	TMMT	Math formulae explained with text words or vice versa	164	1
Different Presentation	DP	Objects with the same meaning but different presentations	733	20
Formula Manipulation	FM	Substituting formulae with a different repr. of semantic content through transformations	228	2
Variation of Subject	VS	Semantics different but argumentative structure similar	959	57

Table 1. Overview of all obfuscation operators and count of annotated cases per obfuscation operator. Comb: cases occurring with other obfuscation. Uniq: cases occurring with only one particular obfuscation.

3.1 Proposed Taxonomy of Mathematical content similarity:

Table 1 presents a summary of all obfuscations. The definitions and rules for each obfuscation operator were fixed gradually, with frequent discussions with experts with experience reviewing documents for zbMATH Open.

A taxonomy needs a precise vocabulary for categories. Terms like formulae and expressions are often used interchangeably [6,25]. Thus, we initially set definitions for obfuscation operator terms. **Formula:** A combination of mathematical symbols formed in accordance with the rules of mathematical syntax. **Maximal Expressions:** A well-formed expression inside a mathematical formula that cannot be extended further without becoming a mathematical statement. For example, $a + 2$ and 0 in mathematical statement $a + 2 = 0$. Non-maximal expressions will be called subexpressions. For the above example, subexpressions would be a , 2 , and 0 . **Synonymous:** Two expressions of almost identical semantics are synonymous when representing the same mathematical object. For example, 1 and 1.0 are synonymous.

Obfuscation Operators: **1. Paraphrasing (P):** Paraphrases convey the same meaning while using different words or sentence structures [5]. We extend the definition to mathematical content paraphrases by placing two restrictions. First, a mathematical content paraphrase requires that all maximal expressions from the source document must be preserved without changes. In addition, relators (greater than, less than, etc.) may be changed according to necessity.

Example Case: { *Src* :The proof reduces to showing that $x \leq 1$. } { *Insp* : It is sufficient to prove that $1 \geq x$. }

2. Insertions and Deletions (ID) A pair of text passages are subject to Insertions if a similar passage adds redundant or insubstantial content. The

definition for deletions is similar, only reversed: insubstantial content from the source passage is deleted.

Example Case: Notice the middle expression in $Insp : \{ Src : \langle x + y, x + y \rangle \leq \|x\|^2 + 2|\langle x, y \rangle| + \|y\|^2 \}$
 $\{ Insp : \langle x + y, x + y \rangle = \|x\|^2 + \langle x, y \rangle + \langle y, x \rangle + \|y\|^2 \}$
 $\leq \|x\|^2 + 2|\langle x, y \rangle| + \|y\|^2. \}$

3. Substitutions (S) We consider a reuse case as a substitution when either of the following occurs: 1. In formulae: replacement of expression e in the source by another expression f , which is related to e by a formula eRf for some relator R , or vice versa. 2. In-text: replacement of textual content in the source by cross-reference to an earlier passage with equivalent content, or vice versa.

Example Case: Content in $Insp$ eventually leads to Src by substitutions. $\{ Src : x + 3 = y - 2 \}$ $\{ Insp : A(x) = B(y), \text{ where } A(x) = x + 3 \text{ and } B(y) = y - 2. \}$

4. Text to Math/Math to Text (TMMT) In the Text to Math obfuscation operator, mathematical content from the source document expressed in text is substituted by mathematical notations. Similarly, Math to Text describes the reverse obfuscation operation.

Example Case: $\{ Src : \text{The force } F \text{ acting on a body equals the product of its mass } m \text{ and acceleration } a. \}$ $\{ Insp : F = ma. \}$

5. Different Presentation (DP) This operator allows two notational changes. First, an expression is replaced by a synonymous expression. Second, if two expressions are single objects, all root and child objects in one expression are matched to synonymous operators and objects in another.

Example Case: A simple change of variable names. $\{ Src : f(x) \}$ $\{ Insp : g(x) \}$

6. Formula Manipulation (FM) Formula Manipulation refers to substituting formulae with different representations of their semantic content obtained through transforming expressions using algebraic identities or statements using logical equivalences or implications. In contrast to the *substitutions* operator, the equivalence of transformed expressions is not announced elsewhere in the document and has to be logically deduced. Examples of methods to transform expressions could be expanding a single term into multiple terms using rules such as distributivity, associativity, etc.

Example Case: The following example involves applying coordinate changes, e.g., expressing a complex number from Src in polar coordinates in $Insp$. $\{ Src : z = a + bi \}$ $\{ Insp : z = re^{i\phi} \}$

7. Variation of Subject (VS) In the Variation of the Subject obfuscation operator, the semantic content of a passage after the application is strictly different from the original passage, but the argumentative structure remains similar.

Example Case: This is a non-mathematical (Included due to space limitation. Please refer to the repository⁴ for a real example) example that illustrates how content might be modified in several places to accommodate the new setting. $\{ Src : \text{I went to [Paris] for a few days and visited the (Eiffel Tower).} \}$ $\{ Insp : \text{I went to [Rome] for a few days and visited the (Colosseum).} \}$

⁴ <https://github.com/gipplab/Taxonomy-of-Mathematical-Plagiarism>

Annotation verification by an expert: Since similar mathematical content annotations were done in a single instance by field experts, verifying if their annotations are consistent is essential. We provided 10% document pairs to an independent expert with a Ph.D. in mathematics to do similar annotations. Creating manual annotations is time-consuming, and it is hard to find experts in multiple disciplines to produce adequate annotations. We provided obfuscation categories defined above to assign to the selected spans. The independent expert annotated the document pairs using TEIMMA [22]. We use Jaccard similarity [4] to find the overlap between the original annotations and the independent expert’s annotations, i.e., for each document pair. On average, there was 80.01% token overlap (both text and math). Similarly, for exact *Case Type* (Whether reuse contains text, math, or both) and exact *Obfuscation Type* matches, there was 95.21% and 74.23% overlap, respectively. For obfuscation types, we further calculate inter-annotator agreement using the kappa value and achieve a score of 0.39, indicating a fair agreement. We understand that agreeing upon the same obfuscations is difficult as it is subjective.

4 Evaluation

Obfusc.→	All			P(78.55%)			ID(44.86%)			S(5.39%)		
Method↓	F1	G	PD	F1	G	PD	F1	G	PD	F1	G	PD
LCIS	0.00	1.00	0.00	0.00	1.00	0.00	0.00	1.00	0.00	0.00	1.00	0.00
GIT	0.05	1.28	0.04	0.05	1.35	0.04	0.02	1.38	0.02	0.00	1.00	0.00
AdaPlag	0.08	1.34	0.06	0.08	1.39	0.06	0.09	1.29	0.07	0.10	1.36	0.08
MABOWDOR	0.17	1.10	0.16	0.13	1.27	0.11	0.04	1.02	0.04	0.00	1.00	0.00

Obfusc.→	TMMT(9.50%)			DP(45.89%)			FM(12.92%)			VS(55.30%)		
Method↓	F1	G	PD	F1	G	PD	F1	G	PD	F1	G	PD
LCIS	0.00	1.00	0.00	0.00	1.00	0.00	0.00	1.00	0.00	0.00	1.00	0.00
GIT	0.00	2.00	0.00	0.02	2.00	0.01	0.01	2.00	0.01	0.02	1.42	0.02
AdaPlag	0.07	1.38	0.05	0.08	1.30	0.07	0.09	1.33	0.07	0.07	1.17	0.06
MABOWDOR	0.00	1.00	0.00	0.07	1.24	0.06	0.02	1.09	0.02	0.09	1.09	0.09

Table 2. Evaluation of plagiarism detection and Math Content Similarity methods on the new dataset. Values in parentheses indicate the percentage of cases of all cases in which a particular obfuscation is present. (Lower G and higher F1 & PD represents better detection)

We use four detection approaches. The first two are the Longest Common Subsequence of Identifier (LCIS) and Greedy Identifier Tiles (GIT) presented in the only prior work on math plagiarism detection using mathematical content by Meuschke et al. [18]. Third, AdaPlag, the winning approach of the PAN plagiarism competition [21]. Fourth, a Math-Aware Best of-Worlds Domain Optimized Retriever (MABOWDOR) [27], which is the best-performing math content similarity approach on the ARQMath dataset [16].

We create a test collection to represent a real-world retrieval scenario. zbMATH Open does not contain full texts, but around 464K entries correspond to documents from the arXiv. We obtained LaTeX sources of these 464K entries from the arXMLiv dataset [10]. We include the corresponding source documents of all annotated 122 inspected documents in the test collection and remove inspected documents from the test collection to avoid exact matches. For each of the four methods, we query our test collection of 464K documents by each of the 122 inspected documents. We combine two steps (source retrieval and detailed comparison) of typical PDS into one, i.e., direct comparison without any hyperparameters tuning. For evaluation, we use the F1 score, i.e., the harmonic mean of precision and recall), Granularity (G), which determines whether a reuse case was detected as a whole or in several pieces, and PlagDet (PD), which combines recall, precision, and G to allow for ranking, are standard evaluation metrics for plagiarism detection [20]. For each of the 122 inspected documents, we take the top 10 most similar documents out of the test collection and calculate evaluation scores. Table 2 shows the evaluation of approaches on the newly constructed dataset. Results indicate that most of the cases from all obfuscation operators remain undetected by existing methods. Improvements can likely be made by hyperparameter tuning.

5 Conclusion

In this work, we established a novel taxonomy of math content similarity. We formed 7 math content similarity types by annotating 122 document pairs identified for potential plagiarism. We analyzed the performance of the best-performing plagiarism detection methods and math content similarity in detecting cases from newly established taxonomy. It was found that the current best-performing methods do not detect most human modifications on math content (Overall PlagDet score of 0.06 and 0.16 for the best-performing plagiarism and math content similarity methods, respectively). The dataset curated in this work and the code of the experiments are publicly available to aid future research on math content similarity detection. A resource and experiments presented in this paper set up a base to develop advanced detection techniques capable of identifying modified math content reuse. This work will accelerate research identifying concealed plagiarism instances in academia and research involving mathematical content similarity.

Acknowledgements

This work was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) - 437179652, the Deutscher Akademischer Austauschdienst (DAAD, German Academic Exchange Service - 57515245), and the Lower Saxony Ministry of Science and Culture and the VW Foundation.

References

1. Mathpix: AI-powered document automation. — mathpix.com. <https://mathpix.com/>, [Accessed 15-10-2023]
2. Alvi, F., Stevenson, M., Clough, P.: Plagiarism detection in texts obfuscated with homoglyphs. In: Jose, J.M., Hauff, C., Altingovde, I.S., Song, D., Albakour, D., Watt, S., Tait, J. (eds.) *Advances in Information Retrieval (European Conference on Information Retrieval)*. pp. 669–675. Springer International Publishing, Cham (2017)
3. Babakov, N., Dale, D., Logacheva, V., Panchenko, A.: A large-scale computational study of content preservation measures for text style transfer and paraphrase generation. In: *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics: Student Research Workshop*. pp. 300–321. Association for Computational Linguistics, Dublin, Ireland (May 2022). <https://doi.org/10.18653/v1/2022.acl-srw.23>, <https://aclanthology.org/2022.acl-srw.23>
4. Bank, J., Cole, B.: Calculating the jaccard similarity coefficient with map reduce for entity pairs in wikipedia. *Wikipedia Similarity Team* **1**, 94 (2008)
5. Bhagat, R., Hovy, E.: What Is a Paraphrase? *Computational Linguistics* **39**(3), 463–472 (09 2013). <https://doi.org/10.1162/COLLa.00166>, <https://doi.org/10.1162/COLLa.00166>
6. Edwards, B.S., Ward, M.B.: Surprises from mathematics education research: Student (mis)use of mathematical definitions. *The American Mathematical Monthly* **111**(5), 411–424 (2004). <https://doi.org/10.1080/00029890.2004.11920092>, <https://doi.org/10.1080/00029890.2004.11920092>
7. Fishman, T.: “we know it when we see it” is not good enough: Toward a standard definition of plagiarism that transcends theft, fraud, and copyright (2009)
8. Foltýnek, T., Meuschke, N., Gipp, B.: Academic Plagiarism Detection: A Systematic Literature Review. *ACM Computing Surveys* **52**(6), 112:1–112:42 (Oct 2019). <https://doi.org/10.1145/3345317>
9. Foltýnek, T., Meuschke, N., Gipp, B.: Academic plagiarism detection: a systematic literature review. *ACM Computing Surveys (CSUR)* **52**(6), 1–42 (2019)
10. Ginev, D.: arxmliv:2020 dataset, an html5 conversion of arxiv.org. hosted at <https://sigmathling.kwarc.info/resources/arxmliv-dataset-2020/> (2020), sIGMathLing – Special Interest Group on Math Linguistics
11. Guidi, F., Sacerdoti Coen, C.: A survey on retrieval of mathematical knowledge. *Mathematics in Computer Science* **10**(4), 409–427 (2016)
12. Hall, S.E.: *Is it happening? how to avoid the deleterious effects of plagiarism and cheating in your courses*. vol. 74, pp. 179–182. SAGE Publications Sage CA: Los Angeles, CA (2011)
13. Hart, K., Mano, C., Edwards, J.: Plagiarism deterrence in cs1 through keystroke data. In: *Proceedings of the 54th ACM Technical Symposium on Computer Science Education V. 1*. p. 493–499. SIGCSE 2023, Association for Computing Machinery, New York, NY, USA (2023). <https://doi.org/10.1145/3545945.3569805>, <https://doi.org/10.1145/3545945.3569805>
14. Kovatchev, V., Martí, M.A., Salamó, M.: Etpc-a paraphrase identification corpus annotated with extended paraphrase typology and negation. In: *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)* (2018)

15. Lovepreet, V.G., Kumar, R.: Survey on plagiarism detection systems and their comparison. In: Computational Intelligence in Data Mining: Proceedings of the International Conference on ICCIDM 2018. vol. 990, p. 27. Springer (2019)
16. Mansouri, B., Agarwal, A., Oard, D.W., Zanibbi, R.: Advancing math-aware search: The arqmath-3 lab at clef 2022. In: European Conference on Information Retrieval. pp. 408–415. Springer (2022)
17. Meuschke, N., Schubotz, M., Hamborg, F., Skopal, T., Gipp, B.: Analyzing mathematical content to detect academic plagiarism. In: Proceedings of the International Conference on Information and Knowledge Management (CIKM) (2017). <https://doi.org/10.1145/3132847.3133144>
18. Meuschke, N., Stange, V., Schubotz, M., Kramer, M., Gipp, B.: Improving academic plagiarism detection for stem documents by analyzing mathematical content and citations. In: Proceedings of the ACM/IEEE Joint Conference on Digital Libraries (JCDL) (Jun 2019). <https://doi.org/10.1109/JCDL.2019.00026>
19. Meyer, J.G., Urbanowicz, R.J., Martin, P.C., O’Connor, K., Li, R., Peng, P.C., Bright, T.J., Tatonetti, N., Won, K.J., Gonzalez-Hernandez, G., et al.: Chatgpt and large language models in academia: opportunities and challenges. *BioData Mining* **16**(1), 20 (2023)
20. Potthast, M., Stein, B., Barrón-Cedeño, A., Rosso, P.: An evaluation framework for plagiarism detection. In: Coling 2010: Posters. pp. 997–1005. Coling 2010 Organizing Committee, Beijing, China (Aug 2010), <https://aclanthology.org/C10-2115>
21. Sanchez-Perez, M.A., Gelbukh, A., Sidorov, G.: Adaptive algorithm for plagiarism detection: The best-performing approach at pan 2014 text alignment competition. In: International Conference of the Cross-Language Evaluation Forum for European Languages. pp. 402–413. Springer (2015)
22. Satpute, A., Greiner-Petter, A., Schubotz, M., Meuschke, N., Aizawa, A., Teschke, O., Gipp, B.: Teimma: The first content reuse annotator for text, images, and math. In: 2023 ACM/IEEE Joint Conference on Digital Libraries (JCDL) (June 2023). <https://doi.org/10.1109/JCDL57899.2023.00056>
23. Schubotz, M., Teschke, O., Stange, V., Meuschke, N., Gipp, B.: Forms of plagiarism in digital mathematical libraries. In: Intelligent Computer Mathematics - 12th International Conference, CICM 2019, Prague, Czech Republic, July 8-12, 2019, Proceedings (2019). https://doi.org/10.1007/978-3-030-23250-4_18
24. Stein, B., Koppel, M., Stamatatos, E.: Plagiarism Analysis, Authorship Identification, and Near-Duplicate Detection PAN’07. *ACM SIGIR Forum* **41**(2), 68–71 (Dec 2007). <https://doi.org/10.1145/1328964.1328976>
25. Torkildsen, H.A., Forbregd, T.A., Kaspersen, E., Solstad, T.: Toward a unified account of definitions in mathematics education research: a systematic literature review. *International Journal of Mathematical Education in Science and Technology* pp. 1–28 (2023)
26. Zanibbi, R., Aizawa, A., Kohlhase, M., Ounis, I., Topic, G., Davila, K.: Ntcir-12 mathir task overview. In: NTCIR (2016)
27. Zhong, W., Lin, S.C., Yang, J.H., Lin, J.: One blade for one purpose: Advancing math information retrieval using hybrid search. In: Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval. pp. 141–151 (2023)
28. Zhou, C., Qiu, C., Acuna, D.E.: Paraphrase identification with deep learning: A review of datasets and methods. *arXiv preprint arXiv:2212.06933* (2022)