

WHY AND HOW TO TRANSPARENTLY DISCLOSE AI USE IN MATHEMATICS RESEARCH

MAXIMILIAN KAIPEL

ABSTRACT. This manuscript advocates for transparent AI disclosure in mathematics research and establishes three guiding principles: (1) Only humans can take responsibility and assume authorship; (2) Disclosure should be as specific as possible; (3) All relevant interactions with AI technologies should be documented. Additionally, it develops a taxonomy for the delegation of mathematics research processes to AI technologies and a $\text{\LaTeX}$ package for implementation.

1. INTRODUCTION

There is an on-going discussion within the mathematics community about possible futures for mathematics in the age of Artificial Intelligence (AI)¹ and a call for the development of guidelines for AI use in mathematics [46]. The purpose of this article is to contribute to this discussion by advocating for transparent disclosure of AI use in mathematics research and by providing concrete and actionable recommendations for the implementation of such a transparent disclosure. As such its target audience is research mathematicians rather than AI labs with access to proprietary AI models [39]. Ideas in this document are intended to inspire the development of new guidelines regarding scientific publishing at a time where scientific publishers are concerned with increasing numbers of AI-assisted submissions [13].

Commercial Large Language Models (LLMs) pose serious ethical issues both generally [8, 10, 27, 53] as well as for mathematics research specifically [1, 5, 46, 60]. As detection of AI-generated text remains unreliable [34, 59], a complete ban on the use of AI technologies for research is currently infeasible and could cause serious reputational damage to unjustly accused researchers. Instead, I argue that scientific journals should enforce the inclusion of maximally transparent AI disclosures in mathematics research articles. Indeed, “transparency” is already listed in the first item of the guiding principles of the UNESCO Recommendation on Open Science [58]. Additionally, some of the leading advocates for AI use in mathematics research have recently urged for transparent AI disclosures [20, 56].

Mathematicians must decide individually whether, and to what extent, they use AI tools in their research, see Section 2.4. As it stands, the use of AI technologies for mathematics research has drastically increased the production of research articles in recent months [28]. As of August 2026, almost 25% of mathematics articles uploaded to the preprint server arXiv disclosed the use of AI to some extent [28]. Without a consensus on what a disclosure of AI use must contain, researchers have found a significant disparity between expectations of AI disclosures and actual AI disclosures in computer science [40]. It is likely that mathematicians experience a similar disconnect between expectations and the reality of AI disclosures. In order to overcome this issue, it is important to discuss stigmatisation of AI use [54] and to reach a communal agreement on acceptable research practices. This document is intended to support this latter process and develops the following three general guidelines for transparent AI disclosure in mathematics:

[Principle 1] Only humans are able to assume responsibility for the quality of mathematics research and may therefore assume authorship of a research article in mathematics.

¹see <https://proofsandprompts.com/>

[Principle 2] AI disclosure should be of the highest possible technological, functional and locational specificity.

[Principle 3] Every interaction with AI which has influenced the final presentation or content of a manuscript should be recorded in full and made available.

These principles, and how to apply them, are discussed in detail in Section 2. As a practical companion to Principle 2, this article develops the *Mathematics AI Delegation Taxonomy* (MAIDeT) for the delegation of mathematics research tasks to AI in Section 3. Other groups of researchers have developed similar taxonomies, without a focus on mathematics, for the sake of transparency in the research process [33, 55]. However, as many branches of (theoretical) mathematics research do not rely on empirical data, they often do not include sections on methodology or ethics, rendering large parts of the existing taxonomies irrelevant to mathematicians. This distinction is amplified by mathematics-specific research tasks such as “Example generation” or “Statement generation”, which fit, at best, ambiguously into existing taxonomies. Without an adequate way of disclosing some of their key tasks, mathematicians are in want of a taxonomy appropriate for their workflows. The MAIDeT, whose focus lies on pre-publication tasks related to research production and manuscript preparation, serves exactly this purpose.


The L^AT_EX-package `mathAIdisclose` is developed in parallel to this manuscript². It supports the implementation of Principle 1, Principle 2 and the MAIDeT for mathematicians using L^AT_EX. It builds on the `cleveref` package³ to enable locational specificity by referencing labels of indexed lemmas, definitions or sections. A brief showcase is included in Section 4.

2. GUIDING PRINCIPLES

Many of the ideas presented in this section have been discussed at length by other mathematicians and scholars. Inspiration for the guiding principles is drawn from existing guidelines instated by academic publishers [12, 16, 21, 61], the work of scholars on research ethics [6, 25, 49] as well as AI-critical declarations by mathematicians [1, 5], which are refined and presented for use in a mathematical context. Most generally, the inclusion of an “AI contribution disclosure” section at the end of the introduction section or at the end of the research article is strongly preferred over the inclusion of an AI use disclosure in an “Acknowledgements” section, which should be reserved for humans [25].

In this section, AI and LLM are used synonymously. At the current time, most interactions with AI technologies for the purpose of mathematics research happen through chat interfaces using natural language. Nonetheless, it is important to highlight that machine learning has already successfully been used for mathematics research before the advent of LLMs [14]. And indeed, one feasible future of AI-assisted mathematics may be to revert to smaller, more specialised and controllable models.

The Human-AI-interaction discussed in this article can be depicted as a cycle



which eventually terminates in the production of a mathematics research paper. A transparent AI use disclosure allows readers to correctly identify the contributions made by AI technologies in this process.

²see <https://ctan.org/pkg/mathaidisclose>

³see <https://www.dr-qubit.org/cleveref.html>

2.1. Who should be author? Presumably driven by their commercial interests, AI companies usually present the name of their LLM, or their company name, as the sole author on their mathematics research articles [11, 42, 45]. Alongside their scientific articles, these AI companies publish press releases attempting to contextualise their findings [2, 43, 44]. In reality, the nebulous reporting of the production of these mathematical results raises more questions than it answers. For instance, OpenAI’s statement completely omits the names of those humans they claim were involved in the preparation of the manuscripts [44]. OpenAI justifies this approach by claiming that “attribution should honestly reflect how a result was produced: claiming human authorship for a proof generated entirely by an AI system would misrepresent both the system’s contribution and the nature of genuine human intellectual work” [44]. Anthropic, in contrast, does name some of their employees in the press release [2]. Inexplicably, Jarred Sumner is deemed to be “in every meaningful sense the paper’s human co-author”, but he is only mentioned in the acknowledgements section of the research article [11]. Moreover, two Anthropic mathematicians are said to “take responsibility for its communication” without assuming any credit for the paper’s findings [11]. Some scholars support this standard, arguing that omitting AI from authorship when it has made a significant contribution unjustly gives credit to the human authors [26, 51].

Others argue that this view is ethically problematic and contradicts widely accepted journal guidelines [25]. They argue that “accountability and credit are two sides of the same coin, and contributors cannot have one without the other” [25]. A large portion of the mathematical community supports this conviction, which is reflected in Value 2 of the Leiden Declaration on Artificial Intelligence and Mathematics (LDAIM). It postulates that results should be “attributable to specific authors who take credit for their discovery and assume responsibility for their correctness” [1]. Indeed, the Policy on Authorship of the Association for Computing Machinery declares clearly that anyone listed as an author must be “an identifiable human being” and “take full responsibility for all content in the published Works” [12]. Moreover, arXiv’s moderation policy states that authors “each individually take full responsibility for all ... contents, irrespective of how the contents were generated” and that “generative AI language tools should not be listed as an author” but instead reported [3]. Because AI technologies are not able to assume responsibility, I argue that they do not have the qualities required to assume (co)authorship of research articles.

[Principle 1] Only humans are able to assume responsibility for the quality of mathematics research and may therefore assume authorship of a research article in mathematics.

It is important that mathematicians using AI tools for research are aware of their responsibility to ensure mathematical rigour and correct attribution of the ideas contained in the manuscript. AI technologies do not replace independent thought and verification, least of all peer-review. For this reason, authors are expected to assume definitive responsibility of the contents of AI-assisted research articles. Authors should declare explicitly that they have verified all AI-generated outputs included, to any extent and in any form, in their research article for correctness and that they assume full responsibility for its contents.

2.2. How specific should AI disclosure be? A maximally transparent AI use disclosure must necessarily be of maximal specificity. It is therefore not sufficient to advise authors to disclose, broadly, *how* they used AI tools in their research. Given the diverse capabilities of commercial LLMs, such a lax guideline would leave too much room for interpretation. Instead, an independent reader should be able to answer the following questions: “Which AI tools were used?”, “What did the author(s) use them for?” and “Where in the document were they used?”. In answering these three questions, authors address the following three complementary types of specificity in AI use disclosure:

technological specificity

functional specificity

locational specificity

Technological specificity refers to the disclosure of the specific AI technology used, including the model and any additional attributes. For example, Anthropic’s Claude Opus 5 with adaptive thinking and max effort or OpenAI’s GPT-5.6 Sol with max reasoning.

Authors may raise two possible oppositions to this advice: The irreproducibility of AI outputs, and AI disclosures providing free advertising for AI companies. Indeed, irreproducibility is an important concern as it also violates the first guiding principle of the UNESCO Recommendation on Open Science [58]. Purely mathematically, one may argue that LLMs are deterministic computer processes, not dissimilar to pseudo-random random number generators. Mathematicians may thus be convinced that theoretical reproducibility of AI-generated output demands a maximally transparent AI disclosures. Of course, LLMs are not deterministic in practice, not only due to irreproducible input seeds, but also due to technological constraints such as the non-associativity of floating-point arithmetic and batching [4, 24, 63].

At the same time, further advertising the mathematical capabilities of commercial LLMs poses a threat to the careers of research mathematicians [1] and to mathematics research itself [5]. This worry appears to be justified, as Google DeepMind has already exploited AI use disclosures in mathematics articles for benchmarking purposes [19]. For this reason, some authors choose not to specify the AI technology used in their research [31]. However, one unethical act does not justify an unethical act in response. I take the stance that neither of those arguments justify a nebulous AI use disclosure, but instead make a further case for maximum transparency.

Functional specificity describes the capacity in which AI influenced the research process and the preparation of the manuscript. Authors may delegate various different tasks, such as example generation or proof verification, to AI technologies and should communicate these transparently. In Section 3, the Mathematics AI Delegation Taxonomy (MAIDeT) is developed for this purpose.

Locational specificity concretises the scale at which AI technologies have impacted mathematics research. For this, view mathematics articles as being divided into three levels: Indexed environments, such as lemmas, definitions or theorems; sections; and the entirety of the article. Locational specificity requires that AI disclosure is as specific as possible without being repetitive. For instance, if AI generates a proof of Lemma 2.5, then the disclosure should say so. If AI was used to reformat a section, but not the entire paper, this should be specified.

[Principle 2] AI disclosure should be of the highest possible technological, functional and locational specificity.

This principle provides transparency for the human processed AI output in the Human-AI-interaction loop. It remains to discuss transparent documentation of the actual interaction between human and technology.

2.3. Why should AI interactions be documented? As it stands, other researchers are not able to reproduce any reported interactions with LLMs, even if they have the authors’ prompts available to them [4, 24, 63]. This makes it important to document the Human-AI-interaction throughout the entire scientific process [32]. Documenting a natural language exchange with AI is often a matter of exporting a chat as a PDF document. Alternatively, written prompts and outputs can easily be screenshot and collected as a coherent documentation. If voice commands are used or audio outputs are generated, recording and transcribing these is the most transparent and straight-forward way to document spoken interactions with LLMs. This documentation can then be added as auxiliary files when uploading the manuscript to the preprint server arXiv and shared with academic journals on submission as their guidelines require.

Transparent documentation also provides a basis for estimating the number of LLM tokens used in the interaction. The inclusion of accurate token use statistics would provide insight into financial and environmental costs of AI interactions [53] and help illuminate the question of sustainability of AI use for mathematics.

[Principle 3] Every interaction with AI which has influenced the final presentation or content of a manuscript should be recorded in full and made available.

Authors may be concerned about stigmatisation when disclosing AI use [54], especially when sharing the entire Human-AI-interaction. However, in terms of transparency, reproducibility and correct attribution of ideas, non-disclosure of AI tool use is unacceptable academic misconduct.

2.4. Visibility of AI-critical positions. There are many ethical issues pertaining to commercial LLMs used in mathematics. Biases [8], exploitation of humans during data set creation [47], privacy and security concerns [62], environment and financial impacts [53], correct attribution [10] and disparities in control and access [50] pose ethical concerns for LLM technologies generally. Within mathematics, the devaluing of certain achievements, such as the construction of counterexamples to long-standing conjectures, the solving of landmark problems without inspiring new research directions and the disruption of existing incentive structures are among the issues addressed in a panel discussion at the Heidelberg Laureate Forum in 2026 [20]. In addition, the misalignment of mathematicians’ interests with those of AI companies [5] may harm the perception of mathematics held by the public, universities and funding bodies [1, 60], resulting in an existential threat to the careers of research mathematicians. These issues, together with those of transparency discussed in Section 2.2, compile an excessive list of reasons for mathematicians to boycott the use of currently available commercial AI technologies for their research [60]. In fact, a large group of mathematicians has already taken such a stance and declared not to use AI in any professional capacity⁴. For those holding this conviction, I strongly recommend including in their research outputs a statement sharing this stance for the sake of visibility of AI-critical positions.

3. TAXONOMY

This section develops a taxonomy for naming and organising mathematics research tasks delegated to AI technologies. Such a repertoire of labels for common research tasks is conducive to transparent AI disclosure [33, 38, 55]. Its goal is not to formalise mathematics research but instead to serve as a foundation for implementing functional specificity, see Section 2.2, in AI disclosure and thus decrease the disparity between researchers’ expectations of AI disclosures and reality [40]. Many scientific journals apply the Contributor Roles Taxonomy⁵ to delineate authors’ contributions to a research output. However, AI technologies do not qualify for authorship, see Section 2.1, thus demanding a new framework for documenting the role of AI technologies in scientific research. Taxonomies for this purpose were developed in [33, 55] and act as inspiration for the *Mathematics AI Delegation Taxonomy* (MAIDeT) discussed in this section. As mentioned in the introduction, the research processes in mathematics differ significantly from those in many other disciplines, rendering existing taxonomies largely irrelevant for them and motivating the development of MAIDeT.

3.1. How the taxonomy was created. The development of this taxonomy follows a combination of the deductive and the intuitive approach to taxonomy creation [41]. It draws from existing taxonomies for AI use [33, 55] as well as from resources on mathematical writing [52], proposed mathematics research workflows [48, 18, 22, 35, 56], reports on “Human-AI collaboration in mathematics” [17, 30] and the author’s own research experience in abstract algebra. MAIDeT distinguishes between macro-level tasks, see Fig. 1, representing mathematics research in five non-linear stages, and micro-levels tasks, subdividing these into concrete steps of the research process. It is important to stress that many mathematicians do not view mathematics research solely as a theorem-production pipeline, but rather as a dynamic endeavour to enhancing human understanding of mathematics [7, 57]. Theorem-proving-above-all-else is a philosophy held by some mathematicians [9] but when treated as universal consensus it presents a misguided narrative which underpins the current

⁴see <https://www.ahmath.org/>

⁵see <https://credit.niso.org/>




FIGURE 1. The macro-level themes of MAIDeT

disruption to mathematics caused by AI companies. Open problems traditionally also acted as indicators for our understanding of mathematics, rather than their solutions being mere desiderata. Nonetheless, inspiration for the formulation of the macro-level themes is drawn from discussions of mathematics research pipelines [15, 56, 57] as well as from the author’s own experience as a mathematics researcher.

The micro-level tasks under *Objective identification* are based on articles on how to identify research questions [22] and the authors’ experience. Inspiration for all micro-level tasks was drawn to a large extent from accounts of mathematicians’ reflections on research before the age of AI [18, 35, 48] and those after [17, 30]. To a lesser extent, ideas contained in the surprisingly vast number of scholarly articles discussing success-optimising AI use for mathematics research are incorporated [23, 29, 36, 37, 64]. It is important to remark that this taxonomy is intended to provide a starting point for a coherent AI disclosure across mathematics. It should be viewed as a dynamic framework, which may grow according to its needs in the future.

The full taxonomy is displayed in Table 1, for which it is necessary to define the intended meaning of some non-standard terms. *Applicability assessment* describes the search for (further) applications of existing results or knowledge. *Axiomatisation* summarises the process of defining a new concept or formulating a statement by declaring a collection of properties of an existing concept, proof or construction as axioms. This includes, for example, removing assumptions made in a theorem which are not necessary for the correctness of its conclusion. *Example generation* encompasses also the generation of counterexamples and nonexamples. *Statement generation* refers to the act of determining the precise formulation of a definition, lemma or theorem in order to convey the intended mathematical meaning. *Proof polishing* is the process of making a mathematical proof easier for others to follow. Finally, *Proof exposition* is the act of conveying the key ideas of a mathematical proof with the aim of maximising understanding while usually compromising logical stringency.

4. L^AT_EX IMPLEMENTATION

In order to facilitate easy-to-use implementation of the ideas presented in this manuscript, the L^AT_EX-package `mathAIdisclose` has been developed in parallel. It builds on the `cleveref` package to enable locational specificity by referencing labels of lemmas, definitions or sections.

4.1. How it works. The package fundamentally offers two disclosure styles: One produces a complete disclosure at the end of the article, one produces a footnote for each disclosure made inside of an indexed environment, such as a lemma definition or section. Depending on the intended disclosure style, the package should be included, using either `\usepackage{mathAIdisclose}` or `\usepackage[footnote]{mathAIdisclose}`. The most important command of the package is `\printAIdisclosure[]`

Macro level	Micro level
Knowledge acquisition	Axiomatisation Example generation Literature search Pattern observation
Objective identification	Applicability assessment Conjecture generation Feasibility assessment Idea generation Novelty assessment Objective formulation Objective subdivision Research gap identification
Proof construction	Statement generation Proof completion Proof generation Proof verification Proof polishing
Code and computation	Algorithm conceptualisation Algorithm implementation Code generation Code optimisation Computation Computation verification Enumeration Formal verification Process automation
Presentation and communication	Context discussion Literature review Proof exposition Reformatting Spell-checking Text generation Translation Trend identification Visualisation

TABLE 1. MAIDeT - A taxonomy for the delegation of research tasks from mathematicians to AI technologies.

which produces a subsection titled “AI contribution disclosure”. If the square brackets contain text, then that text will instead be the title of the generated subsection. Because of how `\cleveref` works, this disclosure needs to be included at the end of the document when not choosing the `\footnote` option. The three main commands for disclosing AI use are the following:

```
\discloseAI[model]{task}   \globalAIuse[model]{task}   \customAIdisclose{text}
```

The command `\customAIdisclose{text}` prints the input `text` in the AI disclosure section. This provides the flexibility to add details which do not fit into the other two methods for disclosing AI use, and allows authors to write custom AI disclosures. Additionally, this command should be

used to affirm that the authors have verified the mathematical correctness of the article and assume full responsibility for its contents. The disclosure intentionally does not provide a standardised clause to this effect in order to guarantee intentionality in such a statement.

The command `\globalAIuse[model]{text}` is intended to disclose the use of AI which has affected the entirety of the manuscript, see Section 2.2. It prints the following sentence in the AI contribution disclosure section:

Throughout the manuscript, the authors disclose the delegation of task to model.

If the command is used additional times, it will add the following sentence immediately after:

They also disclose the delegation of task to model.

The command `\discloseAI[model]{task}` must be included inside of an indexed environment which `cleveref` is able to reference, say in Lemma 3.2. Depending on whether the `footnote` option is included, the package either produces the following sentence in the AI contribution disclosure section

The authors disclose the delegation of task to model in Lemma 3.2.

or prints this sentence as an unindexed footnote on the page that contains Lemma 3.2. Including none of the three commands for disclosing AI in the document, leads to the output of `\printAIdisclosure[]` stating the following:

There is no use of Artificial Intelligence to disclose.

4.2. An example. To illustrate how these commands work consider the following $\text{\LaTeX}$ -code:

```
\begin{example}
\discloseAI[Claude Opus 5 on max effort]{Example generation}
There exists an even prime number, 2.
\end{example}

\globalAIuse[ChatGPT 5.5 Pro on thinking mode]{Spell-checking}

\customAIdisclose{We used Grok 4.6 for creating a blog post on this article.\\
The authors have verified all AI-generated outputs included, to any extent and in
any form, in this version of the article for correctness. They assume full
responsibility for the contents of this manuscript.}

\printAIdisclosure[]
```

If the `footnote` option is not loaded, then its output will be the following:

Example 4.1. There exists an even prime number, 2.

AI contribution disclosure. Throughout the manuscript, the authors disclose the delegation of Spell-checking to ChatGPT 5.5 Pro on thinking mode.

The authors disclose the delegation of Example generation to Claude Opus 5 on max effort in Example 4.1.

We used Grok 4.6 for creating a blog post on this article.
The authors have verified all AI-generated outputs included, to any extent and in any form, in this version of the article for correctness. They assume full responsibility for the contents of this manuscript.

If the `footnote` option is loaded, then its output will be as follows. Note that it creates an unindexed footnote at the bottom of the next page, after Example 4.1 is printed.

Example 4.1. There exists an even prime number, 2.

AI contribution disclosure. Throughout the manuscript, the authors disclose the delegation of Spell-checking to ChatGPT 5.5 Pro on thinking mode.

We used Grok 4.6 for creating a blog post on this article.

The authors have verified all AI-generated outputs included, to any extent and in any form, in this version of the article for correctness. They assume full responsibility for the contents of this manuscript.

There is also a package option `solo` to obtain the singular person version of these declarations. This concludes the showcase of the package `mathAIdisclose`.

Acknowledgements. The author thanks Cliodhna Hughes and Calvin Pfeifer for providing helpful feedback on earlier versions of this manuscript.

AI contribution disclosure. The author has no use of Artificial Intelligence to disclose.

REFERENCES

- [1] Jarod Alper et al. *Leiden Declaration on Artificial Intelligence and Mathematics*. 2026. DOI: 10.5281/zenodo.20302944.
- [2] Anthropic. *Learning more about Claude’s mathematical capabilities*. <https://www.anthropic.com/research/riemann-zeta>. Accessed: 2026-09-01.
- [3] arXiv. *Policies: arXiv moderation*. <https://info.arxiv.org/help/moderation/index.html>. Accessed: 2026-09-22.
- [4] Berk Atil et al. *Non-Determinism of “Deterministic” LLM Settings*. 2025. arXiv: 2408.04667.
- [5] Artur Avila et al. *A Severe Misalignment of AI in Mathematics*. <https://mathandai.org/>. Accessed: 2026-09-12.
- [6] Ahmed S BaHammam. “The transparency paradox: why researchers avoid disclosing AI assistance in scientific writing”. In: *Nature and Science of Sleep* 2025.17 (2025), pp. 2569–2574.
- [7] David Bessis. *The fall of the theorem economy*. <https://davidbessis.substack.com/p/the-fall-of-the-theorem-economy/>. Accessed: 2026-10-01.
- [8] Su Lin Blodgett et al. “Language (Technology) is Power: A Critical Survey of “Bias” in NLP”. In: *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. Ed. by Dan Jurafsky et al. Association for Computational Linguistics, 2020, pp. 5454–5476.
- [9] Kevin Buzzard. *To grieve, or not to grieve?* <https://xenaproject.wordpress.com/2026/10/01/to-grieve-or-not-to-grieve/>. Accessed: 2026-10-03.
- [10] Courtnei Byun, Piper Vasicek, and Kevin Seppi. “This reference does not exist: an exploration of LLM citation accuracy and relevance”. In: *Proceedings of the Third Workshop on Bridging Human–Computer Interaction and Natural Language Processing*. 2024, pp. 28–39.
- [11] Anthropic Claude. *More than two thirds of the zeros of the Riemann zeta function are simple and on the critical line*. <https://www-cdn.anthropic.com/95c246936988e43127bc6b2ceb7077c1dad2d68e.pdf>. Accessed: 2026-09-01.
- [12] Association for Computing Machinery. *ACM Policy on Authorship*. <https://www.acm.org/publications/policies/new-acm-policy-on-authorship>. Accessed: 2026-09-22.
- [13] Gemma Conroy. “How ChatGPT and other AI tools could disrupt scientific publishing”. In: *Nature* 622 (2023), pp. 234–236.
- [14] Alex Davies et al. “Advancing mathematics by guiding human intuition with AI”. In: *Nature, London* 600.7887 (2021), pp. 70–74.

The authors disclose the delegation of Example generation to Claude Opus 5 on max effort in Example 4.1.

- [15] Inc. Education Development Center. *Making Mathematics: Mathematics Research Teacher Handbook*. <https://www2.edc.org/makingmath/handbook/Teacher/IntroductoryExplorations/IntroductoryExplorations.pdf>. Accessed: 2026-09-21. 2002.
- [16] Elsevier. *Generative AI policies for journals*. <https://www.elsevier.com/about/policies-and-standards/generative-ai-policies-for-journals>. Accessed: 2026-09-01.
- [17] Harihusa Enomoto. *I Quit Math, Then Wrote a Paper with AI*. <https://haruhisa-enomoto.github.io/quotient-submodule-equidistribution-essay/>. Accessed: 2026-08-31.
- [18] Pavel Etingof. *PRIMES: How to Succeed in Mathematical Research*. <https://math.mit.edu/research/highschool/primes/succeed.php>. Accessed: 2026-08-31.
- [19] Tony Feng et al. *Towards Autonomous Mathematics Research*. 2026. arXiv: 2602.10177.
- [20] Heidelberg Laureate Forum. *Panel Discussion: AI in Mathematical Research — September 15*. https://youtu.be/H7_d_sgui6o?si=8obmDg0IsYaNYLMm&t=5120. Accessed: 2026-09-21.
- [21] Taylor & Francis. *Using AI tools in your research*. <https://authorservices.taylorandfrancis.com/editorial-policies/using-ai-in-your-research-and-manuscript-preparations/>. Accessed: 2026-09-01.
- [22] P. R. Halmos. “The Heart of Mathematics”. In: *The American Mathematical Monthly* 87.7 (1980), pp. 519–524.
- [23] Jonas Henkel. “The mathematician’s assistant: integrating AI into research practice”. In: *Math. Semesterber.* 72.2 (2025), pp. 117–144.
- [24] Andreas Hochlehnert et al. *A Sober Look at Progress in Language Model Reasoning: Pitfalls and Paths to Reproducibility*. 2025. arXiv: 2504.07086.
- [25] Mohammad Hosseini, David B Resnik, and Kristi Holmes. “The ethics of disclosing the use of artificial intelligence tools in writing scholarly manuscripts”. In: *Research Ethics* 19.4 (2023), pp. 449–465.
- [26] Ryan Jenkins and Patrick Lin. *AI-Assisted Authorship: How to Assign Credit in Synthetic Scholarship*. <https://ssrn.com/abstract=4342909>. Accessed: 2026-09-01. 2023.
- [27] Junfeng Jiao et al. “Navigating LLM ethics: Advancements, challenges, and future directions”. In: *AI and Ethics* 5.6 (2025), pp. 5795–5819.
- [28] Jiashun Jin, Zheng Tracy Ke, and Bingcheng Sui. *The Gold Rush in AI4Math: Where Are We Now?* 2026. arXiv: 2608.24961.
- [29] Haocheng Ju and Bin Dong. *AI for Mathematics: Progress, Challenges, and Prospects*. 2026. arXiv: 2601.13209.
- [30] Michael Kearns and Aaron Roth. *How AI is changing the nature of mathematical research*. <https://www.amazon.science/blog/how-ai-is-changing-the-nature-of-mathematical-research>. Accessed: 2026-08-31.
- [31] Shane Kelly. *Some explicit counter-examples to Weibel’s conjecture*. 2026. arXiv: 2608.16066.
- [32] Graham Kendall. “When using artificial intelligence tools in scientific publications authors should include the prompts and the generated text as part of the submission”. In: *Journal of Academic Ethics* 23.3 (2025), pp. 639–647.
- [33] Mohamed Khalifa and Mona Albadawy. “Using artificial intelligence in academic writing and research: An essential productivity tool”. In: *Computer Methods and Programs in Biomedicine Update* 5 (2024), p. 100145.
- [34] Alistair Knott et al. “Generative AI models should include detection mechanisms as a condition for public release”. In: *Ethics and Information Technology* 25.4 (2023), p. 55.
- [35] Lee Lady. *How Does One Do Mathematical Research?* <https://math.hawaii.edu/~lee/how-to.html>. Accessed: 2026-08-31.
- [36] Chenyi Li et al. *Advancing Mathematical Research via Human-AI Interactive Theorem Proving*. 2025. arXiv: 2512.09443.
- [37] Yuanhang Liu et al. *AI Mathematician: Towards Fully Automated Frontier Mathematical Research*. 2025. arXiv: 2505.22451.

- [38] Theodore Jensen Mary Theofanos Yee-Yin Choong. *AI Use Taxonomy: A Human-Centered Approach*. National Institute of Standards and Technology, Gaithersburg, MD), NIST Trustworthy and Responsible AI (AI) NIST AI 200-1. 2024. DOI: <https://doi.org/10.6028/NIST.AI.200-1>.
- [39] Advisory Group on Mathematics and Artificial Intelligence of the Institute for Advanced Study. *Responsible Release of AI-Generated Mathematics*. <https://agmai.org/general-sep29/>. Accessed: 2026-10-01.
- [40] Arati Mohapatra and Danish Pruthi. *Expectations and Practices around AI Disclosure in CS Research*. 2026. arXiv: 2608.23271 [cs.CY].
- [41] Robert C Nickerson, Upkar Varshney, and Jan Muntermann. “A method for taxonomy development and its application in information systems”. In: *European Journal of Information Systems* 22.3 (2013), pp. 336–359.
- [42] OpenAI. *Finite time Blowup for Navier–Stokes*. <https://cdn.openai.com/pdf/32d9f210-8b73-45e0-91bc-82a30aef8a9a/navier-stokes.pdf>. Accessed: 2026-09-28.
- [43] OpenAI. *On the Navier–Stokes Millennium Prize Problem*. <https://openai.com/index/navier-stokes-solution/>. Accessed: 2026-09-12.
- [44] OpenAI. *Statement on: Ten advances in mathematics and theoretical computer science*. <https://openai.com/index/ten-advances-in-mathematics/>. Accessed: 2026-09-01.
- [45] OpenAI. *Ten Advances in Mathematics and Theoretical Computer Science*. <https://cdn.openai.com/pdf/ten-proofs-oai.pdf>. Accessed: 2026-08-31.
- [46] Markus Pantsar. “The need for ethical guidelines in mathematical research in the time of generative AI”. In: *AI and Ethics* 5.4 (2025), pp. 3657–3668.
- [47] Billy Perrigo. *Exclusive: OpenAI Used Kenyan Workers on Less Than \$2 Per Hour to Make ChatGPT Less Toxic*. <https://time.com/6247678/openai-chatgpt-kenya-workers/>. Accessed: 2026-10-03.
- [48] A-Z Research. *Mathematics Research Methods*. <https://a-z-topics.com/mathematics-research-methods>. Accessed: 2026-08-31.
- [49] David B. Resnik and Mohammad Hosseini. “Disclosing artificial intelligence use in scientific research and publication: When should disclosure be mandatory, optional, or unnecessary?” In: *Accountability in Research* 33.2 (2026), p. 2481949.
- [50] Vishwas Sathish et al. *LLeMpower: Understanding Disparities in the Control and Access of Large Language Models*. 2024. arXiv: 2404.09356.
- [51] Jaime A Teixeira da Silva. “Proposing authorship for artificial intelligence and large language models”. In: *European Science Editing* 50 (2024), e123908.
- [52] Norman Earl Steenrod et al. *How to write mathematics*. Providence, RI: American Mathematical Society (AMS), 1973.
- [53] Emma Strubell, Ananya Ganesh, and Andrew McCallum. “Energy and Policy Considerations for Deep Learning in NLP”. In: *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. Ed. by Anna Korhonen, David Traum, and Lluís Màrquez. Association for Computational Linguistics, 2019, pp. 3645–3650. DOI: 10.18653/v1/P19-1355.
- [54] Yana Suchikova and Natalia Tsybuliak. “The purity myth: Why stigmatizing GAI in academic writing is harmful”. In: *Science communication* 47.5 (2025), pp. 753–761.
- [55] Yana Suchikova et al. “GAIDeT (Generative AI Delegation Taxonomy): A taxonomy for humans to delegate tasks to generative artificial intelligence in scientific research and publishing”. In: *Accountability in Research* 33.3 (2026), p. 2544331.
- [56] Terence Tao. *Mathematics in the age of AI*. 2026. arXiv: 2608.16753.
- [57] William P. Thurston. “On proof and progress in mathematics”. In: *Bull. Am. Math. Soc., New Ser.* 30.2 (1994), pp. 161–177.

- [58] UNESCO. *UNESCO Recommendation on Open Science*. <https://www.unesco.org/en/open-science/about>. Accessed: 2026-09-01. 2023.
- [59] Debora Weber-Wulff et al. “Testing of detection tools for AI-generated text”. In: *International Journal for Educational Integrity* 19.1 (2023), p. 26.
- [60] Max Weinreich. *The crisis of AI-generated mathematics*. 2026. arXiv: 2608.02859.
- [61] Wiley. *Using AI tools in your research*. <https://www.wiley.com/en-us/publish/article/ai-guidelines/>. Accessed: 2026-09-01.
- [62] Yifan Yao et al. “A survey on large language model (LLM) security and privacy: The Good, The Bad, and The Ugly”. In: *High-Confidence Computing* 4.2 (2024), p. 100211.
- [63] Jiayi Yuan et al. “Understanding and mitigating numerical sources of nondeterminism in llm inference”. In: *Advances in Neural Information Processing Systems* 38 (2026), pp. 169819–169851.
- [64] Max Zimmer et al. *The Agentic Researcher: A Practical Guide to AI-Assisted Research in Mathematics and Machine Learning*. 2026. arXiv: 2603.15914.