Report of Summit on *PhD Math Education in the Age of AI*

September 17–18, 2026

Introduction

As we write (September 2026) the mathematics profession is at the beginning of a profound period of change, due to the rise of artificial intelligence. The effects on the mathematics community along many axes are overwhelming. In particular, our students face great uncertainty as they look forward to their education and careers. They have urgent questions for both the short and long term: “Should I be using AI in my research and if so, how?” “What will the mathematics profession of the future look like and will there be a place for me in it?”

The writers of this document are mathematicians who came together for an intensive 2-day meeting to focus on *PhD Math Education in the Age of AI*. We believe in the importance of human mathematicians creating, communicating, and understanding deep mathematics, and see our graduate programs as critical to the continued development of mathematics for humanity. AI tools offer exciting new possibilities that may reshape the discipline, and mathematicians will be needed to fully develop new advancements and directions.

While the highest aims of mathematics PhD programs remain stable, we suggest substantial changes in the implementation of our programs to achieve their goals. These include new emphasis on communication, detailed discussions with advisors about AI use, and changing assessments to not use written work as the primary proxy for measuring understanding. We recommend approaches that we think can help faculty and students navigate these uncertain times to best prepare our students for the future.

Our aim is not to be prescriptive. Our hope is that these recommendations will be a starting point for discussion in PhD-granting mathematics programs about how to reorient PhD programs toward the future of mathematics. The landscape is evolving rapidly; we anticipate that our views and this document will also evolve. To keep the focus on graduate education, this document does not address the broader issues that the rise of AI presents.

Top recommendations

1. Programs should add new training opportunities to teach skills that are becoming increasingly relevant. These may include: collaboration, formulating questions, interdisciplinary work, use of AI tools, and formalization.
2. Skill in communication to both expert and broad audiences should be part of the training in a PhD program and an explicit degree expectation.
3. While homework can provide a useful way for students to learn course material, assessment should be based on in-person written or oral assessments.
4. Programs should provide access to AI tools for all graduate students, with use never mandatory and methods subject to the standards of the advisor and department.
5. Advisors should hold regular detailed discussions with their students about expectations for AI use.

6. Three ethical rules: (i) maintain responsibility for the correctness and understanding of the mathematics that appears under your name, (ii) always disclose AI use, and (iii) obtain explicit agreement from all involved parties before entering material into an AI system.
7. A mathematics PhD program should require rigorous in-person assessment of research plans and progress by multiple faculty at least annually, including meaningful feedback for the student.
8. A thesis should be an original scholarly contribution. This should be understood broadly; for instance, the thesis need not record the first proof of a statement.
9. A PhD should not be awarded primarily on the basis of the text of the dissertation. The text should serve as the foundation for a rigorous thesis defense that requires complete mastery of the thesis content.

We recognize that implementation of our recommendations, and dealing with the many changes in the profession related to AI, will take a lot of time and energy for faculty and students. We suggest that faculty prioritize the changes that seem most useful and important for their programs.

The current landscape

The work of mathematicians and the concepts we have created, understood, and shared have defined the evolution of knowledge and science for thousands of years and continue to do so today. Mathematics trains the mind in logic, abstraction, and how to grapple with technical and conceptual complexity, skills that are valuable far beyond mathematics. While AI tools offer new possibilities that may reshape how mathematicians and others approach problem solving, they will not change the central role of mathematicians in guiding the direction and nature of mathematics, nor will they reduce the importance of training students.

The next few years in mathematics are a time of both uncertainty and possibility. On the one hand, the research process is changing, and it is hard to know exactly what role a graduate student or professor will play in a few years' time. On the other hand, there is a real possibility to contribute deep new ideas, paradigm-changing discoveries, and remarkable applications. AI is a powerful tool. We do not know how AI will develop in the future. If we can harness the power of AI tools to increase our mathematical understanding and creativity—in addition to using them to solve problems—then the future is very bright. AI would then help us to develop deep new mathematics, in which case it is important for humans to understand that mathematics, to reflect on what it means, and to drive its applications. Many more mathematicians will be needed to fully develop new mathematical horizons and to explore applications to science, technology, and beyond. What an exciting time to enter our field!

That said, we do not underestimate the pain caused by rapid change. Deep uncertainty will persist in these next years. We are at the beginning of a turbulent time. The mathematics community is coming together in many forums to navigate through this period. Our efforts will not be perfect; the current storm is too strong to hold that unrealistic expectation. There will be divergent opinions and recommendations as befits a diverse community. And each report, such as this one, will only address some of the issues we face.

Goals of a Mathematics PhD

A mathematics PhD program cultivates, in its students, the ability to formulate and execute investigations of mathematical structures, both abstract and concrete, and to communicate the methods and insights to experts and to the public.

Mathematics PhD programs prepare recipients for a broad range of careers in academia, industry, education, and government, among others. The AMS Ethical Guidelines express the standards for such programs: “a Ph.D. degree is a certification of both mathematical knowledge and independent achievement in mathematics. Institutions are responsible for ensuring both sufficient knowledge by the recipient of important branches of mathematics outside the scope of the thesis and the high level and originality of the Ph.D. dissertation work.” In the age of AI, we need to reevaluate what mathematical knowledge and independent achievement mean, how those goals can be attained, and how we assess them. We further recommend that PhD recipients demonstrate competence in communicating deep mathematical ideas to a wide range of audiences.

Recipients of PhD degrees navigate, absorb, and apply new knowledge at the frontier of mathematics. They demonstrate this capacity by achieving expertise in their chosen domain. They should master the tools used in the development of the subject and show initiative in making research decisions.

Independence and achievement can be demonstrated through creative and original contributions. They may include the discovery of new mathematical phenomena, inventing new methods for analyzing such phenomena, identifying meaningful new questions, finding connections and applications to other fields, resolving open problems, etc.

Students should demonstrate their ability to communicate their research, as well as the mathematical context around it, to experts and colleagues. Another valuable asset is an ability to communicate mathematics to broader audiences.

Learning

Graduate coursework and the transition to independent learning remain key components of a mathematics doctorate. This work should lead students to develop core mathematical knowledge and fluency in complex mathematical arguments, to recognize the modes of thinking in multiple fields, and to understand the interconnectedness of mathematics. Coursework and independent study lay the foundation for students to carry out cutting-edge research, and they instill in them the confidence and perseverance to think hard and develop independent judgment.

There are further goals that we believe should be included in graduate education. These goals are not entirely new, but many of us have not always consciously emphasized them. Ideally a doctoral education also teaches students to ask good questions, helps them work as part of a collaborative community, gives them skills to communicate mathematics both in and outside of the field, trains them to teach and mentor, and opens opportunities outside of academia.

Here are some thoughts about how to work towards these goals.

Communication: Within the mathematical community and outside, communication plays an increasingly important role for our profession. This can be taught in a dedicated class, but it can also be integrated into existing courses. Students, either individually or in groups, can give presentations to faculty and more senior graduate students, gaining comfort explaining what does not work, in addition to what does. Feedback from varied audiences plays a useful role in training.

Collaboration: Beginning with coursework, collaborative practices should be cultivated, creating communities in which students discuss mathematics, share ideas, ask questions, and learn from each other. Study groups and research group meetings (akin to lab meetings in other natural sciences) can integrate collaborative practices into daily mathematical life. Integrating senior graduate students and postdocs into mentoring networks provides students greater resources and trains the next generation of mentors.

Computer tools, including AI: Departments may wish to provide instruction for students in the use of various computer tools for mathematical research. Even students who do not plan to make active use of the tools themselves will benefit from improved literacy about new technology, such as understanding the basic architecture of computer proof assistants and how to determine the meaning of a mathematical definition or statement in a formalized text.

Beyond mathematics: Programs should facilitate the exploration of interdisciplinary work, both within academia and outside. The skills gained in navigating and distilling complex systems, reasoning through hypotheticals, and persisting through challenges are transferable and valued beyond academia.

Homework. Doing homework problems is still an essential way of learning mathematics. On the other hand, assigning a grade just based on written solutions no longer makes sense because it is easy to get solutions from AI tools. Homework could be a small portion of the grade or graded just for effort or not graded at all (but strongly encouraged). Courses should contain opportunities for students to develop their mastery beyond homework, including possibly projects, presentations, quizzes, and written or oral exams.

Asking questions. Faculty should spend more time working with students to develop the skill of asking good questions. We can teach students to ask themselves: What question is the material I am learning answering? Why is this definition introduced, and why does it take this form? Even in first-year classes, we might include raising questions as a consistent part of the students' work. For instance, each problem set could ask students to propose some questions about the topics the class is studying.

Research and AI

Principles

The goals of human understanding and insight in mathematics are eternal. The powerful and far-reaching applications of mathematics can be seen throughout history and will continue into the future.

We propose no hard and fast rules for PhD students on the use of AI in research, since the technology is rapidly changing. Instead, we encourage rapid experimentation to learn best how to use the rapidly evolving tools.

There should be access for all PhD students to use AI in their mathematical research, though we recognize that there can be financial obstacles to this ideal. A student's decision to use AI, as with any other tool, should be intentional, and made in light of their learning and development goals. Just as with other methodological aspects of PhD research, student use of AI in research should be discussed in detail with the advisor. Ongoing clear communication in both directions on past, current, and future AI use is essential.

Specific recommendations

Access to AI tools should be made available for all graduate students to the extent possible.

PhD students should critically consider whether and how AI tools could accelerate or damage their intellectual development. We give below a series of sample uses. There were significant differences of opinion on when to use, how to use, and whether such uses are productive or destructive—the answer can depend heavily on the particular case and goals of the user.

Department-wide forums to hear student questions about AI use, and department recommendations regarding best practices can be helpful to faculty and students. However, these should not substitute for advisor-student conversations. We encourage graduate students and their mentors to engage in detailed discussion about the costs and benefits of each potential use as it arises for the student. The uses below are listed in a rough descending order from those that may helpfully accelerate work to those that may very well prevent intellectual development.

- Making figures
- Literature search
- Cleaning up language
- Proofreading documents/text
- Rapid introduction to a completely new field or topic
- Coding and numerical exploration
- Self-refereeing: pressure-testing your own work
- Autoformalization of a proof
- Deep exploration of an area of expertise
- Brainstorming for general ideas or very specific ideas
- Proof generation (from a potential lemma, to a theorem, to an entire paper)

A critical principle that students and advisors should apply when deciding whether a student should use AI on a task is that AI use should accelerate understanding, not bypass understanding. Students should be advised to be extremely careful before offloading anything that will help them learn how to think. In particular, students should develop the capacity to think for long periods about hard problems, and students and advisors should make sure that AI use does not prevent that.

Ethical issues and best practices: not new and remain important

- Even if you are very “AI-forward” you should deliberately learn how to do things “by hand” without tools. (This is not new! This was true with older tools as well.)
- An author is fully responsible for understanding completely the mathematics they use in what they write.
- An author is fully responsible for citing the proper intellectual history of ideas and giving correct arguments in what they write.

Ethical issues and best practices: AI-specific

There are many ethical issues surrounding the use or non-use of AI, many of which are hotly debated. We record here some recommendations that have strong consensus in the community.

- Always acknowledge AI usage with specifics about the use. Dishonesty and lack of transparency violate the norms of scholarship.
- Asking AI to do literature search, check an argument, or referee a paper does not absolve the author of their responsibility for giving proper citations and correct arguments.
- All involved parties (such as coauthors, advisors, mentors, depending on the situation) need to be consulted and be in explicit agreement before putting material into an AI, even if you specify that the model does not train on your interactions.
- Users should always be aware that, even when there is an option to opt out of LLM interactions being used for training, it may be impossible to guarantee privacy when interacting with commercial systems. This is not the case when interacting with locally stored open-weight models, for mathematicians interested in exploring such options. You should use your own judgment when putting your own material into a model.

New classes and seminars

Experiment with new forms of classes or seminars: preprint seminars in which students are challenged to improve upon papers they are reading; or modern generalized versions of the Kan seminar, a course in which students learn to read mathematics and give talks effectively. Consider having seminars or working groups where sample workflows (for example, useful harnesses) are informally shared. These could be official departmental seminars, or just meetings of like-minded graduate students.

Assessment

We view this moment as an opportunity to adopt more rigorous standards for a mathematics PhD. This will entail thinking creatively about program requirements and their implementation, while reinforcing the important role that faculty have in guiding their students.

We expect future hiring norms in many occupations, including academia, will be based on multidimensional evaluation of the candidate rather than primarily their written products, and we view the recommendations here as beneficial both for assessing student progress in degree programs and for adapting to this new hiring landscape.

Recommendations

The dissertation has in the past frequently been used as a proxy for the kind of mathematical development of a student that we expect: acquiring mathematical knowledge and demonstrating independent achievement. However, with modern AI systems, dissertations are not, when used in isolation, a reliable tool for evaluating students, and PhDs should not be awarded primarily on the basis of the text of the dissertation. We thus recommend regular, multi-faceted evaluations of students in person, by multiple faculty, as the central component of assessment. These evaluations must be rigorous, not pro-forma, and should produce reports that might become part of the student’s dossier for future employment.

There are many ways such assessment can be done, many of which emulate models that are already used in other fields, and some of which we suggest below.

- For students in the stage where the focus is on developing a broad mathematical background, these assessments could take the form of written or oral exams or in-person qualifying exams, much as is commonly practiced.
- For students in the research phase of the program, there can be annual meetings with a thesis committee that rigorously evaluates progress and the student's plans for future work. As part of their scholarship and in preparation for these assessments, students may engage in learning and research seminars, as both an audience member and a speaker, in which the norm will be for vigorous questioning by the other seminar participants.
- A helpful model from other disciplines would be to introduce a formal research proposal, with a written prospectus and accompanying oral exam with the thesis committee. In the prospectus, the student would defend the choice of thesis topic, explaining the context and motivations for the research program as well as describe their plans to complete the thesis.
- One of these regular assessments should be a rigorous thesis defense, requiring complete mastery of the thesis work. Demonstrated mastery of one's own thesis should be an essential requirement of a mathematics PhD.

A thesis is expected to be an independent contribution to the mathematical literature commensurate with the scholarly norms of the discipline. These norms are already in flux, and we expect them to continue to change quickly. It is the responsibility of advisors and dissertation committees to communicate the scholarly norms to their students. We suggest that, while the thesis is expected to be a product of the student's ongoing work throughout their PhD, as demonstrated through regular assessment by a thesis committee, it need *not* be a requirement for the thesis to have recorded the first proof of a particular statement, particularly in an environment in which AI can rapidly and unexpectedly accelerate production.

Participants in the Summit

This report comes out of the work of summit participants in four working groups. The views expressed in this report represent broad agreement but are not necessarily unanimously held.

Mohammed Abouzaid (Stanford University)
Michael Brenner (Harvard University)
Ivan Corwin (Columbia University)
Sergey Fomin (University of Michigan)
Dan Freed (Harvard University)
Javier Gómez-Serrano (Brown University)
Larry Guth (Massachusetts Institute of Technology)
Joe Harris (Harvard University)
Mark Kisin (Harvard University)
Bryna Kra (Northwestern University)
Daniel Litt (University of Toronto)
Melanie Matchett Wood (Harvard University)
Curtis McMullen (Harvard University)
Assaf Naor (Princeton University)
Martin Olsson (University of California, Berkeley)
Tim Perutz (University of Texas at Austin)
Emily Riehl (Johns Hopkins University)
Raphaël Rouquier (University of California, Los Angeles)
Thomas Scanlon (University of California, Berkeley)
Gigliola Staffilani (Massachusetts Institute of Technology)
Ravi Vakil (Stanford University)
Rachel Ward (University of Texas at Austin)
Lauren Williams (Harvard University)
Chenyang Xu (Princeton University)