Advising guidelines*

Elden Elmanto

August 10, 2026

> There's a story
> And then characters will come to you
> Relating events as they choose to
> But all their words and actions come entirely from you
>
> Belle and Sebastian

Prelude

In 2011, I took my first algebraic topology class under Arunas Liulevičius who was most known for solving the odd primary Hopf invariant one problem (the 2-primary Hopf invariant one problem had been solved by Atiyah and Adams) in his 1960 PhD thesis. He had been a student of Saunders Mac Lane and, in a rather unconventional twist, taught us the acyclic models theorem which is a way to prove that certain homology theories (singular, simplicial) are equivalent. It's not typically taught in undergraduate curriculum these days, and probably subsumed by Eilenberg–Mac Lane's characterization of ordinary cohomology theories, but I felt that it was strikingly modern for its time. Anyway, I had asked him — with youthful spry — what he was "doing" these days. He was one to smile with his eyes and he gently said "planting daylilies." Daylilies are not really lilies. One of the interesting things about them is that their flowers last a day, but by sheer numbers, they look like they bloom for an extended period of time. I was quite confused by this reply, but I think his reply was prescient about what I regard as mathematical training in this modern age.

The substance

Since this is an advising document, I would be remiss if I did not go into my own scientific interests. For the foreseeable future, I am interested in exploring motivic phenomena in arithmetic and geometry. In concrete terms, I am interested in cohomological invariants such as algebraic K-theory, motivic cohomology and their derivatives (topological cyclic homology, hermitian K-theory) and non-abelian invariants (torsors, vector bundles, birational properties), and what they tell us about mathematical objects from geometry and arithmetic. Homotopy theory for me is a crucial tool, one that is baked into the very foundations of the subject. Many of the best theorems about motivic cohomology are truly theorems inspired by and/or proved with homotopy theory. Insofar as that is the case, I am interested in various aspects of homotopy theory (e.g. in its unstable, chromatic, motivic and equivariant incarnations).

The question of fit

At my core, I believe that I am a custodian of mathematics. Or, in Liulevičius' terms, a "gardener."[1] There is this not-so-secret garden where mathematical objects live. I am not quite sure myself where they came

*This is by no means a "contract" and is evolving and dynamic.

[1] I also like Freedman's "willow tree" imagery `https://www.youtube.com/watch?v=60X5M1FhmUc`.

from (and we can have lively debates about it) but I am inspired by them and I feel a sense of duty to separate the wheat from the chaff, keep the premises presentable, ensure that the plants are healthy, and nurture the beds of seeds present in this garden[2]. It is also crucial that the produce of this garden be shared with others, in ways that are responsible and respectful to the richness of this garden.

The advent of AI has made this conviction that much stronger. In more concrete terms, I believe that the job of a mathematician is to understand mathematical objects and their attendant phenomena. Problems and proofs have been useful in focusing our attention towards how these objects should reveal themselves and they have been the most tangible yardsticks, benchmarks of sort, for our understanding. But I believe all of this is changing, and changing quickly and mathematical practice deserves a reimagining. All this is to say that if 1) you care about the objects you are studying, 2) are open to doing mathematics beyond problem-solving and 3) are enthusiastic about sharing your mathematics with the world at large, then we might be a good fit.

Learning in the age of AI

There is now overwhelming evidence of the power of LLMs in solving interesting, difficult problems in mathematics. There is even some evidence of LLMs' ability to explain mathematics to humans, and in a way that is better than humans. I think, as many in our profession do, that dramatic and rapid changes will occur in mathematics within the next few years. I do not know what these changes will look like — for example I suspect there will be a massive overhaul of the incentive structure of the profession (e.g. how people get hired, what papers get accepted at top journals, the publishing system itself). Nonetheless, despite these unsettling changes, I believe that the process of learning and doing mathematics will remain valuable (both economically and as a human activity) for the foreseeable future. More importantly, I think nurturing the next generation is one of the best ways to ensure that the profession not only survive but also thrive. To paraphrase Freedman at the end of his talk `https://www.youtube.com/watch?v=60X5M1FhmUc`, what has made a mathematical idea or proof "foolproof" is not some kind of expert certification or a line-by-line check of formal statements, but the hundreds of students and practitioners who sought to use them to make their own results and their own mark on the subject. In other words, I believe that students are the "best chance we got."

Traditionally, a student will first learn about a particular area (or areas) of math before embarking on research. I think AI prompts us to think about what learning mathematics really is, and has been. The guideline I propose is that before you embark on research, you have to take "ownership" of the basics that it entails. I see no sense in prohibiting the use of AI tools in doing this. There are some good advice pieces out there on learning in the age of AI: notably Pavel Etingof's `https://math.mit.edu/~etingof/aiuse.pdf`. The following are some guidelines tailored to my own research taste.

1. First and foremost, engage *deeply* and *slowly* with the mathematics. Understand mathematics beyond the superficial layer, and make sure that it feels real and alive to you. What the latter means is a very personal thing, but there are "symptoms" — being able to make explicit calculations, produce clear and illuminating examples, transport one idea from one context to the other etc.

2. There is a lot of machinery to start in this area, a high "start-up cost". This makes *knowing concrete examples of paramount importance.* The AI is good at producing examples and you can use this to your benefit by prompting it. However, there is nothing better at searing a mathematical concept in your brain than trying to produce an example on your own, even if you fail. Which leads to the next point.

3. Do not be afraid of producing your own mental picture, caricature or examples. Sometimes it is easy to just ask the AI to teach you how to think about math — like a proof or how to unpack your definition. It is important to *have your own mathematical landscape*, even if it is jagged and may not be fully correct; sometimes really good mathematics comes out from the friction between the truth and a lie. This is also where talking to your peers, or experts can be helpful.

[2] Terry Tao uses the child rearing analogy in `https://mathstodon.xyz/@tao/117035751627601466`, which can also be helpful to some people.

4. *Embrace your own human limitations*[3]. That our processing capacity is not as powerful as the AI's is an advantage — we are forced to make mathematics digestible and, in doing so, gain understanding. Experience even shows that this can lead to *surprise*, a commodity in mathematics that will not change even with AI. One helpful way to do this is to make use of your peers. Think about how you can present an idea, or a proof in a way that someone with your background can digest, understand and appreciate.
5. As a piece of somewhat controversial advice, *indulge in digressions*. By this, I mean pursue mathematical detours which come up on your work, at least up to a point. One of the things that AI has challenged the community with is its own notions of productivity; in the age of "mechanized proofs" (or, as Terry Tao calls, of "proof abudnance") I see little point in trying to produce as many papers as possible given a certain time constraint. On the other hand, one of the most powerful engines of mathematical understanding and discovery has been serendipity — a mathematician who accidentally stumbles upon a landscape unknown to their community.
6. Last, but certainly not least, *take ownership and responsibility* of your mathematics. This cannot be emphasized enough. It includes being responsible for your own argument when writing a paper even if you engaged the aid of AI, being able to explain substantive, deep and difficult mathematical ideas to yourself and to other people, and decide your own mathematical journey. One of the nominal values of our field has been agency — the liberty to pursue unbridled curiosity for its own sake. I think we should hold on to this as hard as we can, but we cannot do it without taking agency of our own mathematical life.

Mathematical communication

More than ever, being able to communicate mathematics is of paramount importance for both the sustenance of the field and for your own academic career. If you have been to a colloquium talk at U of T or elsewhere in something not in your area, you will often be frustrated at how incomprehensible some talks are. And these are talks where the speaker is supposed to be communicating to a broad audience! The ability to both convey generalist and specialist knowledge is important. One of the things that people often miss is the historical context of the mathematical talk. Any talk worth its salt is contextualized within a human, mathematical tradition and that is one way to make talk "come alive" especially when speaking to a broad audience. I have personally enjoyed this part of the profession, but I think that we are actually obliged to be able to do this.

Relatedly, to borrow Cathy O'Neil's sober words [8], a proof is a "social construct." It has always been the responsibility of the mathematician to convince their audience and the community at large of the veracity of their claims. This makes the idea of a proof intrinsically social; contrary to how we idealize a proof as the certification of the truth, it is subject to the porous human elements of politics, games of power and status. Nonetheless, this is the superpower of mathematics. That despite all this, we have eventually "gotten there" one way or another and this has depended on a conscientious attempt to clarify ideas and train people to have the competence to produce, examine and have dialogues about mathematics. One of the great shames of modern mathematics brought this to the forefront [1], highlighting the cracks and schisms that can emerge if this practice is broken. Developments in AI and formalization systems are testing this convention to the limits; what should we do with a piece of mathematics that is true and yet completely incomprehensible and incommunicable? My personal take is that we should "do nothing" and that has little to do with mathematics as mathematicians have practiced it.

I should also highlight the new journal `https://www.mathematicaldiscourse.org/` which is a comforting indication about how the community is becoming more intentional about this point.

Concrete steps

Going forward, I aim to run my group as a hybrid between the "classic mathematics model" (one-on-one supervision, a "reading course" at the start of the adviser-advisee relationship, a problem given in the middle

[3] In [2], the authors speaks of the idea of "compression." Much of human activity involves producing "macros" which helps us cut down the cognitive load needed to process mathematics. In *op. cit.*, this idea was mathematically modeled.

of the PhD) and a “science lab model” (more collaborative, weekly meetings). In particular, I want to have regular meetings where someone from the group presents, informally but accurately, a piece of mathematics that they have learned and engaged with, perhaps for an extended period of two to three weeks. I will also try to coordinate projects between members of the group (while reserving a solo project for each group member). My vision is to run the group as a “small seminar” where we try to blur the line between exposition and “new results.” I will also demand more writing, especially of an informal nature, from members of the group. I am still trying to figure out what is best, especially in light of recent development, but I view participation in these group activities as the one mandatory thing that you should be willing to do.

The role of the adviser

At the very least:

1. To give you the necessary resources (like informational, financial ones) to help you succeed.
2. To convey intuition I have about parts of mathematics that I am familiar with.
3. To help you find an interesting and important research project that will form the backbone of your PhD thesis.
4. To introduce you to the community at large and advocate for you, to the best of my abilities and as is fair and appropriate.
5. To give concrete suggestions about what papers to read, who to talk to if you want to learn something, how to give research talks, what career moves to make, where to submit your papers, what conferences to go to and what teaching decisions you should try make.
6. To regularly (at least once a year) check on your progress and your well-being, to the extent that is appropriate.
7. To support your personal career goals — whether it be academic or in industry.
8. To manage the research group and create a safe, fun environment where members can thrive.

As a general rule, I will not take you on as an advisee if I do not believe that we are a match. In particular, this means that even if you are doing a reading course with me, there is no expectation that we will eventually work together unless I think that we do match.

The role of an advisee

I expect you to:

1. Hold yourself to a high standard in research, and in conduct within the mathematics community.
2. This includes giving proper attributions in papers and in talks and having a nuanced understanding of the intellectual history of your own line of investigation.
3. Keep in constant communication with me; I advocate for (reasonably!) frequent vacations and time-off in graduate school, but please let me know if you are taking one. You *should not* feel bad about asking me for time off.
4. Attend regularly (not miss without a real excuse) group meetings, learning seminars and the departmental seminars.
5. Keep abreast of development in the field.
6. Take “ownership" of mathematics, including and especially those that are generated with the aid of AI.

7. Attend conferences and “be visible" as a member of the community — within the department and the subject. One way to do this is to participate in learning seminars that are being offered by students and faculty throughout the department. Do not constrain yourself to just your specific research interests! Explore your curiosity.

What is “enough” for a PhD?

This will likely change in the next few years. Thomas Kahle’s essay is especially relevant [5]. Just like a paper, a PhD has an implicit (enough mastery of a piece of mathematics has been attained by the author) and explicit claim (result X is proved) and AI has “broken the connection between the two.” If you are at all reading this seriously, I trust that you are not interested in “one-shotting” a PhD thesis. Nonetheless, the PhD should align the implicit and explicit claims which means that I expect you to have mastery of your corner of mathematics en route to producing your written thesis. The latter should be written entirely in your own words, and I expect you to develop a sense of mathematical style in doing so. I also expect a PhD thesis to push, in one way or another, the community’s limits of understanding. The latter might not be a technically difficult task, but suggests a somewhat concrete and measurable goal in a PhD program.

Concluding thoughts

At this moment, I cannot in good faith suggest that “everything will be like before.” If you started at PhD program at U of T sometime between 2023 to 2026 you surely have a mental model of what mathematics is. I am unsure of how useful such a mental model is at the moment. On the one hand this is good because mathematics academia was far from perfect. On the other, there is, at present, a period of great turbulence and uneasy change. Many of the issues at stake are actually political [4], and I encourage you to voice your concerns as the most junior stakeholders in the community. I wholly empathize with a resistance movement [3, 6, 7] and such expression of passion is much needed in a field whose values typically align with reason and stoicism. What seems clear to me is that the impact of AI on mathematics can no longer be ignored and it is something that I wrestle with on a regular basis. This document is written, in part, to that end. On the other hand, it is also clear to me that if I truly believe in the importance of a mathematical community, then I have a duty to leave this garden in the best possible state for the next generation and beyond.

Three years ago (2023) I wrote some half-baked version of this document. This was a time when the AI was barely useful for calculating tips, an activity any mathematician would love automated. While things have changed a lot, I still feel happy with how I ended the document, so I will leave it unchanged. It is my answer to the question of “what is mathematics research?”:

In a sense, it is like storytelling. The biggest difference with storytelling is that one is constrained within the rules of logic, but a key part of mathematics research is that your results must be communicable and comprehensible to other people at large. Rather than a disparate, random collection of problems, research (ideally) advances the community’s understanding of certain mathematical objects which the community has deemed important. Sometimes the story presents itself in a clear manner and your job is to explicate what is going on; this does not mean that the story is “easy” as the act of explication can be technically challenging. Sometimes it is not clear how the story proceeds/moves forward and it demands an amount of creativity and trickery to get it going.

Acknowledgments

The original draft of this document was made sometime at the end of 2023. I decided to update and polish it because of recent developments in AI and I apologize to my current PhD students for not producing it earlier. I also want to thank them for enriching my mathematical life these past few years: Giacomo Bertizzolo, Dhilan Lahoti, Mahdi Raifei, Kai Shaikh, and Jason Sin. I would like to thank Aravind Asok, Ben Antieau, Almut Buchard, Jan Cao, Mathilde Gerbelli-Gauthier, Michael Groechenig, Mike Hopkins, Daniel Litt, Siddharth Mathur, Matthew Morrow, Erin Reynolds and Alex Youcis for helpful comments on versions of this document and useful conversations. I hope it is clear that the impact Professor Liulevičius had on me and this document is indelible.

References

[1] Preface to the special issue [Inter-universal Teichmüller theory, I–IV]. *Publ. Res. Inst. Math. Sci. 57*, 1-2 (2021), 1.

[2] Aksenov, V., Bodnia, E., Freedman, M. H., and Mulligan, M. Compression is all you need: Modeling mathematics, 2026.

[3] Chu, T. Mathematicians need to act. `https://tasmin.substack.com/p/mathematicians-need-to-act`, 2026. Blog post, accessed August 4, 2026.

[4] Harris, M. Mathematics is not rational really. `https://siliconreckoner.substack.com/p/mathematics-is-not-rational-really`, 2026. Blog post, accessed August 4, 2026.

[5] Kahle, T. Just "42" tells us nothing. `https://thomas-kahle.de/blog/2026/just-42-tells-us-nothing/`, 2026. Blog post, accessed August 8 2026.

[6] Lazic, V. Where i stand on ai. `https://www.uni-saarland.de/lehrstuhl/lazic/prof-dr-vladimir-lazic/vladimir-lazic-where-i-stand-on-ai.html`, 2026. Blog post, accessed August 5, 2026.

[7] nihilunbounded. `https://x.com/nihilunbounded`.

[8] O'Neil, C. What is a proof? `https://mathbabe.org/2012/08/06/what-is-a-proof/`, 2012.