Career feature

https://doi.org/10.1038/s41587-025-02584-1

An open letter to graduate students and other procrastinators: it's time to write

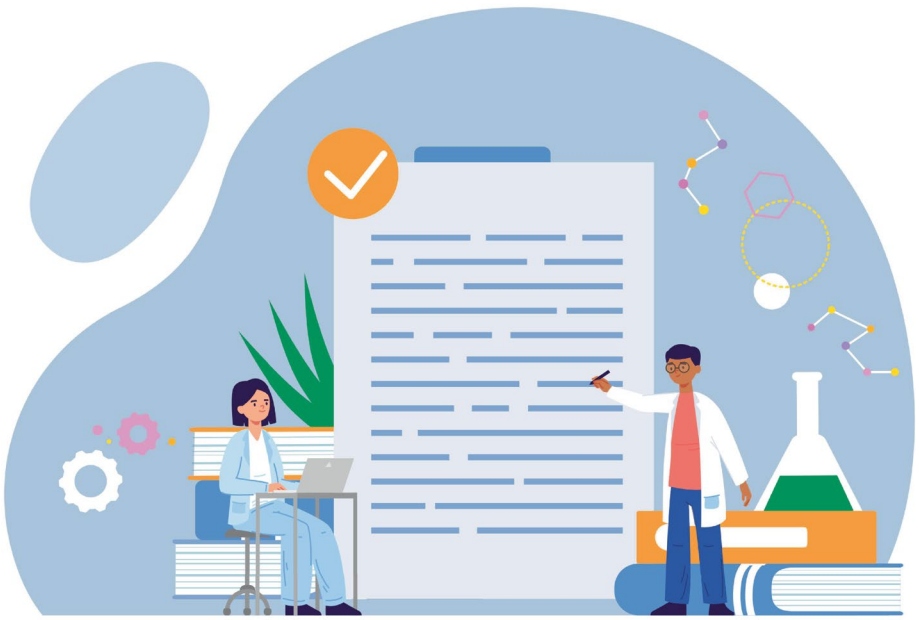

Let me begin with a hard truth. As scientists, writing is our chief activity. It can be argued that it is the only thing we do that matters. It is the singular activity that you as a young scientist partake in that will raise your salary. You may have felt protected from this truth for a time, but this is unsustainable for the long term. Allow me to explain. This letter is for the graduate student who has been asked by peers to begin writeup of a recent series of experiments. This person doesn't know where to begin and has received or absorbed a lot of free and unsolicited advice — including what order to write in, how to arrange their office and schedule to maximize productivity, and perhaps some tips on self-care. As the joke goes, household chores have never held so much allure. Perhaps this person is you. You love science. You love the activity, the technology, the fact that you can see DNA with your eyes on an agarose gel, separated into pieces of different lengths. But you may approach the act of writing science with reluctance, as a nuisance or distraction from the 'real', hands-on, day-to-day work.

Everyone understands the structure of a scientific paper. But when presented with the onerous task of writing a scientific manuscript on your own work, it may seem at times an insurmountable challenge to get started. If you are doing your job properly, you will be writing about something that has never been done before — and in a few cases has never even been thought of before. Your audience needs to understand two things: what is the nature of the question or problem you are attempting to address, and what do you — specifically you — think about it? There are a host of other problems and issues that arise around these two things, but understanding them is key to getting started with writing.

Here we have arrived at the most essential point that should motivate you to read the rest of this letter. If you are in the sciences, at the PhD level, writing is the only thing you do that matters. The most important trait for you to have to have lasting impact on humanity, and indeed to sustain your career, is your ability to have ideas, to gather evidence in support of them, and to publish them. Your job is to have ideas and spread them so that they impact other people's work. Writing is your job. Nothing else. Therefore, everything else that you do is subservient to the activity of writing. Let this revolution take place in your mind.

Writer's block is a myth

There are lots of resources and varying opinions on how to write a technical paper. Whole books have been written on how not to write a paper. Most people I've discussed the topic with are self-assured that they have the correct formula. I'm not here to contradict them or make my own prescriptions. Anyone can edit a paper. That particular exercise — molding existing text into shape — is not your problem. Your problem is getting words on paper in the first place. "I'm a good writer. Why can't I get this thing written? What is wrong with me?" Inevitably, the conclusion will be drawn: "I'm just not up to the task." I only wish that someone had recognized my struggle as a young person and told me the following: There is no such thing as writer's block. There is a procedure to getting started on a new manuscript, and it only takes a few minutes.

Write a title first

There are many valid choices for the order in which to write a paper. It is my personal experience that every paper is different, and deserves to be met on its own terms. This may not be true for large labs with an assembly line of postdocs, labs built on strength of technology and infrastructure. Nonetheless, writing a title down is a critical first step in the thought process that leads to a finished manuscript because it creates a frame of reference to build around.

That frame of reference is the key to the whole thing. But it's important to understand

that when I say to begin with the title, I do not suggest the title should become fixed in stone, lest the entire manuscript collapse like a house of cards under the weight of late-breaking experiments. Quite the opposite: scientific writing is an iterative process in which ideas become more succinct and more cogent each time you pass over them.

So, why begin with the title? I am asking you to think deeply about your project and pull out the golden essence of it, extract it as a single short(-ish) sentence. Think of the singular, major claim at the heart of your work, at the highest level of meaning. What is your intended audience's level of expertise? What is it that they should know, having been exposed to a procession of newly birthed experimental evidence, that they and all others did not, or could not, have known before? Proof is not required at this point; only what you think, boiled down to one or two short phrases, based on what you know.

The hardest part of this task for scientists today is to avoid references to technology. The trend in top-tier journals is toward repeating old conclusions with new technologies. To get into these top journals, many feel that only the latest tech dressed up in the right buzzwords will work. We have no control over what's trendy, but if we want to write papers with lasting impact, we should not reference these technologies in our title even if we have the latest, most expensive spatial omics detectors. Try this as an exercise: pick up an issue of a top-tier journal and try to find titles without reference to the technology used to collect the data. Where such references are made, can you reimagine the paper title without them? Does it have the same perceived impact?

My intention isn't to cast aspersions on papers with technical details in their titles. There are plenty of circumstances calling for inclusion of such terms. The point is for you to practice stating your reasons for writing your paper or thesis independent of those details. While acknowledging that the final title can look very different, the goal is to get you started writing on what truly matters the most to your readers — what is most impactful for them, and for you.

If you can write such a title for yourself, it will serve as a literary guidepost to which you can repeatedly return for direction on whether any given sentence or word is useful, necessary or appropriate. I cannot personally imagine writing without such a guidepost. When you write, you should write first for yourself. You are your own most important audience member. And to satisfy your most discerning reader, yourself, you must pass the highest bar. This is writing for impact.

How to eat a whale

There seem to be many potential origins of this metaphor. To me, it was introduced as a Zen koan. As the saying goes, the only way to eat a whale is to start from the tail. This answer is akin to saying, "You have to start somewhere." That was never my answer, however. Rather, I preferred, "one forkful at a time." All this really means is that there is no other way to tackle the largest tasks in life other than in the increment of work available to you, which may indeed be discouragingly small. But there is hope.

I chose this metaphor specifically to introduce outlines, which you've no doubt been asked to write before. I don't find them intuitive myself, and have sometimes resisted this bit of received wisdom, to be sure. But as I taught myself to write, I found myself nearly always incorporating them in some form. I used an outline to write this letter. The key thing that I would like you to understand about outlines is that they boost your productivity.

Writing in the academy means combining complex ideas, building and staging thought processes, stringing them together like great armies in columns and regiments. Our brains have remarkable capacity to do this, to set these armies against each other in pitched battles, to make these connections over time, to process them while we sleep, to catch us off guard during our morning shower. There is no greater aid to taking on this monumental task, that of capturing the fruits of this nocturnal theater, than an outline. It affords us the opportunity to write little titles to serve as new guideposts for the little mini-essays that are to follow.

Where I differ from what you've likely been told before is that your outline is for you. It is an aid or blueprint to help you eat the whale. As such, not being an exercise or a chore, we should not regard it as something that must be completed, or even take any particular form. The most expedient means to begin is to imagine this less as an exercise in writing, and more as putting together a task list. Anyone can write a simple paragraph with a seed topic plus a few sentences. Our object is to make that a task, and string together enough of these smaller tasks (as paragraphs or subsections) to get something done. You may end up writing down all the tasks you need in one shot to compose your entire paper. You may choose instead to work on a single section at a time. It depends on the project and how sure you are of the material at the outset.

You're probably ready to start writing something, no matter how tangential it seems at the outset. The important thing is that you must break that project into smaller, specific, realistic, attainable tasks. If you're working on something really big, like a thesis, I highly recommend giving yourself deadlines. Now that you have specific, attainable tasks to complete, it should be reasonable for you to set some expectations for yourself. Only with some hard dates set will you keep the boss (you) happy. Keep them happy by rising to meet your own expectations, and by knowing what those expectations are in the first place.

There's no such thing as thinking

Writing is such a difficult thing to get started on that people assume there's some amount of thinking they can invest, and when that mythical grain silo of preliminary thinking has been filled, it will be time to plant the season's crops. The second hard truth I have for you is this: thinking about your project in this way is actually a form of procrastination — quite possibly the most damaging kind. The only kind of thinking that matters in science is structured thinking. The only way to give structure and substance to your thoughts is to write them down. Writing. Is. Thinking.

How many times have you tried to write something down and realized as you were writing that it didn't make sense, at least not precisely, the way it first appeared to you in thought? There is something magical that happens as words travel from your frontal cortex through the motor centers and get translated into complete sentences and paragraphs. They get seen by the other parts of your brain. They can't hide away from the scrutiny of the grammar police in the protected little hovels of your hippocampus. No longer can they fester in your language center, gossiping with each other about how sublime they are. Even whole sentences wilt in the sunlight when considered in a procession of other sentences, or when asked to justify themselves before the withering gaze of the section title.

Manuscripts are not merely important for career advancement. They keep your projects on track. They keep the big picture in focus. They help you understand on a daily basis whether you've done the right controls, and whether you're asking the right questions or accelerating off on a tangent. We're so focused on bagging a paper in a top-tier journal that we've forgotten how to do science, or what science even is any more. Getting your work into a top-tier journal should be secondary to doing top-tier science, because you have

little to no control over the former. So, just do top-tier science! Doing top-tier science requires deep thinking in preparation for the experiments you're doing. Deep thinking means writing.

Just to complete this thought, I would be remiss if I didn't mention that the more you write, the more it will help your performance as a scientist. I keep two notebooks: one a regular lab notebook, the other a sort of private diary on my own scientific thinking. The first type of lab notebook should be self-explanatory, and the vast majority of people reading this have or are required to keep one by their principal investigators. As a computational lab, we keep hybrid lab notebooks with mixed commentary and code. The second type of notebook I only started to keep recently. I don't write in it every day, but whenever I have a new thought or have been inspired to try something different from a public talk by another scientist, I record my ideas. I try to give my thoughts structure as I put them down, as if I were really drafting a document. My memory of these specific thoughts tends to improve with reuse, and what's more, I find that if I use the text to seed the beginning of a new grant or paper draft, the quality of that initial draft is much higher. By the time I start this process, I've already 'thought' about what needs to be said, even at this early stage.

Writing isn't fancy

There's another incarnation of writer's block that says, "My writing doesn't sound brilliant or scholarly." Further, "I can't show anyone my work because it's not polished enough." Don't you believe it. I have been told that my writing isn't good enough. Perhaps this has happened to you, or it has been implied. Get as far away from these people as possible. Anyone can write; it just takes practice. Regardless of your situation, you can still write for yourself, and your voice is great.

As a matter of fact, you contain within your head what people in the information sciences refer to as a large language model (LLM). The sum of your experiences and exposures, including through reading — more about that later — has produced a generative model of word succession probabilities. This model includes information about how phrases are put together, how grammar works (or doesn't), and what makes an effective or convincing argument. Needless to say, you know an effective brief when you see one. And difficult though it may seem at first, you can always edit your own copy to make your arguments more convincing.

As an aside about LLMs, this is why I am opposed to the use of LLMs in my own research group, specifically for writing. I don't mean to be rigid about this, but it is my opinion that if you need an LLM to provide the text for you, it may be because you are unwilling to do the thinking. You may not be able to hide this fact from your reader. There are lots of use cases — for example, in the clinic, where hours of productivity are saved by having AI listen to an interview with a patient and compose notes. But it will always come at some cost. In fact, there is some concern that humans in a few years could lose the ability to think for themselves[1].

There is another barrier to writing that I've encountered, and that is the immigrant or visiting scientist who is not a native English speaker and who may be embarrassed by their less-than-authentic-sounding prose. To these people especially, I want to offer encouragement. As I've said previously, we all have an LLM represented in our language centers. As a non-native speaker, yours is slightly different, and perhaps a little smaller. But every native speaker also has a unique LLM in their own head, representing English, and therefore you could hardly replicate a native speaker's model even if you were native speaker yourself. In fact, most non-native speakers, being forced to rely on more limited vocabularies, have a tendency to get directly to the point, and that's the biggest advantage you can have in writing. If you are a native English speaker reading this, consider acquiring a second language model to improve your English writing skills.

To write for impact means to keep it as simple as possible. Long words and adjectives are not called for in scientific writing. Mark Twain is supposed to have said, "Writing is easy. All you have to do is cross out the wrong words." A rule of thumb is that your drafts should get shorter with each iteration. Don't write for the Nobel laureates in your field. Write for the students — and your parents, because they want to read your writing more badly than anybody. Don't leave them behind.

Why reading is hard

Just to extend this metaphor of the LLM in our heads, I say that reading is hard, because in the act of reading we encounter prose that violates the rules of our own LLM. In the act of violation, it expands the definitions and the probabilistic universe of how words follow one another and how they can be put together in convincing ways. You cannot help but learn new writing skills while reading, and learning is an extremely taxing activity, at least until it becomes a habit. Building your attention span through reading is essential. Fortunately, you needn't read only scientific papers, but you should have a steady diet of them incorporated into your work routine.

Another reason why reading is hard, especially when it comes to technical papers, is that the vast majority of them are poorly written. Don't read a paper if it's poorly written. Is there some essential experiment that you need to understand that is highly relevant to your work? Then do what you need to extract that information and move on.

When you do find a good paper, relish it. Learn from it. Cite it. If you are tasked with reviewing such a paper, reward the authors with a recommendation to publish. Share good papers with your friends. Present them in journal clubs. Contact the authors and ask them your questions. Any good scientist should be flattered by your curiosity.

In our lab, we maintain a virtual journal club. This is a database of papers we've read from title to disclaimer with a star rating system. I also keep a journal of short, one-paragraph summaries I write for the papers I've read. This is enormously helpful in recalling later what the paper was about because being able to see my own words referencing the paper is sufficient to recall a vault of other thoughts I may have had about the paper, like the keys to a bank deposit box.

Speaking is another way of writing

I want to make a final point about scientific communication more generally. If you want to succeed as a scientist, you will be called upon to write grants, prepare figures and — very importantly — to give public talks. The latter include journal clubs, lab meetings, department seminars and formal presentations at national and regional meetings. Each of these has its own peculiar challenges, and each represents an opportunity to apply the principles discussed above.

Whole books have been written about effective visual communication, none more famous than *The Visual Display of Quantitative Information*[2]. I suggest you read them. There are many parallels to the all too brief treatment I've given here. Writing about your ideas will undoubtedly help you to choose the most effective visualizations in support of your thesis.

Sitting through a scientific talk is generally challenging, for the same reason that reading most papers is. They're often prepared without much due given to the basic questions I posed at the beginning of this essay: what is it

you mean to say, and to whom are you saying it? When you fail to communicate effectively, in any medium, you've wasted your time and that of others. Worse still, you've missed out on a precious opportunity to promote your work and that of your colleagues and mentors. The bottom line is: writing is the most essential activity of a scientist because without writing there is no thinking and no real opportunity for exchange in the marketplace of ideas.

Dennis J. Hazelett
Department of Computational Biomedicine, Cedars-Sinai Medical Center, Los Angeles, CA, USA.
e-mail: dennis.hazelett@cshs.org

Published online: 17 March 2025

References

1. Ahmad, S. F. et al. *Humanit. Soc. Sci. Commun.* **10**, 311 (2023).
2. Tufte, E. R. *The Visual Display of Quantitative Information* 2nd edn (Graphics Press, 2001).

Competing interests

The author declares no competing interests.