Find the stable and pull out the bolt

Natalie B. Hogg*

24th February 2026

A little over two years ago, I wrote what might charitably be called a think-piece on the use of ChatGPT in academia, or more specifically, my observations on the use of ChatGPT by a handful of my colleagues in the small niche of French cosmology that I inhabited at the time. On reflection, what I was reacting to was less the use of large language models (LLMs) in research, and more how I felt my colleagues were undervaluing or mistrusting their own creative output. I was writing from the rhetorically precarious position of someone who believes themselves to be creative, values their own creative abilities, and had, at that point, never used or interacted with an LLM in any way, shape or form.

I never claimed to present an unbiased perspective (as I felt was my right, writing on a website of which the domain is my own name), but with the benefit of two years' hindsight, what comes across to me most strongly upon re-reading what I wrote is the emotion behind my words rather than the argument. With this current article, I aim to summarise my opinions about LLM use in a slightly more measured way, and discuss how and why they have changed since that original post.

Beginning at the beginning

First things first: why does my opinion on this matter at all? As previously mentioned, until October 2025 I could proudly claim that my fingers were unsullied by LLM prompting (and indeed, at the time, I felt this was something to be proud of, or even smug about). Then, I got a new job, and discovered that my principles were, disappointingly, much more flexible than I thought when it came to LLM use. I hope, therefore, that my opinions are interesting, as someone who has seen, and indeed lived, on both sides of the fence.

Furthermore, much – if not all – of the public discourse on LLM use in cosmology academia is being carried out by senior scholars, e.g. [1, 2, 3, 4], people whose careers will be largely unaffected, regardless of the way the LLM debate concludes. It is a very unfair characterisation, but I cannot deny the phrase "old man yells at cloud" hasn't drifted across my consciousness more than once since being pulled (initially reluctantly) into the conversation. I think it's important that those with more to lose – more skin in the

*Institute of Astronomy & Kavli Institute for Cosmology, University of Cambridge, UK. nbh25@cam.ac.uk.

game – have their say.[1] We should be paying extremely careful attention to what our undergraduates and PhD students think about this topic.

Lastly, I will focus only on LLM use in cosmology. In other walks of life – medicine, justice, law enforcement, warfare – the use of LLMs *owned by private companies* clearly carries extremely serious ethical concerns. In cosmology, we are mainly worried about how LLMs will change the way we practice our science, and the way we will teach the next generation. Whilst these questions are more existential than ethical, answering them, or at least collectively deciding how we will try to answer them, is nevertheless crucial if we want to ensure that our science not only sustains itself but continues to get better over time.

The rest of this article is structured under a few loose headings. I will give some more context behind my LLM disavowal and rapid change of heart, and then discuss the ways in which LLMs are forcing us, or should force us, to change the way we value and evaluate our work, creativity and peers.

Before

I can't remember the moment when I first heard about ChatGPT, nor when I first saw someone using it. I do remember, however, that I was vocally opposed to its usage – and with no qualifying clauses. I didn't think humanity at large should be making use of LLM technology. In retrospect, this is an uncompromisingly Luddite attitude, opposed to the character of the scientist or the principled Bayesian thinker whose duty it must be to dispassionately assess new methods and results, and consider how they will change or influence their own work. So why did I hold this viewpoint?

I felt strongly about three distinct things:[2]

- LLMs are bad for the environment.
- LLMs are plagiarism machines.
- LLMs erode human creativity.

The environment. By any metric, LLMs are phenomenally wasteful and energy-hungry [5]. Unfortunately, I do not expect this argument to sway any of my colleagues, most of whom regularly consume industrially-farmed red meat and take international flights for business and pleasure, seemingly without a second thought. Is it the case that the climate crisis is too intangible—too huge—too scary—to people, that they cannot allow themselves to worry about it? I can't think of another reason why environmental arguments against LLM use seem to gain so little traction. A colleague of mine espoused the viewpoint

[1] Lately I have come to despise the terms "early career researcher" and the like, which is the group of people I'm really alluding to here. If we count my career as starting from the day I was first paid to do research, it is not much shy of its 10th birthday. Then again, I was actually asked for ID in a pub only a week ago, despite five and change years of postdoc bestowing me with plentiful grey hairs. Perhaps I should not be too eager to throw off the mantle of youth – I may miss it when it's finally gone.

[2] A fourth issue is that the cost of the top-tier models is prohibitive to all but the wealthiest academics, and this does not cleave to the generally accepted goal of widening access to academia. But this is an objection to the companies that are selling us LLM access, not to LLMs themselves.

that we are within our rights as human beings to consume all of the planet's resources as quickly as possible; we will survive in the end, he thinks, because we have always been ingenious enough to do so in the past. I think this is hugely underestimating the magnitude of the climate crisis, and takes a very uncaring attitude towards those who are bearing the brunt of the profligacy of developed nations. One might also wish to ponder the thought that, after 180 million years of global dominance, the dinosaurs thought themselves pretty well set up to live forever too.

Plagiarism. Academics are trained from infancy to know that plagiarism is bad. This is not just because copying a work and passing it off as one's own is lazy, and probably damaging to one's ability to reason and learn; it is bad because it undermines what academia *is*. Isaac Newton put it thus:[3] "If I have seen further, it is by standing on the shoulders of giants". We are the sum of the knowledge that has been amassed before us and the epsilon of our own contribution. If we do not apply the proper attribution to the source of stated facts, methods and arguments, we jeopardise the preservation and transmission of knowledge in its true, correct, undistorted form. In other words, we are what we cite.

The use of an LLM for academic purposes that has been trained on a corpus of text which encompasses the majority of humanity's written word would seem to automatically violate this principle. LLMs can of course provide citations for statements they make, but this is not always foolproof. We call it *hallucination* when an LLM creates false artifacts of the literature; an interesting word choice which implies an underlying sickness instead of innocent mistake or deliberate deception. But an LLM is not merely a transmitter of knowledge – by its very design, it is a transformer, and synthesiser. What is the true provenance of the statements that it makes in response to prompts? By using an LLM, one introduces an opaque and obscuring layer between oneself and the canon, leaving oneself at the mercy of the weights and tokens which have shuffled and shaped it into an often sycophantic and obsequious version of the truth.

Furthermore, and of rather less actual concern to my personal morals,[4] there is the question of copyright. The LLMs currently being adopted at pace in the technical world – Claude, Codex, Gemini – were trained on copyrighted material. Hacked versions of the models can be coaxed to repeat an entire copyrighted book verbatim when prompted with a few words from the first sentence [7]. Copyright exists to assert legal ownership of a work, usually creative, and to ensure the owner can profit from that work before anyone else. It is certain that the companies training LLMs on text scraped from the internet, where copyrighted works may be found relatively easily, did not seek permission to do so from the owners of those works.

And what of the creative output of people who do not have formal copyright of their work? Freelance authors sharing stories online, or people posting essays on sites like Medium or Substack. They are presenting their work for public consumption and criticism, but that is not the same as blanket permission for the transformation of that work by an LLM without proper attribution. A young poet using Facebook[5] to tentatively but proudly share a new sonnet should not have to worry that an LLM will be able to

[3] Ironically, whilst checking the attribution for this, one of my favourite quotes, I discovered that the sentiment predates Newton by some centuries [6].

[4] I am a firm believer in free open-source everything, and have no scruples when it comes to pirating films, for instance.

[5] Again, am I really still an early career researcher if I'm making painfully out-of-date references like this? On which platform are the youth sharing angsty song lyrics nowadays?

reproduce the couplets word for word in another person's computer terminal half a world away, and with no way for that end user to discover the original provenance of those words. This, in my opinion, is far more damaging to society and its individuals than the unabridged texts of bestselling books by millionaire authors making their way into LLM training data.

Creativity. Thus we arrive at the question of creativity. For a cosmology academic, I felt that the main creative skill – beyond devising new research ideas – is to write text which contributes to the literature through its reporting of new findings. As argued above and elsewhere [2], cosmology *is* the literature. However, I believe that many people see this aspect of the job as a chore rather than an opportunity for pleasure.[6] I think this is generally symptomatic of people's lack of faith in their own ability to write. This may be due to simple lack of fluency, or an underexposure to narrative, rather than academic, literature.

It may also be due to the fear of failure, which can manifest as procrastination, or perfectionism. The necessary condition for pursuit of a career in research is to sacrifice oneself to one's passion. To willingly subject oneself to a PhD, or a postdoc, is to have a burning desire to understand *something*, at the expense of financial stability, personal relationships, and often physical and mental health. We write papers and give talks, tying our name and face to our research output. It is intensely difficult to divorce one's ego and emotion from the success or failure of a particular research line.

A way to get around this is to farm one's creative output to a machine. Most of us *really* care about what we do – but unbridled passion is not a fashionable stance. To be seen to care is cringe. We can therefore kill two birds with one stone: remove the ego from the equation, and come across as just a little bit cooler. With an LLM taking care of your writing, you can be dispassionate and urbane, and shield yourself from the hammer blow that is the latest grant rejection, because you didn't write it yourself anyway. I worried that this would start a vicious circle: creation of text now feels easy, therefore creation is easy, therefore creation is not valuable.

These were my main objections to LLM use in academia, and on the face of it, why I publicly and vocally opposed their adoption. On top of these three worries, I also held a private fear. I knew that LLMs were not, and are not, one hundred percent reliable. To use them for science means the scientist is mandated to check and validate everything the LLM produces. But I was afraid that, if I started using LLMs, I would not do this. My approach to work has always been haphazard – I am better at trial and error than sustained, careful reasoning. I make silly mistakes. My patience wears easily. I'm certainly not alone in this – have I not just described human nature? I feared that if I let myself start using LLMs, I could not trust myself to use them carefully. Was I projecting this fear onto others, in my edict that no one should be using LLMs?

After

Despite holding the view that LLMs should not be used, in October 2025, I began to use them in my work. I had started a new job, and joined a group in which LLM use was widespread and encouraged, with money set aside to pay for access to the best models. My personal turning point was during my first in-person meeting with my new boss, Will Handley. He asked me what programming tasks I had on my

[6] As you may be able to divine from the length of this article, I enjoy putting pen to paper.

to-do list. I named one, and within five minutes, he had used Claude Code via its command-line interface to complete the task – a task that I would have probably spent hours or days implementing and testing had I done it myself.

My only prior experience with LLMs having been watching people find bugs in snippets of code by copying and pasting them into ChatGPT's GUI, it was an eye-opening, and almost frightening experience. I can't pretend that I immediately foresaw all positive and negative consequences of these tools in that instant – but the more I have used LLMs (to the point of now training my own), the more I have ruminated on the points discussed above. Why were those arguments insufficient to stop me using LLMs?

The environment. Above, I questioned why the environmental damage caused by LLMs does not give academic users pause for thought. In weighing my own use of LLMs against this point, I can only make two weak arguments. The first is to trot out the old Marxist line that there is no ethical consumption under capitalism. Everything we consume costs energy – heating my home in winter damages the environment, as does buying out-of-season fruit and vegetables, filling my car up with petrol, or taking a shower on a daily basis. Should we therefore not do these things? In our industrialised society, there is almost no way to function *without* doing them. Does prompting an LLM for the purposes of academic research fall under the same category?

My second, equally feeble justification for my LLM use is that I consider my integrated environmental impact to be relatively small for a citizen of a Western nation. I have thirty-one years of vegetarianism under my belt, and I can count on my own digits the number of times I've taken a commercial flight without having to take my socks off. Do I therefore have the right to increase my wastefulness, by using LLMs? Does each individual have a share of energy apportioned at birth, and am I still under budget?

Plagiarism. We can imagine an LLM trained only on free, open-source code; software that has been implicitly or explicitly licenced for re-use and transformation. They would still be excellent code-generation machines, and would no longer run the risk of copyright infringement. Could we, in that situation, accuse the LLM of being a plagiarist? I think not. But this is not the world we live in. Driven in equal measure by scientific curiosity and corporate greed, the Pandora's box has been opened, and it is not only software which has been fed to the LLMs, but prose and poetry and academic research; the sum total and fundamental essence of those fields of research, as we have argued. LLMs, by their nature, transform their training set into something new, the provenance of which can never be fully known, as it is non-deterministic.

The only rebuttal that I can think of to this argument against the use of LLMs is that we *must* check what they tell us, and ensure that each individual statement can be traced to a known source. The only way to *exactly* verify the output of an LLM would be to read the texts they have synthesised, and repeat that synthesis for yourself. This is, of course, an impossible task for any individual person to complete. People like to talk about LLMs as artificial intelligence, and consider their thinking patterns to mimic that of a human being's. This is wrong. In truth, each human being is an individual weight in the LLM that is the literature. If you use an LLM for academic research, you must therefore make a choice: accept that you are plagiarising, or undertake to verify each and every one of its outputs. If the latter, why use one at all?

Creativity. Two years have passed since my initial post, which may be enough time to begin to judge

whether the creative erosion I feared has come to pass. In other people, it's not clear. I have, in the interim, read many papers in passing on the arXiv, and plenty more in detail in my capacity as a reviewer for journals and within collaborations. In general, I find it very easy to tell when a text has been generated using an LLM. Some people have scoffed at this ability, but why shouldn't it be the case? If one can distinguish a passage of Austen from one of Hemingway,[7] it is not surprising that one can tell machine-generated text from human. There is no passion in an LLM's words, and no idiomaticity. Their statements are usually vague, general and overblown – they were probably prevented from becoming too opinionated during their training, and for obvious reasons.

For myself, the argument against using LLMs for the fear of what they will do to my creativity is the strongest. Whilst I have no qualms using them to generate code, I will not allow them to put words in my mouth – words that are inevitably mealy, insubstantial and alien. Why should I expect other people to read something that I did not take time to write? Why should I allow myself to be satisfied with piping mass-produced porridge down my neck when the world's greatest linguistic smorgasbord – the English language – is ready and waiting for me to devour it?

Above, I posed this exact question. I hypothesised that the use of an LLM for text generation is a form of ego protection, a way to emotionally divorce oneself from one's research output. I do not feel insecure about my writing,[8] but I do have insecurities related to other aspects of the job; insecurities that I considered fundamental, such as mental agility. As I will discuss further below, using LLMs as a first-pass through ideas *has* made me more self-confident, and less afraid of my own work. In that sense, they have succeeded at removing my ego from the equation. I am also aware that I have come firmly down on the defence of creative *writing*, as this is a skill I know I possess, and am therefore unwilling to surrender to an LLM. But other aspects of the job are creative too – mathematical creativity, for instance. This is something I have never had and never expect to have – unless I embrace augmentation by LLM.

My use of LLMs

Whilst I would like to pretend that I thought through these arguments in detail before starting to use LLMs, in actual fact my firmly held principles were brushed aside by nothing more than greed – the realisation that with LLMs I could work faster than my peers. That was the carrot, and the stick was that I'd accepted a job where LLM use was in the job description. I was obviously not forced to do this, but the allure of working in a place like Cambridge meant I did not hesitate for a second.

However, having taken the plunge, I began to realise I rather enjoyed the swim. Why? As mentioned above, I found that LLMs like Claude Code have started to impart a remarkable confidence in my work. I find them to be a safe space for asking stupid questions – questions that I probably learnt the answer to as a second-year undergraduate, but never actually understood. They allow me to spitball, to make mistakes quickly, to sanity check a thought process. In this way, using an LLM is like discussing an idea with an

[7] I may underestimate how widespread this skill is; see my previous point about the lack of exposure to narrative, rather than academic, text. Could you taste the difference between Châteauneuf-du-Pape and St Emilion, if you've only ever drunk the house red?

[8] The reader can judge for themselves whether this confidence is merited or not.

infinitely patient colleague who has bottomless knowledge of every topic in the world.[9] As someone who finds discussing science with colleagues really difficult, the opportunity to simulate that experience in a safe environment is seductive.

More than that, even. Historically, a large part of my job has been to write code. That is true of many cosmologists. We are blessed with an abundance of data, and must rely on numerical and statistical methods to infer things about our Universe from that data. As an undergraduate, my numerical methods course was taught in Fortran, a fact I have worn in later life as a badge of honour. I taught myself Python during my Master's project. I've dabbled in HTML and CSS. I have written a lot of code, and I consider myself to be reasonably good at it. But LLMs are better. I'd bet that they're better than most cosmologists at writing code. What does that mean for the practice of cosmology?

LLMs have forced me to re-evaluate the kind of work I value

The first realisation I had about LLM use in cosmology academia is that they are exposing a cognitive dissonance: that because a work is difficult, it must be valuable; and the corollary, that easy work is valueless. As stated above, I took pride in my ability to grind out Fortran syntax; I had spent years working out my programming muscle, and because that muscle became strong and beautiful, I cherished it, to the point where it was easy to lose sight of what I was actually *doing* with that code.

LLMs expose the truth: no matter how difficult the programming may be, the programming is not the point. The point of the code is science; but many of us spend our days acting like programmers, not scientists. With an LLM in hand, software that once would have taken me weeks, even months, to develop, can be working and testable within a matter of hours. This genuinely delivers on the promise of LLMs: that they will accelerate science, and remove grunt work.

This may be a difficult truth for some to face, including myself. How many of us have entered or ended up in this field precisely *because* modern cosmology is so reliant on software? The number of people who leave the field to successfully ply their trade as programmers in industry, be that under the title of data scientists, software engineers, quants and the like, speaks to how well training as a cosmologist prepares one for a career typing code. This is a shame; but it would be even more of a shame if we, as a field, did not grasp the opportunity LLMs are offering us with both hands. We have a brief window of time to restructure our work around *science* instead of drudgery.

And what of the objection that LLMs make mistakes? If a student is not sentenced to two years in the Fortran mines during their undergraduate degree, how will they be able to verify the code an LLM writes? There are two answers to this: the first is that the code an LLM writes is testable. Provided you have the machinery to compile and execute the code, you can see for yourself whether it works, and whether its output is reasonable.

The second is more fundamental: we should be teaching our students critical thinking. In the past, learning to write code was important, because writing code was important. If what will become important in

[9] I am blessed to have one or two real-life colleagues who embody this ideal. I would much rather have them on tap than an LLM, but they have better things to be doing with their time than answering my endless questions.

the future[10] is checking the output of an LLM, then we should be teaching students how to do that. Critical thinking is usually taught as part of the humanities: how to closely read and critique an essay for instance, or how to analyse the subtext in a passage of prose. My memory is not flawless, but I do not recall any similar instruction during my undergraduate physics degree. I suppose I was expected to learn how to do it myself – but instead of expecting students to intuit this, why not impose it?

LLMs are forcing me to understand why I value some forms of creativity over others

In thinking through my concerns about the erosion of creativity, I was forced to face another uncomfortable truth. I almost exclusively focused my discussion about creativity on the practice of creative *writing*, because I thought this was the skill most threatened by LLM use. It would be fairer to say that this was the skill *which I possess* that I thought was most threatened by LLM use. Here is another small shard of coloured glass in the cathedral window of human nature: we value the things we're good at.

It is not just that I think I'm good at writing; writing is a fundamental part of my thought process. Without writing, I understand almost nothing, even my own work. I write to myself as others perhaps talk to themselves. I have thousands of pages of 'work diaries' which I have been amassing since the middle of my PhD, in which I record everything I'm doing. I reason through writing. I therefore defend my right to write above all else. To give up writing – to surrender it to a machine – would be a swift and terrible lobotomy.

But by thinking about what LLMs are, and how they are being used in cosmology academia, I have been forced to reckon with the notion that not everybody values the written word in the same way that I do. Clearly, writing is not the only form of creativity. If it is drudgery for someone, who am I to deny them the possibility to ease that pain?

My original fear regarding loss of creativity tends along the same lines as the concern that LLM use will destroy people's ability to think and learn. Such statements have been made in recent and distant history about the internet, about dictionaries, and about printed books. LLMs will not destroy our ability to think, or to be creative; rather, they will teach us to think and learn and create in new ways. In a letter to writer and composer Heinrich Köselitz, in which he discussed his recent adoption of a typewriter, the philosopher Friedrich Nietzsche expressed it thus: "Our writing equipment takes part in the forming of our thoughts." [8]. As long as we are conscious of how tool use affects our output, I am satisfied that creativity will remain broadly unharmed.

LLMs should force us to carefully examine the metrics by which we measure success

As we have seen, LLMs have the potential to accelerate the practice of cosmology, thanks to removing any and all friction associated with writing and testing code. An important consequence of this is that

[10] The very, very near future. So near, in fact, that it is the present.

cosmology academia must now undergo radical change. In the past, we have valued results that are the first, the fastest, the most cited. Rarely do we take the long view, and honestly assess work on its quality, depth and contribution to the literature.

This mode of life is of course driven in large part by external pressures. Much as we cosmologists like to think of ourselves as free from earthly desires, we do need to live and pay for our science.[11] In order to calculate who should receive money, and how much, there must be a yard-stick by which we measure an individual's contributions to the science. We then assign the money to those who have contributed the most, or those we think have the potential to contribute the most in the future based on past performance.

The question then becomes about the yard-stick we choose. We have talked at length about the importance of the literature to cosmology. It would therefore seem an obvious choice to judge an individual by their contribution to that literature. This is where we run into problems. What aspect of a paper translates into material worth? The word count, or the number of figures, or equations? Clearly not. It has to be some meta-aspect of the work.

In cosmology, we have chosen to use the number of citations a paper has received as a proxy for its quality, but the reliance on citations and other publication-related metrics like h-indices is in jeopardy in the age of LLMs. Citations are often gained by the paper that showed things first, rather than the paper that showed things properly. If LLM use means that you will *always* be first, citations become a way to track how good a researcher is at prompting, and tell you next-to-nothing about the quality of the science presented.

However, citation count was a flawed metric even before the advent of LLMs. I have had the opportunity to work on a number of quite distinct topics, which nevertheless all fall into the broad bucket of cosmology. My works on beyond-standard-model cosmology (dark energy models, distance duality relation violation) have garnered far more citations than my works on strong gravitational lensing. There is no correlation of citation count with publication date or venue, and there is certainly no correlation with the quality of the science. It just so happens that the beyond-ΛCDM community like to cite each other with abandon,[12] whereas the strong lensing community is far more restrained (or perhaps we just read less).

What we can conclude here is that LLMs are forcing us to face the fact that, as a field, we do not do well at assessing ourselves and our peers. Very often we work on bad ideas, write bad papers, and get cited by others doing the same. LLMs will accelerate this process, but the underlying problem – the comparative lack of value placed on good science in the short-term – existed long before LLMs. We must summon our collective willpower to reform the field, and consciously decide to reward good ideas, rather than rapid execution. How we can actually achieve this is not yet clear to me, but I do not think it requires the exorcising of LLMs from our scientific practice.

[11] Another bitter pill we may have to swallow in the near-term is that politicians will get tired of us justifying the practice of fundamental research just because it is fundamental. Why do we do cosmology? [2].

[12] To cherry-pick an example which supports my point, compare the bibliography of this work [9] with this work [10].

Conclusions

The best thing that we're put here for's to see;
The strongest thing that's given us to see with's
A telescope.

Robert Frost

I have talked about my arguments against LLMs, and why these arguments failed to stop me from using LLMs. I discussed environmental concerns, and questions of plagiarism and creativity. I described why I like using LLMs, and how LLMs should force cosmologists to think carefully about how we evaluate our science in an era where production of code and text is trivial.

Having laid out my arguments, and thought about them carefully, I end this essay in a cautiously optimistic frame of mind. All things considered, my remaining anxiety about LLMs seems to boil down to verification. Can we trust ourselves to thoroughly vet every single output of the LLMs we use? In theory, we must commit to this, if we are to avoid contaminating the literature with not only things that are wrong, but things that are right yet unsourced. In practice, however, that Rubicon of trust was already crossed in cosmology long ago. When was the last time you repeated an N-body simulation just to see if the original authors did it right? What LLM use must really force us to do is consider whether we are happy and satisfied with this status quo.

In general, I like to think that I can trust my colleagues, and should therefore extend this trust to encompass their use of LLMs. On the other hand, I am currently privileged to be working at one of the best seats of learning in the world, where I am surrounded by talented individuals who give careful thought to all that they do. Is it fair to extrapolate my faith in their actions to the global population in current higher education? I also have no contact with undergraduates, and little experience of teaching. I cannot judge how LLM use will affect this sphere; but the centring of critical thinking in the curriculum would seem to be crucial.

As cosmologists, we have seen further by standing on the shoulders of giants. We have seen even further than that thanks to the incredible tool that is the telescope, which augments our human sight from peering around our own backyard to imaging the deepest reaches of space and the first syllable of recorded time. We did not shun that tool – we adopted it with open arms, hungry to see what it could do, and what it could tell us about our place in the Universe. Wielded judiciously, and honestly verified, LLMs have the chance to become the next tool to radically enhance our vision of the cosmos.

Acknowledgements

This work was enriched by discussions with James Alvey, Matt Grayling, Will Handley, David Hogg, Julien Larena, Chris Lovell, James Nightingale, and David Yallup; and by opinions shared in written form, both publicly and privately, by David Hogg, Juna Kollmeier, Hiranya Peiris and Roberto Trotta. I am grateful to Pierre Fleury for help with my red wine analogy.

The title is taken from The Fascination of What's Difficult by William Butler Yeats.

References

[1] Roberto Trotta. The indiscriminate adoption of AI threatens the foundations of academia. *Nature Astronomy*, 9:1748–1749, December 2025.

[2] David W. Hogg. Why do we do astrophysics?, February 2026. arXiv: 2602.10181.

[3] Juna Kollmeier. On Explicit Utility Functions in Physics: Compute for Compute, Physics for Understanding. https://thejunaverse.substack.com/p/on-explicit-utility-functions-in. Accessed 24/02/26.

[4] Scott Dodelson. Evolving Dark Energy and AI. https://scottdodelson.substack.com/p/evolving-dark-energy-and-ai. Accessed 24/02/26.

[5] Nidhal Jegham, Marwan Abdelatti, Chan Young Koh, Lassad Elmoubarki, and Abdeltawab Hendawi. How Hungry is AI? Benchmarking Energy, Water, and Carbon Footprint of LLM Inference, May 2025. arXiv: 2505.09598.

[6] Robert K. Merton. *On the Shoulders of Giants: The Post-Italianate Edition*. Chicago University Press, 1993.

[7] Ahmed Ahmed, A. Feder Cooper, Sanmi Koyejo, and Percy Liang. Extracting books from production language models, January 2026. arXiv: 2601.02671.

[8] Nicholas G. Carr. *The Shallows: How the Internet Is Changing the Way We Think, Read and Remember*. Norton, 2010.

[9] Sunny Vagnozzi. Seven Hints That Early-Time New Physics Alone Is Not Sufficient to Solve the Hubble Tension. *Universe*, 9(9):393, 2023.

[10] Philip Holloway, Aprajita Verma, Philip J. Marshall, Anupreeta More, and Matthias Tecza. On the detectability of strong lensing in near-infrared surveys. *Monthly Notices of the Royal Astronomical Society*, 525(2):2341–2354, October 2023.