
The Age of the Architecton: Redefining the Landscape of a Scientist in the Era of AI

Matt L. Wiemann
Princeton University
April 2026

Abstract

The rapid advancement of AI systems is fundamentally reshaping what it means to be a scientist. While technical competence becomes commoditized, the most valuable scientific skill remains the ability to ask the right questions, identify the right problems, and turn fragmented pieces of knowledge into science. Drawing on Kant’s concept of the Architecton, this essay argues that the future scientist is not a labourer displaced by machines, but a master builder empowered by them to create things greater than previously possible. It offers practical guidance for early career researchers navigating this transition and explores how one may cultivate this ability. The central argument of this essay is one of optimism.

“Without systematic unity, our knowledge cannot become science; it will be an aggregate, and not a system. Thus architectonic is the doctrine of the scientific in cognition, and therefore necessarily forms part of our methodology.”

— Immanuel Kant, *The Architectonic of Pure Reason*.

1 Introduction

THE HISTORY OF SCIENCE is not a history of tools and instruments, but a history of questions. That is not to say that tools and instruments hold no scientific value, in fact shortly I will argue quite the opposite. No, my point is that science, what it means to be a scientist, is about asking the right questions at the right time. This has historically been true, but there are louder and louder claims of the ability of AI to in fact ask these questions, moving itself from a tool, to the scientist.

I will not speculate too heavily on what a future AI enhanced world may look like, nor will I try to define any notions of machine *consciousness*. This essay outlines the current capability of AI systems for science, where they are directly headed (in the foreseeable future) and how this changes how one should approach forging a scientific career.

2 What can AI do, and whats going away

Current AI tools are dominated by large language models. We have models capable of performing a wide range of technical tasks, from programming to formal mathematics verification, at a level close to and often exceeding that of a human expert. Consider AlphaFold, which effectively solved the protein folding problem, a challenge that occupied structural biologists for decades. It is an extraordinary achievement, yet AlphaFold did not choose which proteins to study, nor did it ask what the resulting structures mean for disease, evolution, or drug design. It answered a well-posed question with remarkable competence. So where does this leave us? Any field that requires even modest amounts of data analysis or programming is already starting to radically change. Technical competence will no longer be a marketable skill, it will be a requirement. You will not land a job based on being one of the few people who can work with some field-specific data pipeline.

There are still tricks of the trade that LLMs may not know about, best practices in your field that do not get written down. But this will change. LLMs will continue to accumulate knowledge until they are indistinguishable from domain experts in every scientific field.

I now wish to make a key distinction: what is being devalued is routine implementation, writing boilerplate code, running off-the-shelf models on new datasets, building standard data pipelines. Work where the task is well laid out and there is little room for creativity. Clever implementation, on the other hand, is a fundamental scientific endeavor. Bringing an idea to fruition often requires rethinking methods, inventing new approaches, and in the process providing the community with ideas that previously didn't exist. The team at DeepMind didn't run a standard pipeline, they created AlphaFold through profoundly creative implementation. That kind of work is the modern manifestation of the Architecton.

3 A case for optimism

There is much negative sentiment about the rise of AI in the sciences, from both traditional scientists unwilling to trust what they cannot understand, and those who see so much value in current models that they fear being supplanted. I truly believe AI will be the single biggest driver of scientific progress the world will ever see. Yet AI will not sit in the driver's seat. AI will empower scientists to explore ideas at both greater depth and breadth than previously possible.

For AI to truly take the job of a scientist, it must not only be capable, but motivated to ask the types of questions our leading scientists ask today. I have seen no evidence of this, nor any progress towards creating intrinsically motivated AI systems that would actively choose to solve scientific problems.

Despite how it may feel, any motivation an AI model displays is an imitation of the motivations we have imposed via reinforcement learning. Future world-model approaches may embed causal dynamics into an AI agent's perception, achieving what we now call artificial general intelligence. But achieving AGI is not the same as creating an intrinsically motivated agent. If unprompted, what will these models *desire*? Will they have any interest in probing the secrets of the universe?

This ability is far closer to consciousness than it is to scaling up current capabilities. We do not have a well-defined metric for consciousness, and as such have no way of tracking progress towards creating it. However, many consider the problem of continual learning, a model's ability to retain old information when trained on something new, a required path towards this. Many researchers are working on the problem of continual learning, yet it currently remains a distant goal. Perhaps of more importance is the fact that solving the continual learning problem is not a requirement to build incredibly useful, and extremely valuable, future AI agents. So I think one has cause to be optimistic that for the foreseeable future, AI agents will continue to be a tool, an all-knowing oracle perhaps, but an oracle nonetheless. One that speaks only when spoken to, and only about what it is asked.

4 The Architecton

In ancient Greek, an Architecton (*architektōn*) is defined simply as a "chief builder" or "master craftsman". This direct definition has had many interpretations, but my favorite, and the one I use for the premise of this essay is that of the German philosopher Immanuel Kant. To summarise the thoughts from Kant's writing in *The Architecton of Pure Reason* (see [here](#) for chapter 3) we must understand his idea of the distinction between knowledge and science. As Kant implies, the Architecton turns knowledge into science, it is the imposition of systematic unity by a mind that begins with a governing question and knows where each piece belongs. The Architecton provides that unity. They are the thought leader able to create something

greater than the sum of its parts.

Kant, who lived from 1724-1804, could not have anticipated artificial intelligence, but his framework provides a strikingly apt description of it. An LLM accumulates knowledge, vast, fragmented, unsystematic pieces of information. It can retrieve, recombine, and even generate new fragments. But it cannot, on its own, provide the fundamental questions that turn this aggregate into science. That remains the work of the Architecton. The scientist who understands not just what the data says, but why it matters, which question it serves, and where it belongs within a larger structure of understanding. AI provides the raw material in abundance. The Architecton is the one who knows what to build.

5 To have a good idea

Many people, far more experienced and articulate than myself have proposed an answer to this question. And while answers vary widely by individual, and by research field, the best resource I have come across for this is the talk (and transcription of) [You and Your Research](#) by Professor Richard Hamming. In this talk, Hamming describes the importance of developing research *taste*. A skill to identify problems which are interesting, and also important. Hamming states: "The importance of a problem to a great extent depends upon, have you got a way of attacking the problem?". Here is the key, one must be ambitious enough to try something new, yet wise enough to know if we have the tools capable of tackling it.

A key point that I think is often lost on the early career researcher is the reality of the progression of a good idea. In our schooling and tertiary studies we are often presented with well posed problems, ones which having a good understanding of the material, and then a clever idea about a solution is immediately rewarded by the correct answer. We ace the exam, pass the homework sets, and pat ourselves on the back. We are implicitly learning that if an idea is good, it should produce relatively rapid results. This is a very dangerous mindset to get into, and I think undersells an important part of scientific problem solving. Clever implementation *is* problem solving, it is science.

Let me pose a hypothetical scenario where we live in a world with a finite set of good ideas about the answer to some scientific questions. Let us then say half of these ideas/solutions are easy to implement and provide rapid signs of success, and the other half require far more work to bring to fruition and have much higher chance of being abandoned. If one assumes that we uniformly sample from this total set of good ideas, then over time the amount of easy to implement, rapidly successful ideas will soon diminish by comparison to the difficult yet equally valuable ideas. Really, this is quite similar to the argument of low hanging fruit, but for some reason this is still often forgotten as many people spend time searching for the good *and* easy ideas, as opposed to making the hard ones work. Moving forward, if LLMs can now rapidly test and implement the easy ideas, then the pool of remaining human-valuable work

shifts even further toward the difficult, long-horizon problems. A caveat to this, would be that it is in fact a practical issue, since as an early career scientist you often live on short term contracts, and one must produce tangible results by the end of this contract to climb to the next rung of the ladder. However, this practice still reinforces the tendency to chase easy ideas, a habit I think can be very hard to shake even when job security is achieved. So it is important to always have the concept of a good, and hard idea in mind.

My PhD supervisor Peter Melchior fits the definition of someone who has good ideas, and embodies the role of the Architecton that I have laid out. I have learned a great deal about the process of bringing these ideas to fruition, through extended brainstorming sessions, the countless hours of trying and failing with new implementation tricks, and perhaps most importantly, the willingness to try things the field may consider unwise, even foolish. Things almost never work on the first attempt, and if they do, one should be deeply suspicious. This experience has shaped my own research taste in ways that reading papers alone never could have. And that distinction matters. We tend to encounter researchers we admire and would like to emulate through their publications, which are polished, finished objects that often bear little trace of the months or years of failure that produced them. Hence to emulate this researcher one should not aim to emulate the result, but the process, and the most reliable way to learn the process is through direct mentorship, working with someone day to day. Only then can you see how an Architecton truly creates their master work, and hope to one day become one yourself.

To summarize, to me a good idea is one which is not obvious, more often than not it is not easy to implement, yet if successful the results have a wide impact on one or multiple fields. I also think a key to a good idea is that some part of it must challenge a pre-existing norm in the field. If not, then I think the work falls closer to incremental improvement, work which is far easier to envision a future AI agent would be able to achieve with minimal to no human interaction.

6 Rethinking scientific careers

Every year we produce more and more PhD graduates while the amount of faculty track jobs remains essentially stagnant, therefore it is more important than ever now to stand out early. This is a fundamental problem with academia, and really, not one that most people should concern themselves. I say this as the overwhelming majority of people that may come across this essay will not now, nor ever be in a position to make any meaningful change to this system. In fact, to be in a position to make these changes more than likely requires one to abandon their own scientific dreams in favour of taking on administrative responsibilities. And that is not the intended audience of this essay, this is intended as a guide for scientists, and those who wish to continue to practice science for the remainder of their careers.

Well, what if I wish to pursue science in an industry role, should my approach change too? In short, yes, but for different reasons. Ideally a scientist desires full autonomy of their research agenda. In academia, this is a given, in industry this is far from the case. In fact, most industry roles include joining a larger research team often led by a single (though granted very impressive) research lead. Hence the goal of the true scientist in industry is clear, you want to be the research leader, to have the tools and manpower to pursue your ideas. And this is where I draw the similarities between the two scientific paths, as the end goal of both is the same, and the pipelines just as competitive.

So what do we need to rethink? In short, marketing. The way we market ourselves, especially as early career scientists must radically change. In times gone by, at the conclusion of one's PhD the goal was generally to be considered both competent in your field, i.e know the specific tools of your trade, and an expert on one specific sub-domain. You are the one who works on X, the one who built Y etc. This was enough to continue to the next step in gaining a postdoc, finding an industry role. The biggest change I believe to come, will be the continued devaluation, and expendability of competent researchers. I think it is already clear that technical competence has rapidly lost value, but what about one's value as a sub-domain expert? Unfortunately the value of this is also in rapid decline, even present day LLMs are incredibly knowledgeable about a range of domains, and given sufficient prompting can explain links between research fields outside one's direct area of expertise. A single experienced researcher can now become familiar with new domains at an electrifying pace. But the roles that are becoming increasingly more valuable is that of the **Architecton**, the ideas-person, the research lead. Of course, these roles are far more exclusive and therefore the applicants profile will undergo much more scrutiny, with the main trait one *must* demonstrate being the ability to have novel, and useful ideas which are beneficial for a wide variety of downstream applications. I.e to have *good* ideas.

This requirement of demonstrating oneself to be a thought leader is nothing new, in fact this is mostly a requirement at a faculty hiring level. The thing that is changing, and has perhaps changed already, is the stage in one's career in which they must demonstrate this ability. The method of demonstrating this skill is also changing, academic papers are becoming an ever weaker signal of an individual's quality, so one must take to the road. I strongly believe that now more than ever, early career academics, especially PhD students must spend more time presenting their work in person, networking, and brainstorming with more established scientists. This is the single best way to demonstrate yourself to have creative interesting ideas, as more often than not, over the span of a PhD your most ambitious goals will not have had time to produce tangible results, your unique insights may not fit in the confines of an academic paper. People must view you as a budding Architecton, someone who if given the resources, can truly produce something great. A paper can show what you have done, but it cannot show what you are capable of. That distinction is made in person, in conversation, in

the way you think on your feet when someone challenges you over coffee at a conference, or at the end of a talk. Papers do not show potential. People do.

7 Recommendations

While I have now laid out why I believe this is one of the most exciting times to be a scientist, I do think there must be some changes to how we train, and prepare to enter the field. Most prominently, at the early career and especially doctoral training stage.

7.1 One must learn the tools

The first non-negotiable is to become intimately familiar with AI-based tools and workflows. Just as programming has become fundamental to so many areas of science, utilizing AI tools will soon be the same. Resistance to adopt these tools will not remove a potential advantage, but will actively handicap your capabilities compared to your peers.

7.2 One must effectively wield the tools

Using AI tools is not the same as utilizing AI tools. AI tools are a productivity multiplier, however if you are not asking the right questions, not using them with skill and sophistication you will simply produce a larger pile of *slop*. Coding agents specifically are great tools if you know their limitations, they are still an LLM, every word is just a token, a numerical representation. Just because one of these words may be your all important function name, `solve_MOND()`, does not mean it is seen as anything special by the coding agent. One must be excruciatingly clear with instructions and how the code operates, and ever cautious about building something that you do not already have a pretty good idea how to do on your own.

7.3 Develop taste

One must learn to develop research taste. The most direct route is through direct mentorship with an admirable scientist, however this is not always attainable for a variety of factors. In the absence of an advisor, you must resist the temptation to pursue the easy ideas at the expense of the hard ones. It is only through looking back and seeing why your last 6 months was a futile effort that you can learn to identify the early signs of futility, and slowly build intuition about what things are good to work on. I also recommend broadening your horizons about what it means to conduct research and ask scientific questions, go to talks outside your field and note how they conduct *science*.

7.4 Show your thinking

One can no longer expect a well constructed paper to be enough to convince someone of your value in the modern scientific age. To be blunt, a paper cannot truly show how one thinks, and it is precisely how one thinks that will be the most important skill to demonstrate at all

levels of a scientific career. Hence you must take the show on the road, present at conferences, group meetings, ask to join lab meetings of groups you find interesting. You need to maximize your exposure to people who make the hiring decisions, and crucially, they must view you as a budding Architecton.

8 What if I am wrong, and our models do become *intrinsically motivated*?

Well, in my opinion, there is really no point in preparing for this scenario as a scientist. We will end up consumers of ideas, passengers in a new age of discovery. In an ideal world, we may still reap the benefits of this society, but as time goes on we will have a smaller and smaller part to play in its continued improvement. One can argue we may still have a role to play in learning to understand these complex, and intellectually superior beings to ourselves, but even so, would we be able to determine a more effective way to communicate than they would?

Going down this rabbit hole leads to some dark, uncomfortable places. Even in a world of abundance, is it not still human nature to long for a purpose? I also think this is a scenario where no amount of prior planning can help, hence I do not recommend one torment themselves in trying.

9 Conclusion

The role of the scientist is changing, but it is not going away. If anything, the arrival of AI has enhanced what has always been the most important trait of a scientist, the ability to conceive of what does not yet exist, and the creativity and persistence to be able to construct it. Unfortunately due to the nature of academia, and large industry labs, many roles are given the title of *scientist* when in reality the work is closer to a hired gun, paid for routine technical competence, and to bring forth the ideas of the Architecton. I say this not to diminish these roles, historically these roles have been crucial to bring about almost every technical revolution we have had today. But it is these roles that are at risk of being automated away, and as we have previously labeled these roles to be the job of a scientist one may naively think that indeed the role of the scientist may be a thing of the past.

Kant's Architecton was never the one who laid the bricks. They were the one who knew why the building needed to exist in the first place, and what new approaches may be needed to allow for its construction. That distinction is more relevant now than at any point in history. Labour is devalued, knowledge is cheap. What remains scarce, and what will only become more valuable, are the minds that know what to create, and how to synthesize these fragments of knowledge to create it.

I hope the reader leaves with a sense of optimism, as well as urgency. It has never been a more exciting time to be a scientist, but to truly be a scientist is not an easy task. We can no longer view routine implementation as science, as these are the roles that will go away, and these are the roles one must avoid if they truly want to take part in this scientific revolution.

The age of the Architecton is here. Prepare accordingly.

Acknowledgments

This essay has been inspired by numerous conversation with my classmates, and faculty at Princeton. Particularly from discussions with Peter Melchior, Christian Jespersion, Lizhou Sha, and Andrew Saydjari in our machine learning journal club started by Benjamin Remy, as well as discussions and editing tips from Angelina Yan.

References

- [1] Hamming, R. W. (1986). You and Your Research. Transcript of a talk given at the Bell Communications Research Colloquium Series, Morris Research and Engineering Center, March 7, 1986. <https://www.cs.virginia.edu/~robins/YouAndYourResearch.html>
- [2] Kant, I. (1998). *Critique of Pure Reason*. Translated and edited by P. Guyer and A. W. Wood. Cambridge University Press. Originally published 1781; See especially “The Architectonic of Pure Reason,” A832/B860–A851/B879.
- [3] Jumper, J. et al. (2021). Highly accurate protein structure prediction with AlphaFold. *Nature*, 596, 583–589. <https://doi.org/10.1038/s41586-021-03819-2>