FEATURES

THE LAST ASTRONOMERS

Amid a flood of AI advances, astrophysicists are questioning the soul of their field

JOSHUA SOKOL

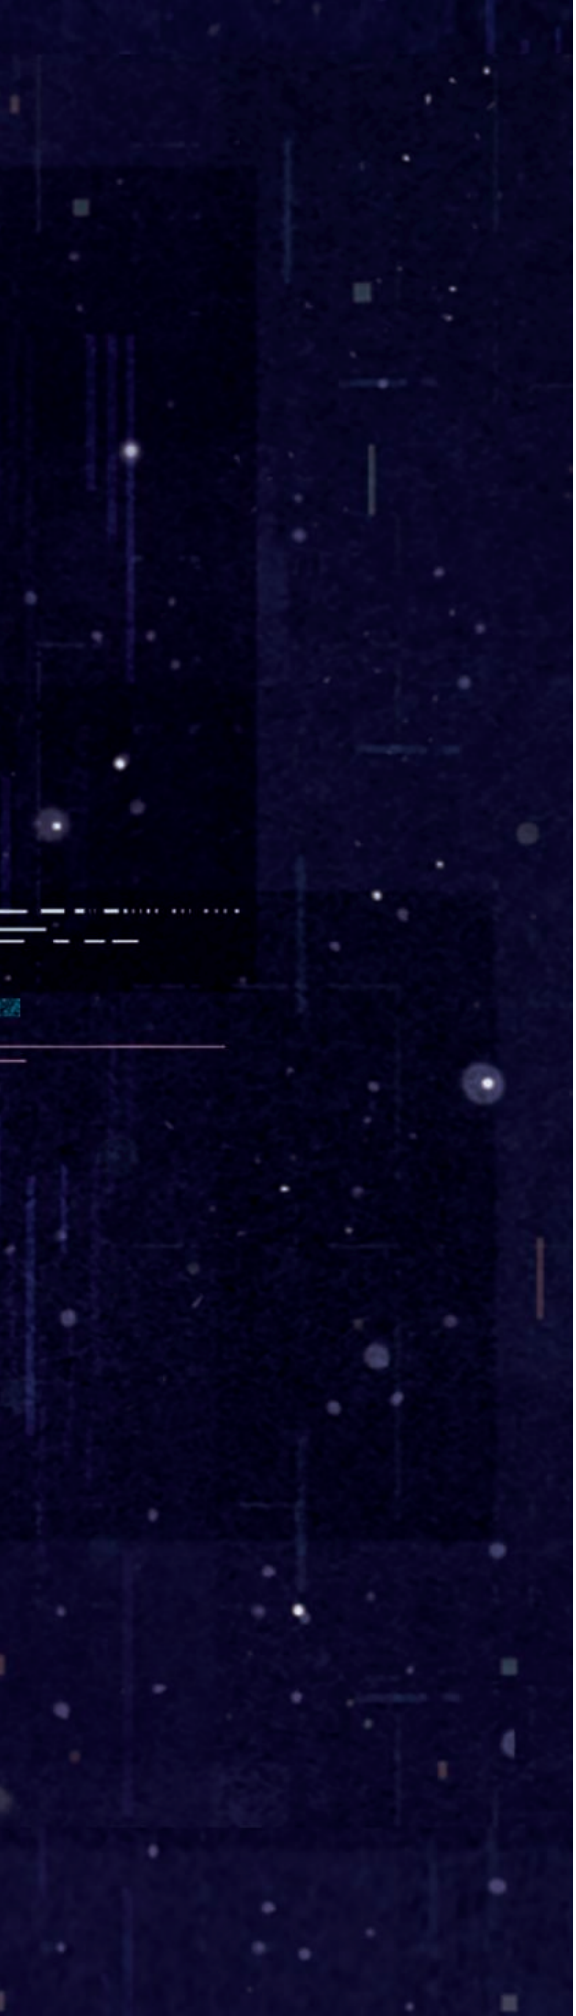

ILLUSTRATION: CHIARA VERCESI

One afternoon in April, Cecilia Garraffo settled down at the head of a conference room table in Cambridge, Massachusetts, and gazed out at what might be the last astrophysicists of their kind.

The walls of this room had, in the past, reverberated with the din of thousands of other groups of scientists. Now, as streaks of sunlight poured in, the discussions turned to nonhuman collaborators. One by one, the gathered researchers discussed how they planned to apply machine learning to problems in astronomy. Observing an interstellar comet. Discerning wispy filaments of galaxies at the universe's largest scales. Developing a new "tokenizer" that can translate astrophysical images into a form more readable by artificial intelligence (AI). "Sometimes models will be overconfident," Garraffo warned a junior team member.

Afterward, as everyone filed out, black hole researcher Daniel Palumbo made a brief announcement. Representatives from AI chipmaker NVIDIA were on campus in search of scientists who wanted to solve problems using their hardware. To anyone who might need extra processing power, "today's the day," he said.

The Center for Astrophysics | Harvard & Smithsonian employs more than 600 astronomers and other staff, making it one of the world's largest concentrations of professional stargazers. Garraffo heads its AstroAI group, which is charged with leading the center's approach to applying machine learning to various problems. In just 4 years since they first proposed this specialized team, Garraffo and her colleagues have forged collaborations throughout the building and with industry teams such as Google DeepMind and Anthropic.

Originally, their goal was to use machine learning and AI to remove the technical barriers of math and computation while preserving what Garraffo considers the fun part of physics: honing scientific questions. Their toolbox did not include chatbots. Despite the buzz around ChatGPT, which was released within months of AstroAI's first proposal, Garraffo had thought her group would steer clear of the headline-grabbing tool and other large language models (LLMs).

At least, until recently.

Now, stories of miraculous progress were starting to spread across the institution. As her MacBook Air pinwheeled to a crawl thanks to an AI agent running software locally, Garraffo's colleague Alyssa Goodman showed me a data-fitting problem. She wanted to understand how the spiral arms of a distant galaxy were moving. But isolating just that motion from other patterns imparted into her data by the spin and the geometry of our own Galaxy had thwarted her group for years. She asked ChatGPT, which resolved the problem in a few minutes. Now, her research group was planning to write several papers on the resulting data set, "the single best map of spiral arm kinematics ever—like, by a factor of 100."

Conversations comparing these tools with human researchers, once an underground whisper, have grown into a deep rumble in astrophysics departments around the world. Many are turning over aspects of their research practice—searching the literature, developing code, writing proposals to use telescopes, doing first-pass "reads" of their peers' submitted proposals, and actually solving problems—to agentic AI systems such as Anthropic's Claude or OpenAI's Codex. Major institutions such as the Space Telescope Science Institute and the Institute for Advanced Study have rushed to hold meetings on LLMs and AI agents in science. In March, Anthropic published a blog post from Harvard University physicist Matthew Schwartz about his experiments in "vibe physics." After supervising Anthropic's Claude model closely enough to catch its many fabrications and bluffs, he forced it to generate in 2 weeks a real, publishable physics paper that he claimed would normally take a year. Schwartz's sense was that the "AI grad student" approximated a second-year grad student at Harvard. Give AI 12 more months, Schwartz extrapolated, and LLMs' capabilities may rival those of postdocs.

Representatives of the AI companies, who seem to view astrophysics problems as public relations boons that offer compelling showcases of their models, boast that their technologies will soon achieve supremacy over actual theoretical physicists, astrophysicists, and cosmologists. Some even make mechanizing the study of the night sky a selling point. In February, as Elon Musk began seeking to take SpaceX public to raise cash for orbiting data centers, his company explained its true goal as "scaling to make a sentient Sun to understand the Universe."

Already, by making it faster and easier to produce professional-seeming papers, AIs threaten both to overwhelm journals and peer reviewers and to take opportunities away from junior scientists. But far upstream of that, many scientists interviewed by *Science* sense a phase change underway. Many fear that if unleashed in all parts of the scientific process, AI tools could lead to nothing less than the death of astrophysics as a human endeavor. "A lot of people think that it's too late to intervene—we're done," says David Hogg, a computational astrophysicist at New York University (NYU).

ALTHOUGH FEARS OF A ROBOT takeover are inspiring soul-searching across society in general, astrophysics is a strange, special case. It relies on reservoirs of prestige, deep public appeal, and invocations of wonder: Human beings and funding agencies alike just seem to like studying the stars, and astronomical progress has long been thought of as a synecdoche for human progress in general. Yet because astrophysics is already mostly data science and math, many of its juiciest problems may be low-hanging fruit for LLMs.

At the same time, understanding the night sky promises little in the way of economic applications or human lives saved, so progress at all costs doesn't seem as imperative as it might for fields such as drug discovery. If any scientific discipline could stake out a nuanced, human-centric relationship with powerful new AI tools, it would be astrophysics. So how is that going so far?

Early signs are, in a word, mixed. In September 2025, a guest speaker at the NYU physics department ran an AI agent in real time in the background. As he spoke, the system—called Denario and built by a group at the Flatiron Institute, a privately funded research center dedicated to advancing science using computational methods—generated entire scientific projects. It scoured journals, spun out ideas, carried out analyses, and extruded professional-seeming scientific papers—some obviously goofy, some plausible—that popped up on the screen behind him. With tools like this and beyond, he said to an audience of mostly grad students, you no longer need grad students. Why wait months for a young human scientist to do a project when an AI can give you the answer within an hour?

> "What I love is to chase the truth."
>
> **Cecilia Garraffo**
> Center for Astrophysics
> Harvard & Smithsonian

"It was a very bizarre talk," says Matthew Daunt, a seventh-year grad student at NYU, because even Denario's most impressive outputs didn't look all that scientifically useful. "And yeah, the comment of 'You don't need grad students anymore' kind of pissed me off."

Hogg, who also works at the Flatiron, had a one-on-one meeting with Daunt right after the lecture. "He was like, 'I'm not cattle,'" Hogg says.

The incident inspired Hogg to think about the moral value of AI systems studying the universe—and the value of the humans that such systems may replace. Simply banning the use of AI models would require an impossible amount of policing, he felt. But if astrophysicists simply sit back and let these models "cook," their hundred-thousand–fold advantage in speed would mean scientists would be quickly buried under an unreadable avalanche of machine-only studies.

It seemed to him that finding a middle path between these extremes, in which LLMs help science and scientists alike, involved first wrestling with a question that seems less in the realm of physics and more like metaphysics. In February, he posted his paper to a preprint server. Its title: "Why do we do astrophysics?"

Hogg offered an answer. It wasn't to solve the cosmos but to grapple with it: a journey, not a destination. Graduate students aren't supposed to be only the means of the science—doing work for senior scientists—but an end in themselves, molded into competent scientists by doing that work. In so doing, today's students would become the latest link in an unbroken chain of practice, going all the way back to the first among us who looked up agog at a sky full of stars. "Anyone working in astrophysics," Hogg wrote, "is someone who wants to *do astrophysics*, not someone who wants to *learn the answers*."

There are only a few thousand astrophysicists worldwide. Perhaps 100 email responses came in within a few days after Hogg posted the paper. Some correspondents disagreed. Many thanked him for dragging the conversation from coffee shops into the literature. "Flatiron fell apart," Hogg says. In frenzied conversations, many of his own colleagues were compressing their internal timelines for when AI might bring major changes to their science from years to weeks—and fretting about the consequences.

Hogg had suggested I could visit the Flatiron to hear some of these conversations in person. But after weeks of emailing, a member of the Flatiron's communications team turned me down. The institutional politics around AI and systems such as Denario, they told me on background, had become so fraught that my request stirred up a hornet's nest of internal conflicts.

ASTRONOMY IS NO STRANGER to advances in quantitative methods, having arguably given birth to them in the first place. Thousands of years ago, scholarly orders from Mesopotamia to Mesoamerica to China and beyond served as intermediaries between their societies and the cosmos by combining extensive bookkeeping with clever mathematical algorithms—what

ILLUSTRATION: A. MASTIN/*SCIENCE*

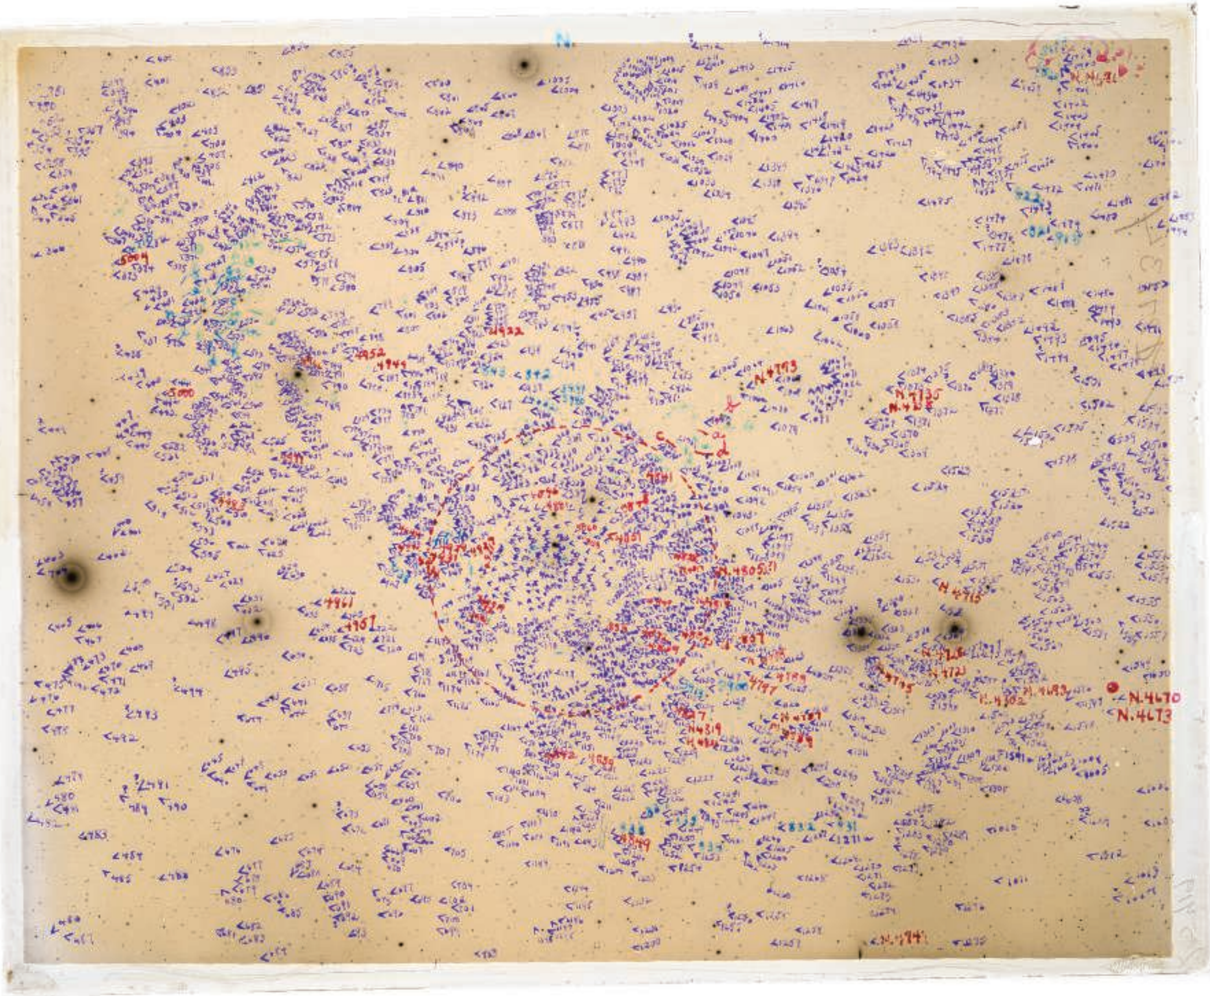

Annotating a 1934 photographic glass plate by hand, Harvard University "computer" Muriel Mussells Seyfert found 431 galaxies within the dashed red circle marking the center of the Coma Cluster.

PHOTO: PLATE MC27134, 16 MARCH 1934/HARVARD PLATE STACKS/CENTER FOR ASTROPHYSICS | HARVARD & SMITHSONIAN

might now be called data science. In the second century B.C.E., Hipparchus accessed and reprocessed centuries of Babylonian records to develop theories of solar and lunar motions. At the beginning of the 1600s C.E., Johannes Kepler used Tycho Brahe's observational data to fit mathematical laws to the motions of the planets.

As the supply of astronomical data exponentially increased, astronomers and mathematicians pioneered many of the first "computers": human beings scribbling on paper. One early program employed teams of hairdressers put out of work by the execution of the aristocracy in the French Revolution to calculate logarithms. By the 1820s, mathematician Charles Babbage, dreaming of automating mental work the way the steam engine had automated physical labor, designed the first programmable computing engine to make astronomical calculations.

By the end of the 19th century, photography had began to be applied to the study of the night sky, leading to another explosion in data and computing needs and rendering a human eye on a telescope eyepiece irrelevant. Photonegatives of distant galaxies piled up in basements in places such as the Harvard Observatory. More scientific data could be extracted, astronomers understood, if only these images could be studied in detail.

In the Harvard Plate Stacks, three levels of hulking black cabinets store some 600,000 thin glass plates from this era. To process all these data from the cosmos, curator Thomas Burns explained, Harvard astronomers brought in a new workforce of woman computers. The women got a chance to make heroic contributions to a field that would otherwise have excluded them. But they were also exploited. Some, from working-class backgrounds, needed a paycheck too much to negotiate. Others, from wealthy families, didn't need the paycheck at all.

Burns descended a narrow spiral staircase and fetched a glass negative from 1934. Photons from thousands of galaxies around the Coma Cluster had traveled 300 million light-years to Earth, where a chemical reaction in the film had trapped their energy like fossils in amber. The glass, illuminated against a white backlight, was covered in tiny markings in different colors. First a "computer" named Muriel Mussells Seyfert had labeled countless galaxies. Then others had revisited the same frame in other passes, leaving lasting traces of a layered dialog.

Burns explained that Seyfert, uncredited, had helped calculate the mass enclosed within a red circle around the center of the cluster. The result corroborated an odd contemporary proposal that the galaxy cluster harbored some extra, unaccounted-for "dark matter." Another astronomer, Vera Rubin, came across these same underappreciated data in the early 1950s. Perhaps inspired by this hint, she would later go on to find even firmer evidence for the existence of dark matter.

Another three-quarters of a century later, in March, NVIDIA announced its new Space-1 Vera Rubin Module, designed for use in swarms of orbiting data centers. Their proliferation could pollute astronomical images—and the AI models they enable could render astrophysics itself unrecognizable. I asked Garraffo about the scenario in which her field produces fewer human stories like the one captured in this annotated glass plate. What if people—such as Seyfert, or Rubin—no longer make these kinds of discoveries themselves?

"What I love is to chase the truth," Garraffo replied, considering. Perhaps the moment of human insight frozen on the plate was romantic, but the work had clearly also been tedious—and to her, an elegant truth of the universe such as general relativity was more romantic. If tools exist that would lead to such truths faster and better, she argued, astronomers should use them. "Pretending, 'Let's cover the truth so people feel like *they* can discover it'—it's very narcissistic, in a way."

Garraffo isn't sure when AIs might surpass human capabilities, but a recent experience convinced her they haven't yet. In her early career, as a pen-and-paper theoretical physicist, she had worked on alternate versions of Albert Einstein's theory of gravity, general relativity, which is known to be incomplete. Garraffo wanted to solve the equations for one alternative, called Einstein-Gauss-Bonnet, and find an exact description within the theory of the shape of spacetime around a rotating black hole. It was a practice case, she thought: If you could figure out how to do that, the same conceptual tools might work across other alternate theories of gravity, giving scientists a new way to probe and evaluate new theories that aim to usurp Einstein's. She trained a neural network to intuit its way to an approximate, numerical solution. But this kind of network was not an LLM, and it couldn't rewrite its inner workings out into a neat, concise mathematical answer.

Then she asked the best publicly available versions of both Claude and ChatGPT to solve the Einstein-Gauss-Bonnet equations for a rotating black hole analytically and produce the kind of crisp answer she sought. It was a very hard but well-posed mathematical problem, and she was optimistic it could work. "The models I tried failed miserably. Miserably," she said. Claude, in back-and-forth exchanges, claimed it *could* find an interesting result for the problem, then wrote up an entire impressive-looking paper. But as she read the paper, she realized it was just stating existing ideas in byzantine ways.

If the LLMs had been able to make progress on the problem, and if AI systems stopped improving right at that level, "that to me, maybe selfishly, would be ideal," Garraffo said. With tools of that strength, humans like her would still need to identify the relevant questions, ask them, interpret the answers, and have the fun of making sense of what that meant for physics. "Then we have to press the brakes."

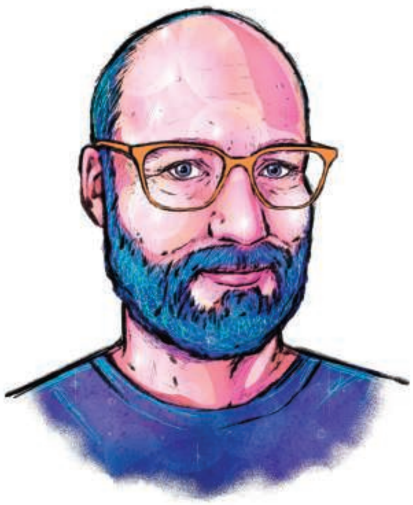

> **Anyone working in astrophysics ... is someone who wants to *do astrophysics*, not someone who wants to *learn the answers*.**
>
> **David Hogg**, New York University

But the power of the models may soon surpass that level. "We all collectively came to the realization that these tools are about to take over," postdoc Rodrigo Córdova Rosado told me in his office earlier this spring, flanked by a LEGO X-Wing on top of a shelf of physics textbooks. In about 2 hours, he told me, he had used Claude to interpolate between the best existing general relativity textbooks and write a new one, complete with sample problems and figures generated by scripts it wrote in the math software Mathematica. He was now checking the book for errors. (On the side, as a student of the Osage language and a member of the Osage nation, he had been experimenting to see whether AI might be able to create an Osage-English dictionary.)

"This is a wave," he said. "If you do not surf it, it feels like you're gonna get drowned by it."

DURING MY VISIT, Garraffo invited team members who wanted to talk more philosophically about AI to come to her office. Seven scientists packed in, two more joined remotely, and they began an impromptu roundtable discussion.

The conversation grew loud. Do AI agents lack good "taste" in scientific problems? Were these tools democratizing, helping non-English speakers participate in science—or were they antidemocratic, by making science dependent on subscriptions to some of the wealthiest corporations in the world? Were astronomers now free to care less about computer science and more about the universe? Would the fact that LLMs are inherently probabilistic—and not unerring engines of mathematical logic—make the math and statistical analyses they spit out untrustworthy? "Is it insane," one grad student exclaimed, "if I don't see this as such a big issue?"

Both in that room and outside it, many of astrophysicists' concerns lump into two families. First, many researchers fear LLMs' ability to effortlessly conjure text, code, and highly technical analyses threatens the current metrics, incentives, and trust networks that knit "astrophysics" into a cohesive, reliable body of knowledge.

For example, in a publish-or-perish world, publishing papers at a fast rate and racking up citations has been treated as a proxy for the quality of a scientist's output, which in turn determines who gets resources to do research at all. But now publishing an incremental but impressive-looking paper seems cheaper and easier than ever. "LLMs are forcing us to face the fact that, as a field, we do not do well at assessing ourselves and our peers," Natalie Hogg, a cosmologist at the University of Cambridge, wrote in a February blog post.

Ethan Vishniac, editor-in-chief of journals published by the American Astronomical Society (AAS), is already dealing with this issue. Since LLMs first became popular, he says, paper

ILLUSTRATION: A. MASTIN/*SCIENCE*

submissions have surged, making it harder to find reviewers. Much of the increase from established astronomers has since leveled off, but paper submissions from outside the known community have continued to increase. In April, for example, a "20-year-old working alone from Chennai," India, claimed in a blog post that an end-to-end science pipeline of his own creation had conceived of, executed, and written a study on a distant source of radio signals. That paper was accepted in March by an AAS journal but was later rejected for not disclosing the nature of its AI use, his post said.

Traditionally, AAS editors have responded to every submitter with personal feedback. But now, Vishniac says, the correspondence sometimes gets so "weird" that he suspects he's talking to agents, not people. Rather than democratizing science, these tools may soon force journals in the opposite direction. "The quantity of things of low quality can strangle the system," Vishniac says. "And the only solution to that is to do pretty much arbitrary gatekeeping."

The second cluster of worries involves what researchers outside of astronomy have termed "deskilling" or, more bleakly, "cognitive surrender." What if AI-dependent astrophysicists, especially young ones, lose or never build their own math, coding, and reasoning skills? "At that point, we've just completely selected for the disappearance of science in 50 years, because nobody will know how to do anything," Córdova Rosado said.

What AI enthusiasts dismiss as now-optional "grunt work" is exactly where graduate students develop the skills and intuition to do high-level science, Minas Karamanis, a cosmology postdoc at the University of California, Berkeley, argued in a March blog post. "Every hour you spend confused is an hour you spend building the infrastructure inside your own head that will eventually let you do original work," Karamanis wrote. Once a student crosses the hard-to-see line where the machine lets them skip thinking for themselves, "you haven't saved time. You've forfeited the experience that the time was supposed to give you."

NONE OF THE SCIENTISTS interviewed for this piece seemed confident projecting whether, in a few months or a year, their concerns will seem embarrassingly overblown or already painfully out of date. But in a field already threatened in the United States by sharp cuts to science funding and threats to the visas of foreign-born students, the situation for young scientists is "nightmarish," Vishniac told me. "Why is this happening in 2026?" Hogg said, rhetorically. "Can we postpone the AI revolution in astronomy?"

The issues are closest to home for the supposedly replaceable grad students. Poking my head into offices at the Center for Astrophysics, I offered to anonymize any grad student who wanted to participate in this article. Just one, a physicist, came forward. The student said their peers talk about these issues all the time. They're aware of Schwartz's "AI grad student" post, which, the student said, "implies that everyone sitting in that office is obsolete until they become a third-year grad student." But they were skeptical they could be so easily replaced.

And of course they, too, were using LLMs and agents. The grad student said their peers "are trying to [be] realistic: 'This is about to be our future, this is what our careers are going to be.'"

During the discussion in Garraffo's office, one of the most humanist, tech-skeptical voices had been Rafael Martínez-Galarza, AstroAI's deputy director. The next day, in his office, he showed me how he was using Claude to contribute to a code repository and to build a machine learning model tailored to solar physics. He had learned how to code and supervise students before AI arrived, he said, putting him in a "sweet spot" to profit from these tools without blindly trusting them.

"I didn't want to say this yesterday in the room because there were students there," he volunteered, "but I can see how a lot of things that I sit here to discuss with the students, I can just transform into a prompt and get the results much faster. But do I want that?"

External pressures to be more productive might push him toward an LLM, he admitted. And he did share some of the stated goals of AI engineers who want their models applied to astrophysics. He, too, deeply wanted to understand the universe and see humans exploring the Solar System. "I suspect a lot of why astronomy and cosmology are so catchy is because at the very end, it is really about meaning, right?" he said. "I see value in the process of matter turning into neurons trying to understand itself. I think it's beautiful—it's almost poetic."

But he conceives of that process, science, as a social activity, by and for humans. It would be more rewarding, he thought, to spend months solving a problem with a human than minutes with even the most brilliant chatbot. Even if answers to the cosmos could spill out in a sudden, oracular revelation, why not delight in getting to ask the universe our own questions and working through them together?

"I don't see why we would rush colonizing the Galaxy or even understanding the universe, if that means a step back in the human experience," he added. "I don't see the point." □

"I see value in the process of matter turning into neurons trying to understand itself. I think it's beautiful—it's almost poetic.

Rafael Martínez-Galarza
Center for Astrophysics
Harvard & Smithsonian

ILLUSTRATION: A. MASTIN/*SCIENCE*

Joshua Sokol is a freelance journalist in Raleigh, North Carolina. *Science*'s AI in Science Reporting Initiative is supported by Ray Rothrock & family.

The last astronomers

Joshua Sokol

Science **392** (6802), . DOI: 10.1126/science.aej3885

Science (ISSN 1095-9203) is published by the American Association for the Advancement of Science. 1200 New York Avenue NW, Washington, DC 20005. The title *Science* is a registered trademark of AAAS.