

# AI Reaches for the Stars

*Stella S. R. Offner*

*Artificial intelligence is rapidly being integrated into a variety of tasks in astronomical research, from model fitting to anomaly detection. While AI offers huge potential to accelerate discoveries about the universe, its current capabilities and trustworthiness fall short. I discuss the road to full AI integration into astronomical workflows: why the field of astronomy is a particularly fertile ground for AI development, the potential revolutionary gains possible for scientific analysis and discovery, how AI developments within astronomy can accelerate AI advancements more broadly, and the current roadblocks to fully AI-integrated astronomy.*

In recent years, progress in artificial intelligence has exploded, achieving previously unimaginable advances in image and language processing. Meanwhile, engagement with AI methods within scientific communities like astronomy has grown exponentially on two fronts. Over the past decade, formal use by researchers to apply AI methods to data analysis, including to classify objects, estimate parameters, and identify outliers, has increased tenfold.<sup>1</sup> Meanwhile, the recent rise of powerful public large language models (LLMs) such as OpenAI’s ChatGPT, Anthropic’s Claude, and Deepseek has sparked a quiet revolution in the adoption of AI to summarize research papers, draft and revise text, and write code. A 2024 survey of nearly five hundred astronomers by the American Astronomical Society found that 33 percent of respondents had used AI as a writing assistant and 48 percent had used AI to program.<sup>2</sup> AI is rapidly becoming not only common within but also an indispensable component of research workflows.

While embrace of AI by the full astronomy community remains somewhat distant, the future is clear. AI will evolve beyond being just another useful tool to become an intelligent assistant or “copilot” that is fully integrated in all aspects of the scientific process: hypothesis generation, information gathering via telescope observations or archival database queries, data processing and analysis, and data presentation in journals. The resulting AI boost to productivity and efficiency has the power to accelerate the research process; projects that currently take months or years will be completed in days.

My own journey into AI began in the early 2010s via an interest in image classification. Astronomy, as a fundamentally observational science, has a unique relationship with image processing. Preeminent telescopes like the Hubble Space Telescope (HST)

and Spitzer Space Telescope have generated huge mosaics of the sky, capturing the ethereal beauty of star clusters, galaxies, and nebulae. The Spitzer Galactic Plane Survey panned across the midline of our Milky Way Galaxy to produce an enormous image of 390,000 x 6,000 pixels, peppered with supernovae shells, wispy irradiated dust clouds, and winding dark lanes where stellar nurseries enshroud newborn stars.<sup>3</sup> These complex details are readily identified by eye, but classical automated pattern searches fare poorly, since there are many overlapping, faint features, and even bright sources are complex and heterogeneous. This astronomy data challenge inspired many citizen science projects, including Galaxy Zoo and the Milky Way Project, which, by the mid-2010s, had engaged hundreds of thousands of volunteers to visually identify and classify tens of millions of celestial objects.<sup>4</sup>

My own research investigates the complex energy signatures that young stars produce in the form of high-velocity jets, wind-blown bubbles, radiatively driven shells, and supernovae. These features, captured in high-definition images taken by the Hubble, Spitzer, and James Webb Space Telescopes, shape the stellar birth environment, including how efficiently stars form and when star formation stops. As a computational astrophysicist, my primary research instruments are supercomputers rather than telescopes. By the 2010s, numerical simulations of star formation had reached a volume, level of complexity, and accuracy to make ideal training sets for AI models.

Unlike human visual identifications, the outputs of computer simulations offer reliable, labeled data, while theoretically circumventing human biases. My foray into image analysis began with *random forests*, an old but robust approach to classification that is based on decision trees.<sup>5</sup> Although this method performed comparably to visual identification, it was not flexible enough to automatically identify complex stellar feedback features.<sup>6</sup> I turned to convolutional neural networks (CNNs), ultimately adopting an architecture able to identify structures via both spatial and wavelength correlations.<sup>7</sup> My research group developed the Convolutional Approach to Shell Identification in three dimensions (CASI-3D) to rapidly comb through large observational datasets and find stellar feedback at the level of individual voxels, contouring features to provide three-dimensional maps of their locations and properties.

Today, due to the complexity of astronomy data, many analyses continue to be conducted “by eye.” Galaxy Zoo is still going strong! However, a broad range of astronomy tasks, from outlier detection to model fitting, are being automated through AI methods.

**S**o what exactly does it mean to use AI in astronomy research? The rapid growth of AI capabilities and use have generated a plethora of definitions that reflect AI application in different contexts. Here, I define AI in the broadest sense: intelligent machines that learn to complete tasks without being explicitly programmed.<sup>8</sup> *AI machines* could refer to intelligent robots, which learn from their

environment; however, in the context of science, AI refers to intelligent *algorithms*, which are trained to perform tasks without being programmed with explicit rules. Just as ChatGPT learned to write poetry from reading large numbers of poems, sophisticated astronomy AI models trained on massive observational datasets learn to differentiate galaxies, stars, and accreting blackholes without being given instructions for how to identify these sources.

AI algorithms, like humans, have the ability to learn symbolic representations of concepts.<sup>9</sup> Some AI methods – namely, those architectures based on artificial neural networks, like CASI-3D – are direct analogs to the human brain, updating according to new input information and responding to positive or negative feedback (rewards or punishments) from the environment. AI methods also include LLMs, as well as vision transformers and diffusion models, like the algorithms underlying DALL-E, a commercial image generation model. Machine learning (ML), the subfield of AI commonly applied in the context of scientific data analysis and modeling, refers to algorithms trained to make data-driven decisions and predictions.<sup>10</sup> Common ML tasks include classification (as with random forests) and parameter estimation (regression).

**T**o understand the full potential for AI within astronomy today and why astronomy is a unique sandbox for AI requires an appreciation for the history of astronomical data and collaboration. The astronomy ecosystem is uniquely conducive to rapid AI development due to the standardization of practices that promote frictionless reproducibility: freely shared data, open-source code, and driving challenges – three conditions that accelerate data-driven technological advances.<sup>11</sup> This inclusive research environment was cultivated through decades of internationally coordinated efforts to digitize observational data, standardize data formats, and promote free open-source tools.

By the mid-1800s, a large number of astronomical observatories were operating around the world, recording images of the night sky on photographic plates. These plates, which were made with glass or metal and covered with a light-sensitive coating, could reach sizes of 30 cm x 30 cm; they were both a poorly transportable data record and owned by the observer, who had little incentive to loan these data to other researchers. Photographic film began to supplant plates in the 1950s, although data distribution and sharing remained challenging. In the 1960s, advancements in computing and the rapidly growing wealth of data incentivized efforts to digitize telescope observations.<sup>12</sup> Early preparations for the design and launch of the Hubble Space Telescope accelerated these efforts and, by the mid-1990s, astronomy had fully entered the digital era.<sup>13</sup> A variety of subsequent efforts have scanned and digitized photographic plates, further expanding the public archive of astronomical data.<sup>14</sup>

Despite the growing collection of digital data, access was also inhibited by heterogeneous data storage formats, the challenges of manipulating the stored data

without a shared software toolkit, and the need to combine observations across the electromagnetic spectrum from different facilities, each with bespoke data recording practices. In the 1970s, this state of affairs motivated the development of a standard data storage format. In 1982, the International Astronomical Union urged the adoption of the Flexible Image Transport System (FITS), a syntax for recording, sharing, and archiving data.<sup>15</sup> More than forty years later, FITS remains the default global format for storage of all astronomical data products.

In the absence of open-source software and public data repositories, the adoption of the FITS format was not sufficient to produce universal data sharing. Over the next two decades, community pressure and funding from NASA and the National Science Foundation (NSF) gradually promoted the creation of open-source community tools, leveling the field for data access and exploration. These developments were further accelerated by large astronomical survey projects, including those led by the HST, which sparked the creation of extensive digital archives or “virtual observatories” that further accelerated the democratization of astronomy data access. Every year since 2003, more than half of the papers published with HST data use archival data rather than new observations.<sup>16</sup>

In today’s competitive AI landscape, in which datasets are monetizable, collaboration and free exchange within astronomy are further facilitated by the low economic value of astronomical data. The late astronomy enthusiast and computer scientist James Gray joked that he “liked working with astronomers because their data is worthless.”<sup>17</sup> In other words, because astronomy data are free from the privacy, ethical, legal, and national security issues of other datasets, astronomy offers an AI playground. This culture of open data is reinforced by decades of national and international funding to support transnational astronomy infrastructure and research, which requires open data access and community sharing. Federal support has historically been buoyed by an enthusiastic public, inspired by the wonder of the night sky and universal appeal of the question: How did the universe come to be?

**C**ontemporaneous with the recent revolution in AI, progress in observational capabilities and computing has opened new frontiers for astronomical data analysis. Astronomy datasets are increasingly large, complex, and heterogeneous (“multimodal”), creating fresh barriers to data access and exploration. Emerging AI methods offer enormous potential to address inefficiencies and bottlenecks that touch all aspects of the astronomy research workflow.

## **Hypothesis Generation: Sparking New Research Directions**

The rapid advancement of LLMs promises a radically new approach to search and synthesize online knowledge, including scientific journal articles, catalogs of archival data, and even databases of future observing campaigns (such as abstracts

and specifications of accepted telescope proposals, which are publicly available in astronomy). Due to their extensive training data, general LLMs like GPT-5 can expertly answer high-level astronomy questions. In tandem, there is increasing interest in specialized domain-specific LLMs. Focused models tuned for narrower application are potentially computationally cheaper and have better-defined and more-transparent inputs. In 2023, the first open-source astronomy-domain LLM for literature review, AstroLLaMa, was trained on three hundred thousand paper abstracts.<sup>18</sup> While more prototype than assistant, AstroLLaMa was shortly followed by AstroLLaMa-2 and Pathfinder, which included larger training sets containing partial or full journal articles as well as various journal metadata.<sup>19</sup> These larger astronomy models perform significantly better on the search and synthesis of astronomy journal archives. To date, the performance of the models emerging from academic research lags behind that of state-of-the-art general, commercial models like GPT-4o (*o for omni*) and Claude.<sup>20</sup>

Nonetheless, value is added by the development of domain-specific features. Pathfinder is accompanied by a tool to visualize the astronomy research “landscape.” The topology of this literature map allows researchers to visualize topic relationships and quickly identify areas of missing work or emergent interest. Digestion of the full corpora of knowledge enables such models to suggest novel research directions.<sup>21</sup> However, this capability remains an active area of research, as the promise of LLM-suggested research directions still falls short compared with those proposed by expert researchers.<sup>22</sup> Despite the current deficits in LLM complex reasoning, there is cause for optimism: the steep upward trajectory of AI methodology suggests LLM ideation will soon equal or surpass human hypothesis generation, thereby advancing research frontiers in new and unexpected directions.

## Automating Telescope Operations: Collecting & Wrangling Big Data

Multinational, multibillion dollar observatories, like the recently launched James Webb Space Telescope, are the flagship research laboratories that drive astronomical discoveries. The capabilities and power of such facilities have been steadily accelerated by computing and technological innovations. Next-generation facilities, such as the ALMA (Atacama Large Millimeter/submillimeter Array) Wideband Sensitivity Upgrade, next-generation Very Large Array, Square Kilometre Array, and U.S./European Extremely Large Telescopes, will produce a firehose of rich data, far exceeding current processing capabilities.

Despite advancements in detector/instrument capabilities, the data-processing pipeline – from light detection to user delivery – is not fully automated. Data quality assurance, which includes vetting the impact of changing weather conditions and occurrence of instrumental artifacts, is substantially human-limited. The current reliance on human effort by domain experts will not scale with the data volume

expected from near- and long-term planned facilities without a paradigm shift in methodology.

Upon delivery to the researcher, the large volume (petabyte-scale) of data products hinders efficient scientific analysis, even in the context of archival research. For example, a single observation using the next-generation ALMA Wideband Sensitivity Upgrade detector will increase dataset sizes seventy-five-fold, recording light information across 1.2 million wavelength channels. In under a decade, observing projects for individual investigators will regularly generate *trillion* voxel terabyte-scale datasets, exceeding the capacity of typical PCs. Moreover, such cubes are often noisy and signal-sparse and thus challenging to parse for relevant features, especially for researchers without easy access to high-performance computing resources. Such data may have thousands of atomic and molecular features, many of which will be faint or have a structure that is not known *a priori*.<sup>23</sup> Emerging AI techniques for online analysis, outlier detection, and signal processing are critical to streamline data workflows by efficiently extracting salient features that indicate operational anomalies in the data collection or that are scientifically interesting for further study.

### **Accelerating Simulations: Fast Calculations of Astrophysical Systems**

Machine-learning methods offer enormous potential to accelerate the modeling of astrophysical systems. Supercomputers are the laboratories of astrophysics research. Since it is not possible to travel outside our solar system to collect samples or to recreate the extreme conditions of space in laboratories on Earth, supercomputers fill a vital experimental niche, enabling researchers to solve the equations of the universe and explore how different conditions and physical laws shape our reality.

Simulations running on the world's largest supercomputers have been used to calculate how the early universe expanded, model the evolution of galaxies, and study the formation of stars and planets.<sup>24</sup> Solving the fundamental equations of nature, however, is not cheap: these calculations typically require many millions of processor hours. Machine-learning methods offer several promising shortcuts to the traditional approach that directly solves the underlying physics equations. Neural networks may contain billions of tunable parameters, which provide the flexibility to describe complex physical systems. Given sufficient input training data, neural networks act as extremely high-dimensional fitting functions, able to represent complex, nonlinear dynamical behavior. Meanwhile, regression-based approaches aim to discover new, more-efficient expressions of the underlying physical laws, thereby allowing much faster calculations. AI methods can also be used to map the problem into a different mathematical space of lower or higher dimension, leading to a more efficient representation that is more computationally tractable.

AI acceleration is not without pitfalls, however. Accuracy depends on vast amounts of high-quality, labeled training data, which requires running large, very

expensive suites of simulations. This catch-22 only works in favor of AI applications if the simulations and data are available *a priori* and do not exceed the cost of calculating the solution directly. This points to training foundation models, which are sufficiently powerful to generalize beyond some narrow scope. In addition, many proof-of-concept studies demonstrate accuracy using idealized, low-dimensional problems (such as two-dimensional, smooth fluid flows and swinging pendulums). Architectures that perform well on simple benchmarks may not achieve the needed accuracy to model the high dynamic range and extreme conditions of astrophysical systems.<sup>25</sup> In addition, processes like star formation and galaxy evolution include phenomena that span timescales from days to billions of years. Classical techniques to circumvent short timescales, such as travel times imposed by the speed of light, produce significant spatial coupling across the problem domain, incurring large communication costs. Promising AI applications for astrophysical modeling involve replacing the most costly and modular components of direct simulations, like stellar evolution, gravitational dynamics, and chemical reactions.<sup>26</sup> Molecule formation in dense gas, or “astrochemistry,” is a prime example, since chemical reactions may occur in seconds, but the resulting distribution of species is a crucial input to the gas temperature, which coevolves with other processes that span thousands of years. The dynamics of few-body stellar/planetary interactions is another promising application, since the outcomes may be represented by statistical sampling, which is otherwise costly.<sup>27</sup>

Current research has only scratched the surface of the full potential for modeling astrophysical systems. Next-generation AI models able to generalize to a broad range of physical conditions – while performing with the fidelity of classical models – have the power to transform how we model the universe.

### **Supercharging Data Analysis: Fitting, Explaining & Predicting Physical Parameters**

Science is fundamentally grounded in statistics: all measurements have measurement errors. Consequently, comparing models to data requires a rigorous estimation of measurement confidence. The degree of the uncertainty in a given observation makes the difference between, say, an accelerating universe and a decelerating universe.<sup>28</sup> Testing and validating theoretical models requires a detailed model not only for a given parameter but for the possibility or *likelihood* of obtaining different results.

While the statistical frameworks concerned with estimating likelihoods predates AI/ML, AI offers fresh, efficient approaches for estimating measurement expectations. Estimating likelihoods for astronomical parameters is particularly challenging, since observations suffer from the fundamental limitation that physical properties – like density, temperature, or mass – cannot be directly measured but instead must be *inferred* from the characteristics of the observed light. Mapping

between observations and theoretical predictions requires assuming a physical model for the system that itself contains uncertainties. This “inverse” problem is common throughout science, but it is particularly vexing in astronomy, where physical scales span many orders of magnitude, spatial and velocity information are missing, evolution occurs on timescales dwarfing human lifetimes, and some of the key model parameters concern substances – dark energy and dark matter – that do not even emit light.

AI techniques combined with detailed simulations of physical systems, or “simulation-based inference,” allow complex likelihoods to be quantitatively estimated even when the underlying system properties are uncertain. Simulation-based inference techniques have been particularly powerful in driving progress in the field of cosmology, where AI methods combined with thousands of simulations of the universe and observational surveys encompassing millions of galaxies are able to tightly constrain the nature of dark matter and the expansion rate of the universe.<sup>29</sup> AI-based inferences thereby allow astronomers to confidently rule out sets of theoretical models more quickly and more robustly than with classical statistics.

AI successes in cosmology are partially owed to the relatively simple initial conditions of the early universe, which can be described by a small number of parameters that are (relatively) easy to sample and model. The situation becomes less tractable when the physical conditions of a given system are poorly constrained and no first-principles mathematical description exists, as in the subfields of star and planet formation. However, even in the simplest physical problems, the universe is complex; simulations, which are critical for AI inference, remain approximations of the physical world, or to quote the late statistician George Box, “all models are wrong.” As computing power continues to grow and simulations become even larger and more realistic, AI advances are needed to provide better metrics to quantify how simulations diverge from reality.

But even imperfect models advance our understanding of the true nature of the universe. Mismatches between models and measurements reveal our ignorance and misconceptions, with AI pointing the way to more accurate representations of our reality. In other words, while all models are wrong, even flawed models are useful.

**W**hile astronomy research clearly benefits from foundational advances in AI, the astronomy domain in turn offers enormous potential to spark AI innovation. Continued AI growth demands rich, accessible data, substantive novel challenges, and expert talent, all of which astronomy provides in abundance.

Astronomical data, unlike data in many scientific fields, are public and plentiful. Flagship observatories produce extremely high-dimensional and low signal-to-noise data, providing large and rich datasets able to power AI development. These astronomical data are inherently complex (multimodal), presenting a challenging application space that offers extensions to a broad range of other fields.

In contrast, public data that are siloed by facility, researcher, or application due to heterogenous storage conventions are not truly accessible. Improper or incomplete data documentation creates further impediments to use. Efficiently training AI models requires standardized data annotated with clear, complete metadata. The field of astronomy, having globally adopted a uniform data format (FITS) decades ago that encompasses multiwavelength observations from facilities around the world, is significantly ahead of the curve compared with other scientific disciplines.

In addition, the nonpropriety nature of astronomy data allows AI researchers to experiment and develop new capabilities without cumbersome legal, ethical, and financial restrictions. Because astronomy research occurs in an open-source, collaborative ecosystem, it provides unparalleled opportunities to develop and test AI methods in a “safe” public domain.

Beyond plentiful, rich, interoperable data, human capital is an essential fuel powering the AI revolution. AI advances depend critically on the availability of an extensive AI-skilled workforce. Astronomy, through the wonder of the night sky, is a natural gateway into science, technology, engineering, and math (STEM) fields. Astrophysics research, which engages a broad range of undergraduate students, graduate students, and post-PhD researchers, provides intensive training in wrangling, modeling, and predicting on large, complex datasets. Consequently, academia is an important pipeline of talent for the private sector; more than half of junior astronomy researchers ultimately transition to industry jobs.<sup>30</sup> The universal appeal of astronomy and its unique research ecosystem are a springboard to accelerate the training of the future AI workforce.<sup>31</sup>

In recognition of these synergies, in 2024, the National Science Foundation and Simons Foundation partnered to fund two five-year AI institutes focused on astronomy: the NSF-Simons AI Institute for Cosmic Origins (CosmicAI) and the National AI Institute for the Sky (SkAI). These two institutes expand the National AI Institutes program, which has funded twenty-nine institutes to date “to foster long-term, fundamental research in AI . . . and reinforce the foundation of technical leadership in AI.”<sup>32</sup> These new AI institutes will confront pressing challenges in astronomical research, while expanding the frontiers of AI capabilities and building infrastructure to train the next-generation workforce.

**D**espite the tremendous promise of AI approaches, the road to fully AI-integrated scientific workflows faces formidable obstacles. In addition to various technical challenges, progress is stymied by a lack of educational infrastructure, the inherent interdisciplinary nature of AI application, and a variety of sociological factors, including significant community skepticism, all of which have hampered AI adoption. For example, despite the widespread public embrace of LLMs as writing assistants, in a 2024 survey, 78 percent of astronomers reported being either somewhat or very concerned about using AI to write scientific papers,

while 50 percent were either somewhat or very concerned about using AI to analyze data.<sup>33</sup> This trepidation is rooted in (legitimate) apprehension about errors produced by AI hallucinations and the often inscrutable nature of AI predictions, concerns that are buoyed by continuing deficits in AI literacy, common among even the highly educated academic workforce.

Despite the rapid growth of AI within the public and private sectors, formal AI training has lagged behind the rising demand for expertise. Moreover, there is little advanced training within the context of AI in science.<sup>34</sup> Interested science students must seek out classes in other departments, such as in computer science, math, or data science. However, these courses are often difficult for nonmajor students to access, tend to focus on ML theory over application, and rarely include scientific, much less astrophysical, examples.

Developing AI educational resources and academic infrastructure to train a next-generation AI-skilled workforce is a central mandate of the NSF AI Institute program. Recent steps to expand scientifically relevant AI training include the NSF-Simons CosmicAI–sponsored Graduate Certificate of AI and ML (CAIML) program, which offers low-cost, online graduate courses in fundamental AI methods, including an astrophysical data application course.<sup>35</sup> Online programs like CAIML help bridge educational gaps while traditional academic undergraduate and graduate programs in the physical sciences more gradually incorporate domain-specific AI training within their standard curricula – a process that can take years.

Meanwhile, technological challenges remain at the interface of AI and scientific research. Black boxes are fundamentally anathema to science; AI adoption goes hand in hand with AI *trustworthiness*. While scientists are among the most trusted public figures, with 75 percent of international survey respondents trusting the scientific method, significant wariness about AI in science remains.<sup>36</sup> High-confidence scientific results have several key attributes, which are often missing in AI solutions. A trustworthy result is *reproducible*, meaning that it is consistently verifiable by independent observers. Scientific measurements must also be *interpretable*; *why* the outcome occurs must be predictable according to a series of logical arguments. Trust also requires *explainability*, meaning that the measurement is embedded within a broader, coherent narrative. For example, measurements of planetary motions around the Sun are *reproducible*. Kepler’s laws of planetary motion provide an *interpretable* framework for the planet positions, while the Newtonian laws of physics *explain* the behavior of planets in the context of the cosmos.

Among the sciences, sociological factors provide an additional confidence boost for astronomy research. Astronomy is an old science, accessible to any with the patience and curiosity to monitor the night sky by eye or with a small backyard telescope. Meanwhile, astronomy research is almost entirely federally funded, subject to little external financial pressure or corporate influence – another perk of such “worthless” data.

The rigorous mathematics underlying astrophysical theories provides a natural framework for reproducibility, interpretability, and explainability. While AI models also rest on a firm mathematical foundation of statistics and linear algebra, they trade simplicity for flexibility. AI models, like the complex workings of the human brain, violate Occam's razor: the simplest AI model is not necessarily the best one. Advances in mathematical theory are needed to provide explainability for high-parameter AI models, with implications for understanding the complex biology of human thought. In the meantime, while AI technologies hold vast promise to accelerate astronomical discoveries, social and cultural factors resist change; in other words, some of the fault is not in our stars but in ourselves.

**J**ust as the first programmable computers in the 1940s sparked a renaissance that ultimately revolutionized how we do science, AI advancements are rapidly broadening the scope of the possible. In the near future, astronomy researchers will be able to conceptualize, define, and execute research projects using trustworthy, efficient, interpretable, and robust AI methods. Astronomers will be able to interact with trustworthy LLMs that are aware of the full canon of astronomy data and research. Fast, robust AI methods will expedite data collection and modeling, while interpretable AI techniques will allow insightful and meaningful conclusions from the analysis.

Prior innovations in astronomical imaging have seeded practical advancements in diverse areas of society.<sup>37</sup> The development of adaptive optics systems to confront the distorting effects of turbulence in Earth's atmosphere led to improvements in retinal imaging and laser eye surgery.<sup>38</sup> Astronomers were the first to recognize the exciting scientific possibilities of charge-coupled devices (CCDs) – chips that are indispensable components of today's cell phones – and the drive to capture larger, higher-resolution images of the cosmos accelerated CCD technology.<sup>39</sup> AI advancements, sparked by our innate human desire to explore the cosmos, offer similar potential for technology transfer.

Despite this bright future, current astronomy AI applications primarily accelerate rather than revolutionize research. To date, there is no clear astronomy analog to the game-changing protein-folding AlphaFold AI model. On the horizon, foundation models trained on the massive next-generation observational and simulation datasets hold vast potential to induce a phase change in research workflows, not only sparking discoveries but changing how research is conducted. However, more work is needed to incorporate and provide uncertainties in model predictions, which are a critical aspect of advancing scientific progress. In the distant future, foundation model-enabled rapid source detection and AI-driven decision-making coupled to telescope infrastructure will be able to seamlessly integrate astronomical datasets, telescope operations, and computational data analysis. This future "AI astronomer" will sort through massive astronomical archives to form and explore novel hypotheses by peering into the depths of the universe through AI-optimized telescope eyes.

Any transformative innovation inevitably raises the question of human value: Will AI astronomers replace human astronomers? The history of technological progression suggests that AI will indeed change how new discoveries are made, demanding new skill sets and workflows. But human astronomers will remain essential. Future scientific progress will be catalyzed by increasingly integrated human-computer collaborations. AI will amplify humanity's reach toward the stars, copiloting the next generation of astronomical discoveries: galaxies forming during the cosmic dawn of the universe, the gravitational chirps of merging black holes, exotic worlds circling distant suns, and even, possibly, the holy grail of astronomical searches, the elusive signals of life beyond our solar system.

---

#### ABOUT THE AUTHOR

**Stella S. R. Offner** is Professor of Astronomy at the University of Texas at Austin. She is also the Co-Director of the Center for Scientific Machine Learning at the Oden Institute for Computational Engineering and Sciences and is the PI and Director of the NSF-Simons AI Institute for Cosmic Origins. She has recently published in such journals as *Nature Astronomy*, *The Astrophysical Journal*, and *Astronomy & Astrophysics*.

#### ENDNOTES

- <sup>1</sup> Daniela Huppenkothen, Michelle Ntampaka, Matthew Ho, et al., "Constructing Impactful Machine Learning Research for Astronomy: Best Practices for Researchers and Reviewers," arXiv (2023), Figure 1, <https://doi.org/10.48550/arXiv.2310.12528>.
- <sup>2</sup> L. Prato and B. Gaensler, private communication, 2025.
- <sup>3</sup> "GLIMPSE (Galactic Legacy Infrared Mid-Plane Survey Extraordinaire)," Spitzer Space Telescope, <http://ftp.astro.wisc.edu/glimpse> (accessed December 15, 2025).
- <sup>4</sup> Galaxy Zoo, <https://www.zooniverse.org/projects/zookeeper/galaxy-zoo>; Milky Way Project, <https://www.zooniverse.org/projects/povich/milky-way-project>; and Chris J. Lintott, Kevin Schawinski, Anže Slosar, et al., "Galaxy Zoo: Morphologies Derived from Visual Inspection of Galaxies from the Sloan Digital Sky Survey," *Monthly Notices of the Royal Astronomical Society* 389 (3) (2008): 1179–1189, <https://doi.org/10.1111/j.1365-2966.2008.13689.x>.
- <sup>5</sup> Tin Kam Ho, "Random Decision Forests," in *Proceedings of 3rd International Conference on Document Analysis and Recognition* (IEEE, 1995), 278–282, <https://doi.org/10.1109/ICDAR.1995.598994>.
- <sup>6</sup> Duo Xu and Stella S. R. Offner, "Assessing the Performance of a Machine Learning Algorithm in Identifying Bubbles in Dust Emission," *The Astrophysical Journal* 851 (2) (2017), <https://doi.org/10.3847/1538-4357/aa9a42>.

- <sup>7</sup> Yann LeCun, Yoshua Bengio, and Geoffrey Hinton, “Deep Learning,” *Nature* 521 (7553) (2015): 436–444; Colin M. Van Oort, Duo Xu, Stella S. R. Offner, and Robert A. Gutermuth, “CASI: A Convolutional Neural Network Approach for Shell Identification,” *The Astrophysical Journal* 880 (2) (2019): <https://doi.org/10.3847/1538-4357/ab275e>; and Duo Xu, Stella S. R. Offner, Robert Gutermuth, and Colin Van Oort, “Application of Convolutional Neural Networks to Identify Stellar Feedback Bubbles in CO Emission,” *The Astrophysical Journal* 890 (1) (2020), <https://doi.org/10.3847/1538-4357/ab6607>.
- <sup>8</sup> John McCarthy, Marvin L. Minsky, Nathaniel Rochester, and Claude E. Shannon, “A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence, August 31, 1955,” *AI Magazine* 27 (4) (2006), <https://doi.org/10.1609/aimag.v27i4.1904>.
- <sup>9</sup> Sergios Theodoridis, *Machine Learning: A Bayesian and Optimization Perspective*, 2nd ed. (ScienceDirect, 2020).
- <sup>10</sup> Viviana Acquaviva, *Machine Learning for Physics and Astronomy* (Princeton University Press, 2023).
- <sup>11</sup> David Donoho, “Data Science at the Singularity,” arXiv (2023), <https://doi.org/10.48550/arXiv.2310.00865>.
- <sup>12</sup> Edward Kibblewhite, “Automated Measurement of Astronomical Photographs” (PhD diss., University of Cambridge, 1971).
- <sup>13</sup> W. Patrick McCray, “The Biggest Data of All: Making and Sharing a Digital Universe,” *Osiris* 32 (1) (2017), <https://doi.org/10.1086/693912>.
- <sup>14</sup> The Harvard Plate Stacks, “Our Collection,” <https://platestacks.cfa.harvard.edu/our-collection> (accessed December 16, 2025); and Harry Enke, Taavi Tuvikene, Detlef Groote, et al., “Archives of Photographic PLates for Astronomical USE (APPLAUSE): Digitisation of Astronomical Plates and Their Integration into the International Virtual Observatory,” *Astronomy & Astrophysics* 687 (2024), <https://doi.org/10.1051/0004-6361/202348793>.
- <sup>15</sup> McCray, “The Biggest Data of All,” citation therein.
- <sup>16</sup> Barbara A. Mikulski Archive for Space Telescopes, “HST Reports: Paper Counts,” <http://archive.stsci.edu/hst/bibliography/pubstat.html> (accessed December 16, 2025).
- <sup>17</sup> Alexander S. Szalay, “Jim Gray, Astronomer,” *Communications of the ACM* 51 (11) (2008): 58–65, <https://cacm.acm.org/research/jim-gray-astronomer>.
- <sup>18</sup> Tuan Dung Nguyen, Yuan-Sen Ting, Ioana Ciucă, et al., “AstroLLaMA: Towards Specialized Foundation Models in Astronomy,” arXiv (2023), <https://doi.org/10.48550/arXiv.2309.06126>.
- <sup>19</sup> Rui Pan, Tuan Dung Nguyen, Hardik Arora, et al., “AstroMLab 2: AstroLLaMA-2-70B Model and Benchmarking Specialised LLMs for Astronomy,” arXiv (2024), <https://doi.org/10.48550/arXiv.2409.19750>; and Kartheik G. Iyer, Mikael Yunus, Charles O’Neill, et al., “pathfinder: A Semantic Framework for Literature Review and Knowledge Discovery in Astronomy,” *The Astrophysical Journal Supplement* 275 (2) (2024), <https://doi.org/10.3847/1538-4365/ad7c43>.
- <sup>20</sup> Sebastian Antony Joseph, Syed Murtaza Husain, Stella S. R. Offner, et al., “AstroVisBench: A Code Benchmark for Scientific Computing and Visualization in Astronomy,” arXiv (2025), <https://doi.org/10.48550/arXiv.2505.20538>.
- <sup>21</sup> Yangqiaoyu Zhou, Haokun Liu, Tejes Srivastava, et al., “Hypothesis Generation with Large Language Models,” arXiv (2024), <https://arxiv.org/abs/2404.04326>.

- <sup>22</sup> Chenglei Si, Tatsunori Hashimoto, and Diyi Yang, “The Ideation-Execution Gap: Execution Outcomes of LLM-Generated versus Human Research Ideas,” arXiv (2025), <https://doi.org/10.48550/arXiv.2506.20803>.
- <sup>23</sup> Ryan A. Loomis, Karin I. Öberg, Sean M. Andrews, et al., “Detecting Weak Spectral Lines in Interferometric Data through Matched Filtering,” *The Astronomical Journal* 155 (4) (2018), <https://doi.org/10.3847/1538-3881/aab604>.
- <sup>24</sup> Jonah C. Rose, Paul Torrey, Francisco Villaescusa-Navarro, et al., “Introducing the DREAMS Project: DaRk mattEr and Astrophysics with Machine Learning and Simulations,” *The Astrophysical Journal* 982 (2) (2025), <https://doi.org/10.3847/1538-4357/adb8e5>; Philip F. Hopkins, Michael Y. Grudic, Kung-Yi Su, et al., “FORGE’d in FIRE: Resolving the End of Star Formation and Structure of AGN Accretion Disks from Cosmological Initial Conditions,” *The Open Journal of Astrophysics* 7 (2024), <https://doi.org/10.21105/astro.2309.13115>; and Dávid Guszejnov, Michael Y Grudić, Stella S. R. Offner, et al., “Effects of the Environment and Feedback Physics on the Initial Mass Function of Stars in the STARFORGE Simulations,” *Monthly Notices of the Royal Astronomical Society* 515 (4) (2022), <https://doi.org/10.1093/mnras/stac2060>.
- <sup>25</sup> L. Prato and B. Gaensler, private communication, 2025.
- <sup>26</sup> Guangping Li, Zujia Lu, Junzhi Wang, and Zhao Wang, “Machine Learning in Stellar Astronomy: Progress up to 2024,” arXiv (2025), <https://doi.org/10.48550/arXiv.2502.15300>; and Aldana Grichener, Mathieu Renzo, Wolfgang E. Kerzendorf, et al., “Nuclear Neural Networks: Emulating Late Burning Stages in Core-Collapse Supernova Progenitors,” *The Astrophysical Journal Supplement* 279 (2) (2025), <https://doi.org/10.3847/1538-4365/ade717>.
- <sup>27</sup> Philip G. Breen, Christopher N. Foley, Tjarda Boekholt, and Simon Portegies Zwart, “Newton Versus the Machine: Solving the Chaotic Three-Body Problem Using Deep Neural Networks,” *Monthly Notices of the Royal Astronomical Society* 494 (2) (2020): 2465–2470, <https://doi.org/10.1093/mnras/staa713>.
- <sup>28</sup> Adrian G. Adame, Janette Aguilar, Steven Ahlen, et al., “DESI 2024 VII: Cosmological Constraints from the Full-Shape Modeling of Clustering Measurements,” *Journal of Cosmology and Astroparticle Physics* 2025 (2025), <https://doi.org/10.1088/1475-7516/2025/07/028>.
- <sup>29</sup> ChangHoon Hahn, Pablo Lemos, Liam Parker, et al., “Cosmological Constraints from Non-Gaussian and Nonlinear Galaxy Clustering Using the SimBIG Inference Framework,” *Nature Astronomy* 8 (2024): 1457–1467, <https://www.nature.com/articles/s41550-024-02344-2>.
- <sup>30</sup> Patrick Mulvey and Jack Pold, “New Astronomy PhDs: What Comes Next,” American Institute of Physics, January 25, 2023, <https://www.aip.org/statistics/new-astronomy-phds-what-comes-next>.
- <sup>31</sup> Keith Poletti, Stella S. R. Offner, Rachel A. Ward, “Modeling Turbulent and Self-Gravitating Fluids with Fourier Neural Operators,” arXiv (2025), <https://doi.org/10.48550/arXiv.2507.23662>.
- <sup>32</sup> U.S. National Science Foundation, “National Artificial Intelligence Research Institutes,” <https://www.nsf.gov/funding/opportunities/national-artificial-intelligence-research-institutes> (accessed December 16, 2025).
- <sup>33</sup> Andrew Ferguson, Marisa LaFleur, Lars Ruthotto, et al., “The Future of Artificial Intelligence and the Mathematical and Physical Sciences (AI+MPS),” arXiv (2025), <https://doi.org/10.48550/arXiv.2509.02661>.

<sup>34</sup> Ibid.

<sup>35</sup> Computer & Data Science Online, The University of Texas at Austin, “Graduate Certificate in Artificial Intelligence and Machine Learning,” <https://cdso.utexas.edu/caiml> (accessed December 15, 2025).

<sup>36</sup> Viktoria Cologna, Niels G. Mede, Sebastian Berger, et al., “Trust in Scientists and Their Role in Society across 68 Countries,” *Nature Human Behaviour* 9 (2025): 713–730, <https://doi.org/10.1038/s41562-024-02090-5>.

<sup>37</sup> Joyful E. Mdhului, Charles Takalana, Ramasamy Venugopal, et al., “Astronomy as a Strategic Driver for Sustainable Development,” *Nature Astronomy* 9 (2025): 936–939, <https://www.nature.com/articles/s41550-025-02602-x>.

<sup>38</sup> Engin Akyol, Ahmed M. Hagag, Sobha Sivaprasad, and Andrew J. Lotery, “Adaptive Optics: Principles and Applications in Ophthalmology,” *Eye* 35 (1) (2020): 244–264, <https://doi.org/10.1038/s41433-020-01286-z>.

<sup>39</sup> NOIRLab, “The Charge-Coupled Device: Revolutionizing How Astronomers See the Universe after 50 Years of Development: What’s on the Horizon?” April 22, 2024, <https://noirlab.edu/public/blog/50-years-ccds>.