

How Do We Build AI to Push the Frontiers of Scientific Discovery?

Anima Anandkumar

Scientific progress is limited not by a lack of new ideas but by the time and cost involved in physical experimentation. Scientific discovery is a needle in the haystack problem: it does not help if AI gives you a vastly bigger haystack. Without knowing if any of the ideas work, an AI system that designs experiments just increases the effort required, since performing the experiments to validate the ideas is the real bottleneck. In my view, AI's most transformative impact in enabling scientific discoveries lies in reducing the need for such experiments. To get there, we need to build AI models that are able to granularly simulate and understand physics at all scales, rather than just abstractly reason in the textual domain. In this essay, I explore what methods like Neural Operators have already helped achieve, what still needs to be done, and what lies ahead.

Can AI generate truly novel and surprising scientific discoveries and inventions? My research has focused on this challenge for almost a decade. Before my arrival at Caltech, I led AI research at Amazon Web Services, building some of the first AI products on the cloud. Many of my colleagues at Caltech had a healthy dose of skepticism about AI and whether it could have a lasting impact in scientific domains, where data are nowhere as plentiful as images and text on the internet. This challenge of developing AI algorithms for scientific advances became my research focus both at Caltech and during my time at NVIDIA.

Fast forward to this era of large language models, and AI is making significant strides in scientific domains. Language models have achieved gold medal-level performance in major math contests. They have fundamentally changed how many people write software code today. There have also been extensive efforts to use them for scientific discovery. For instance, “AI scientists” have shown some success in specific cases, such as predicting chemical synthesis and drug repurposing.¹ These AI scientists – systems built on language models, sometimes with some additional domain-specific fine-tuning – attempt to generate new ideas and hypotheses based on their textual knowledge. And while such language models may be part of the process of scientific discovery, they are not the complete answer.

Textual models do not address the critical bottleneck of scientific discovery. Scientific progress is not limited by a lack of new ideas but by the cost and time

needed to test them through physical data acquisition and validation in lab experiments. Scientific discovery is a needle in the haystack problem: it does not help if AI gives you a vastly bigger haystack. Without knowing if any of the ideas work through experimental validation, the volume just increases effort. AI's most transformative potential for enabling scientific discoveries therefore lies in reducing the need for such experiments.

To reach this potential, the hypotheses proposed by AI need to have some level of certainty of succeeding in the physical world. This is unlikely to come solely from textual data. For AI to have physical understanding, we need to train it on data from the physical world. There have been many recent success stories of training AI on domain-specific data, with outcomes such as discovering new functional enzymes and predicting coronavirus variants of concern, protein dynamics, and extreme weather events.² One common theme in all these examples is the availability of sufficiently large (curated) datasets. What about areas in which such datasets are unavailable or too difficult to collect? Moreover, even in the above areas, can we go beyond the regime of training data? This is not just a theoretical question. For example, if we are trying to predict extreme weather events in the context of significant climate change, such data are not available.³ In protein folding, many important protein families are not well represented in the current training data.⁴ Further, many bold claims of AI-driven discovery have seen pushback. For instance, the revelation that AI had discovered “2.2 million structures . . . which escaped previous human chemical intuition” was tempered by experiments showing that many of the AI-discovered materials did not function as predicted.⁵

Discovery is about going beyond what is known; it cannot just rely on examples. Standard data-driven AI struggles to extrapolate beyond the training distribution and is typically overconfident and miscalibrated. This implies that purely data-driven AI for scientific discovery has significant limitations, and initial successes may be attributed to either pure luck or being close enough to the training data, making it an interpolation rather than an extrapolation problem.

For AI to make truly novel scientific discoveries, AI models need a broader physical understanding beyond a narrow data distribution. This can come from the integration of a variety of sources and domain constraints – such as the governing laws of physics and symmetries – with data-driven AI.⁶ In addition to improving extrapolation capabilities, the knowledge of physics allows for verification of AI-discoveries, telling us what works in the real world and minimizing or even eliminating the need for physical prototyping. For instance, many physical phenomena are described by partial differential equations (PDE) such as fluid dynamics, plasma evolution, and quantum mechanics. We can train AI models to follow these guardrails of physical laws and domain constraints, making their discoveries more trustworthy.

To build AI that understands the physical world, we need to consider the governing processes that may span from subatomic to planetary scales and involve spatio-temporal evolution of processes across different scales and domains. Interactions between Earth's complex physical processes span more than twelve orders of magnitude in space and time, from micrometers at the molecular scale to thousands of kilometers at the planetary scale.⁷ Analyzing how coronavirus spreads in respiratory aerosol, for example, spans from atomic-scale electronic structure to whole aerosol particle morphology.⁸ For modeling such physical phenomena, we need AI that can capture fine details accurately without being computationally intractable. Standard neural networks such as transformers and convolutional neural networks are ill-equipped to capture multiscale phenomena since they assume fixed-sized inputs and outputs. While such neural networks can learn on discrete distributions such as text or images at a fixed resolution, they are not sufficient to handle multiple scales of interaction in scientific domains. This limits us from using data sources that are available at multiple resolutions and varying levels of fidelity – for instance, weather and climate models run at different resolutions – and limits us from imposing the physics constraints of governing laws at multiple scales.

I led the team of researchers that invented Neural Operators as a universal framework for learning mappings between function spaces.⁹ The relationship between neural networks and Neural Operators mirrors raster versus vector graphics: the former uses a fixed number of pixels and gets blurry when zoomed in, while the latter uses functional representation that retains sharpness even when zoomed in. Neural Operators learn targets that are continuous functions and that can be discretized to any resolution, while achieving super-resolution that goes beyond training data. They can learn from data at multiple scales and incorporate the knowledge of mathematical equations to fill in the finer details when only limited-resolution data are available. Neural Operator-based models can seamlessly assimilate experimental and observational data and adapt to new information. Moreover, since AI models are differentiable, they enable inverse design, supporting exploration and optimization under diverse design and manufacturing constraints.

Neural Operators have already achieved success in many scientific domains. I am proud to have led an interdisciplinary team from NVIDIA, Caltech, Berkeley, and other universities that trained FourCastNet, the first high-resolution, fully AI-based weather model.¹⁰ FourCastNet has competitive accuracy while being tens of thousands of times faster than traditional weather models. This surprised many weather scientists, given the general consensus that AI was still in its infancy and that a number of fundamental breakthroughs would be needed before deep learning could be competitive with traditional weather models, which could take years or even decades.¹¹ FourCastNet led Huawei and Google DeepMind,

among many others, to pursue AI-weather models and for weather agencies such as the European Centre for Medium-Range Weather Forecasts to host AI-weather models alongside traditional forecasts, allowing the community to compare the effectiveness of AI models in action.¹²

FourCastNet also correctly predicted many extreme weather events, including Hurricane Lee and Beryl, several days before traditional weather models.¹³ This was surprising since there are not many instances of hurricanes and other extreme events in the training data. How can AI weather models do so well even with rare events? One possible explanation is that hurricanes carry the distinctive physical signature of warm ocean air rising into the storm, which could be easily learned with only a few samples.¹⁴ It is noteworthy that the overall training of FourCastNet used only about fifty thousand samples, which is minuscule compared with the web-scale training of language models. In other domains, we have seen similar instances of physical phenomena that were thought to be prohibitively complex but do not actually require an inordinate amount of training data.¹⁵ But access to *high-quality* curated data is another important common feature among these success stories. In the case of weather, the open dataset ERA5 consists of roughly forty years of reanalysis weather data, which combine raw satellite observations with physics. In this sense, the AI-weather models are already physics-informed, and thus have better generalization abilities, which makes them practically useful.

The claims of AI improving our ability to predict extreme weather events must be carefully tested, and a few success stories with recent hurricanes do not suffice to make a definitive scientific claim. How can models prove themselves? Extreme weather is a chaotic system. Popularly known as the “butterfly effect,” small changes in initial conditions of a system can lead to rapid divergence in trajectories and radically different outcomes. This means that an already-chaotic system like extreme weather is fundamentally unpredictable beyond a certain time horizon using deterministic forecasts.¹⁶ Prediction models can overcome this unpredictability either by incorporating probabilistic outputs, as with diffusion models, or by considering statistical ensembles over initial conditions. The latter is more lightweight, and due to the speed of deterministic AI-weather models such as FourCastNet, we can support large ensembles that are out of reach for traditional weather models. Further, since we can test for physical validity of each ensemble member and filter out the nonphysical ones, we can better align with ground-truth physics. We train on such ensembles to accurately capture the tail event probability in the latest iteration of training for FourCastNet (version three).¹⁷ This achieves calibrated probabilistic performance on extreme weather with speed and accuracy. Using physically realistic large-ensemble weather forecasts up to sixty days ahead, FourCastNet3 surpasses traditional and AI-weather models in medium-range skill, speed, and stability, while providing trustworthy predictions. This has immense promise for early-warning systems for catastrophic weather events worldwide.

For an AI model to make a broader impact, it needs to have a broader physical understanding, as opposed to being just a narrow surrogate limited to its training regime. Surprisingly, FourCastNet functions not only as a weather model for the short to medium term but is also impactful for subseasonal and seasonal modeling and even long-term climate modeling.¹⁸ Given that the pretraining of FourCastNet involves one-step prediction, typically over six hours, with some additional multistep fine-tuning over a small number of steps in autoregressive rollouts, it is surprising that FourCastNet is able to extrapolate to much longer timescales of several months to years. This capability does not exist in other AI-weather models, which experience blowups over long rollouts. This long-term stability can be attributed to the Spherical Fourier Neural Operator architecture employed in FourCastNet.¹⁹ This design incorporates multiscale learning into an equivariant neural architecture that enforces spherical geometry (is equivariant to rotations on the sphere). In contrast, other AI-weather models use standard neural networks, such as transformers or graph neural networks, which assume a rectangular domain, leading to buildup of distortions near the poles and eventually to catastrophic blowups. Geometry-informed AI has been more robust for extrapolation in other areas as well, such as in molecular and protein modeling.²⁰ This demonstrates the importance of AI being informed by domain geometry and symmetries, and incorporating them into the neural architectures to achieve broad generalization.

Careful design of FourCastNet has given it broad capabilities, earning it the label of foundation model. But the extrapolation capabilities of such models are not boundless. Some AI weather models are unable to forecast “gray swans,” the extreme weather events that are so rare they do not exist in the training set.²¹ Yet such events are possibilities in the context of climate change, despite the absence of observational data.²² Similarly, in other areas such as protein folding, the predictions degrade with proteins outside the training set.²³ As discussed earlier, true extrapolation may not occur using purely data-driven methods. We need physics-informed AI to push those boundaries further.

While the machine-learning community has focused on data-driven approaches, the applied math and computational sciences community has prioritized modeling physical phenomena through partial differential equations and other mathematical equations, and applying numerical methods for solving them. Compared with machine-learning methods that can be considered top-down, mathematical modeling is bottom-up and provides a mechanistic explanation of the observed phenomenon. PDEs are popular in mathematical modeling because they relate across space and time and can model diverse phenomena such as fluid dynamics, material deformation, and quantum chemistry. However, numerical solvers for PDEs are computationally expensive since they require a fine grid to capture the multiscale nature of dynamics of many phenom-

ena. Thus, many real-world phenomena are intractable to simulate accurately even though their governing laws can be written down as succinct PDEs. For instance, Schrodinger's equation models the behavior of all matter, but to solve it exactly using current numerical methods for even a small molecule of about one hundred electrons would take more than the age of the universe. As another example, today's global climate models are unable to resolve clouds and other fine-scale phenomena, which are the biggest source of uncertainty for climate change predictions. Resolving them with current numerical methods is estimated to require several orders of magnitude more computing than is available today.²⁴

Given such impossible computational requirements for exact numerical solutions of PDEs, many approximations have been developed, including density functional theory (DFT) for quantum chemistry and large eddy simulations (LES) for fluid dynamics.²⁵ However, deep domain expertise may be needed to choose the appropriate level of approximation for a specific problem, and such approximations still may not be optimal. In many scenarios, even the approximate numerical solutions are still too expensive to be practical, such as for DFT on protein dynamics, among other simplifications that may not be well principled. Can AI come up with novel approximations that are both significantly cheaper and accurate? The revolution in computer vision and image processing has already shown this is possible: while methods prior to deep learning that used hand-engineered features struggled to reach good accuracy, feature learning with neural networks has been far more successful. Similarly, Neural Operators can learn faster solvers for PDEs compared with human-engineered numerical methods. Additionally, they can incorporate experimental and observed data to reduce modeling errors. We already see this with the Neural Operator-based weather model described above, which is tens of thousands of times faster than physics-based forecasts that solve PDEs, and which also surpasses them in accuracy since it can overcome modeling errors.

An alternative approach to data-driven methods for solving PDEs is Physics-Informed Neural Networks (PINNs), which attempt to solve PDEs in a data-free manner by imposing the PDE residual as the loss function.²⁶ PINNs can be viewed as test-time optimization, and the neural network serves as the *ansatz* for optimization. However, for many time-varying practical problems, the optimization landscape is often intractable, and even when they converge, PINNs are typically not faster than numerical solvers.²⁷ In contrast, we can train Neural Operators to learn the solution operator of PDEs via supervised learning by generating solutions to PDEs using existing numerical simulations. Such a data-driven approach overcomes the optimization challenges of PINNs. However, it is often expensive to generate sufficient high-quality data to train AI. A more-nuanced approach could combine simulation data from a variety of numerical solvers at different levels of approximation as well as incorporate the knowledge of physics, either as PDE residual loss function or through symmetry and conservation laws. Neural Operators are ideal

for this, since the data from numerical solvers at different levels of approximation are typically at different resolutions; for example, the direct-numerical simulation in fluid dynamics is on a finer grid than is large-eddy simulation, and the PDE residual can be imposed as a loss function at different resolutions.

Neural Operators present a principled AI approach for solving PDEs and modeling processes on a continuum, since they are not restricted to one resolution or scale.²⁸ They learn operators that are mappings between function spaces. They are discretization-invariant, meaning they work on any discretization of inputs without retraining. Neural Operators are universal approximators in function spaces: they can fit any continuous operator, such as solutions of PDEs at all resolutions.²⁹ This grid-free approach of Neural Operators allows us to seamlessly train on data at multiple resolutions. We can also impose physics losses, such as PDE residuals, at a higher resolution than the training data. Such a physics-informed Neural Operator approach can potentially surpass the accuracy of the training data.³⁰ Neural Operators can therefore holistically combine data and physics at multiple scales, harmonizing the bottom-up physics-based modeling and top-down data-driven learning.

The development of Neural Operators comes from first principles. Let us start with the simple setting of linear PDEs, for which the solution operator can be expressed as a linear integral operator with a given kernel, also known as Green's function. These kernel integrations can be learned using graphs, Fourier transforms, or attention mechanisms. We can then extend this idea for learning general nonlinear operators by composing learnable linear integral operators with nonlinear activations. This general framework of Neural Operators can universally approximate any nonlinear continuous operator, including solutions to nonlinear PDEs. Many popular neural-network architectures can be converted into Neural Operators that work at any resolution – such as the Fourier Neural Operator, Graph Neural Operator, and Transformer Neural Operator – yielding a rich family of models that are expressive and carry good inductive biases relevant to scientific domains.³¹

In many areas, Neural Operators have already replaced numerical solvers at rates that are four to six orders of magnitude faster. In addition to weather forecasting, this approach has been transformative in other areas such as modeling plasma evolution in nuclear fusion for detecting disruptions and modeling carbon dioxide storage in reservoirs for climate change mitigation, as well as detecting deviations from assumed physics, such as tipping points in climate change that present significant deviation from current conditions.³² Further, Neural Operators enable inverse design since they are differentiable and can both provide guidance for iterative design improvements and generate new designs from scratch. For instance, Neural Operators optimally designed a novel medical catheter that reduces bacterial contamination by two orders of magnitude, addressing one of the most common

sources of health care–related infections.³³ Neural Operators accurately simulated the bacterial density in fluid flow and designed shapes within the catheter to prevent bacteria from swimming upstream into the human body. The AI-designed catheter required only one set of physical experiments for validation and removed the need for the laborious trial-and-error loop typical of design optimization. In inverse lithography, Neural Operators produced better mask quality while being six times faster through progressive self-training, which is not feasible with a traditional solver.³⁴ In lung ultrasound imaging, Neural Operators can work on raw radio frequency signals directly and be physics-informed, yielding much better accuracy in lung aeration levels, as opposed to the current practice of subjective B-mode evaluation.³⁵ Diffusion Neural Operators can learn better priors from data when solving inverse problems.³⁶ Neural Operators are also helpful in deriving efficient control policies through reinforcement learning on hard problems such as drag reduction in turbulent systems.³⁷ Thus, the approach of using physically valid AI-generated designs, discoveries, and control policies using Neural Operators has been transformative in many areas.

Neural Operators also offer a foundation for novel scientific discovery involving multiscale phenomena. Combining this with symbolic reasoning can further boost capabilities for extrapolation. For instance, proving theorems about phenomena such as fluid dynamics would entail symbolic reasoning. PINNs have helped make progress on the famous Millennium Prize problem in fluid dynamics, but the symbolic analysis has so far been done by humans. Can we build AI that can do both?³⁸ There has been exciting progress in theorem-proving using AI, with AI systems turning in gold-level performances at the International Math Olympiad contest for high school students. Several of the winning AI systems made use of verifiers like Lean, a formal proof system that can check every proof step proposed by language models. This ensures hallucination-free reasoning, which is essential for AI to prove long theorems: in the absence of verifiers, errors can build up catastrophically over long reasoning chains. I led a team of researchers in the development and release of the first open-source framework combining language models with Lean – LeanDojo – which has ushered in this revolution in math AI.³⁹ This progress in formal reasoning is complementary to physical modeling with Neural Operators, as it allows us to refine PDEs and their numerical approximations both for data generation and physics losses to train Physics-Informed Neural Operators, while guaranteeing their correctness.

Neural Operators can be thought of as an end-to-end approach to using AI in scientific research. Conversely, more-traditional approaches have retained numerical workflows with traditional solvers and augmented them with AI, typically as an additive correction term. Popular in domains such as fluid dynamics, plasma physics, and climate modeling, this is known as “closure

modeling.”⁴⁰ More generally, in areas like molecular modeling, it is known as coarse graining and renormalization. This process, which predates the use of machine learning, corrects approximate coarse-grid numerical solvers using closures built from physics, such as large eddy simulation for fluid dynamics.⁴¹ Thus, there are two paradigms for physics-informed learning: replace with AI or augment with AI. The former incorporates physics constraints into end-to-end learning while the latter augments existing numerical solvers with learned closures. Is there a clear winner?

In terms of computational efficiency, end-to-end learning with Neural Operators offers significant speed gains over traditional solvers (four to six orders of magnitude), along with differentiability for inverse problems. On the other hand, with closure modeling, we typically do not get the full benefits of faster models and, often, the orchestration of numerical solvers with machine learning leads to slowdowns rather than speedups. Closure-model systems are also typically non-differentiable and therefore not suitable for efficiently solving inverse problems through gradient descent.

With regard to training data requirements, one would expect closure models to be more data efficient, since they would not need to learn from scratch due to augmentation with coarse-grid numerical solvers. We proved a surprising result to the contrary: learning-based closure models fare poorly in both data efficiency and training stability compared with end-to-end learning. Because closure models only have access to training data on a coarse grid, the learning closures are an ill-posed problem. Further, corrections made by the closure model are only allowed additively on the same coarse grid as the numerical solver, which limits expressivity and does not allow nonlinear mixing between scales. We proved that Neural Operators not only circumvent the limitations of standard closure models but also achieve statistically optimal estimation of long-term statistics of chaotic systems.⁴² To achieve data efficiency, we progressively trained the Neural Operator at multiple resolutions: first with coarse-grid solver data, which is cheaper to obtain, then by fine-tuning with a small amount of fine-grid simulations and physics constraints. Viewed under this lens, the Neural Operator approach can be seen as a generalization of the closure model, removing its limitations, in which coarse graining is possible through nonlinear mixing across scales and learning can happen progressively using multiresolution, multifidelity training samples and physics knowledge.

The rapid advances in AI occurring today *can* enable novel discoveries, but not without strong algorithmic foundations. Discovery is a journey into the unknown and cannot rely solely on AI trained on known samples; infusing physics into AI is essential. Since our world has physical interactions at multiple scales, we need AI that can traverse resolutions and scales. Neural Operators offer a principled foundation for extending standard neural networks to function spaces to enable multiscale physics-informed learning. A foundation model for physics

built on Neural Operators will be able to solve multiscale, multiphysics problems in a variety of domains. Similar to how complex sentences are formed by the arrangement of and interactions between words and phrases, the real-world environments in which science and engineering researchers work are made up of combinations of different physical systems. A foundation model with physical understanding would be able to learn such a compositionality and attain emergent capabilities to simulate complex physical phenomena and generate novel designs previously out of reach. The development and use of Neural Operators provide a strong foundation for building AI systems with broad physical understanding across scales and domains, equipping us to tackle the full complexity of the real world for scientific discovery.

ABOUT THE AUTHOR

Anima Anandkumar is the Bren Professor of Computing and Mathematical Sciences at the California Institute of Technology. She previously served as Senior Director of AI Research at NVIDIA and Principal Scientist at Amazon Web Services. She has recently published in such journals as *Nature*, *Science Advances*, *Nature Reviews Physics*, *Science Robotics*, and *Journal of Machine Learning Research*.

ENDNOTES

- ¹ Zhiling Zheng, Ali H. Alawadhi, Saumil Chheda, et al., “Shaping the Water-Harvesting Behavior of Metal-Organic Frameworks Aided by Fine-Tuned GPT Models,” *Journal of the American Chemical Society* 145 (51) (2023), <https://pubs.acs.org/doi/10.1021/jacs.3c12086>; and Ali Essam Ghareeb, Benjamin Chang, Ludovico Mitchener, et al., “Robin: A Multi-Agent System for Automating Scientific Discovery,” arXiv (2025), <https://arxiv.org/abs/2505.13400>.
- ² Théophile Lambert, Amin Tavakoli, Gautham Dharuman, et al., “Sequence-Based Generative AI-Guided Design of Versatile Tryptophan Synthases,” bioRxiv (2025), <https://doi.org/10.1101/2025.08.30.673177>; Maxim Zvyagin, Alexander Brace, Kyle Hippe, et al., “GenSLMs: Genome-Scale Language Models Reveal SARS-CoV-2 Evolutionary Dynamics,” *The International Journal of High Performance Computing Applications* 37 (6) (2023): 683–705, <https://doi.org/10.1177/10943420231201154>; Zhuoran Qiao, Weili Nie, Arash Vahdat, et al., “State-Specific Protein-Ligand Complex Structure Prediction With a Multiscale Deep Generative Model,” *Nature Machine Intelligence* 6 (2) (2024): 195–208; and Thorsten Kurth, Shashank Subramanian, Peter Harrington, et al., “FourCastNet: Accelerating Global High-Resolution Weather Forecasting Using Adaptive Fourier Neural Operators,” in *PASC ’23: Proceedings of the Platform for Advanced Scientific Computing Conference* (Association for Computing Machinery, 2023), 1–11, <https://doi.org/10.1145/3592979.3593412>.

- ³ Y. Qiang Sun, Pedram Hassanzadeh, Mohsen Zand, et al., “Can AI Weather Models Predict Out-of-Distribution Gray Swan Tropical Cyclones?” *Proceedings of the National Academy of Sciences* 122 (21) (2025): e2420914122, <https://doi.org/10.1073/pnas.2420914122>.
- ⁴ Christina Bergonzo and Alexander Grishaev, “Critical Assessment of RNA and DNA Structure Predictions via Artificial Intelligence: The Imitation Game,” *Journal of Chemical Information and Modeling* 65 (7) (2025): 3544–3554.
- ⁵ Anthony K. Cheetham and Ram Seshadri, “Artificial Intelligence Driving Materials Discovery? Perspective on the Article: Scaling Deep Learning for Materials Discovery,” *Chemistry of Materials* 36 (8) (2024): 3490–3495; and Mark Peplow, “AI Is Dreaming Up Millions of New Materials. Are They Any Good?” *Nature* 646 (8083) (2025): 22–25, <https://doi.org/10.1038/d41586-025-03147-9>.
- ⁶ Maziar Raissi, Paris Perdikaris, and George E. Karniadakis, “Physics-Informed Neural Networks: A Deep Learning Framework for Solving Forward and Inverse Problems Involving Nonlinear Partial Differential Equations,” *Journal of Computational Physics* 378 (2019): 686–707; and Zongyi Li, Hongkai Zheng, Nikola Kovachki, et al., “Physics-Informed Neural Operator for Learning Partial Differential Equations,” *ACM/IMS Journal of Data Science* 1 (3) (2024): 1–27, <https://doi.org/10.1145/3648506>.
- ⁷ Kurth, Subramanian, Harrington, et al., “FourCastNet: Accelerating Global High-Resolution Weather Forecasting Using Adaptive Fourier Neural Operators.”
- ⁸ Abigail Dommer, Lorenzo Casalino, Fiona Kearns, et al., “#COVIDisAirborne: AI-Enabled Multiscale Computational Microscopy of Delta SARS-CoV-2 in a Respiratory Aerosol,” *The International Journal of High Performance Computing Applications* 37 (1) (2022), <https://doi.org/10.1177/10943420221128233>.
- ⁹ Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, et al., “Fourier Neural Operator for Parametric Partial Differential Equations,” paper prepared for the International Conference on Learning Representations (ICLR 2021), January 12, 2021, <https://openreview.net/forum?id=c8P9NQVtmnO>; and Kamyar Azizzadenesheli, Nikola Kovachki, Zongyi Li, et al., “Neural Operators for Accelerating Scientific Simulations and Design,” *Nature Reviews Physics* 6 (2024): 320–328.
- ¹⁰ Kurth, Subramanian, Harrington, et al., “FourCastNet: Accelerating Global High-Resolution Weather Forecasting Using Adaptive Fourier Neural Operators.”
- ¹¹ Martin G. Schultz, Clara Betancouty, Bing Gong, et al., “Can Deep Learning Beat Numerical Weather Prediction?” *Philosophical Transactions of the Royal Society A* 379 (2194) (2021), <https://doi.org/10.1098/rsta.2020.0097>.
- ¹² Matthew Chantry, Zied Ben Bouallegue, Linus Magnusson, et al., “The Rise of Machine Learning in Weather Forecasting,” European Centre for Medium-Range Weather Forecasts, June 20, 2023, <https://www.ecmwf.int/en/about/media-centre/science-blog/2023/rise-machine-learning-weather-forecasting>.
- ¹³ Gregory Barber, “AI Hurricane Predictions Are Storming the World of Weather Forecasting,” *Wired*, September 27, 2023, <https://www.wired.com/story/ai-hurricane-predictions-are-storming-the-world-of-weather-forecasting>.
- ¹⁴ Kerry A. Emanuel, “The Theory of Hurricanes,” *Annual Review of Fluid Mechanics* 23 (1) (1991): 179–196.
- ¹⁵ Qiao, Nie, Vahdat, et al., “State-Specific Protein-Ligand Complex Structure Prediction With a Multiscale Deep Generative Model.”

- ¹⁶ Julia Slingo and Tim Palmer, “Uncertainty in Weather and Climate Prediction,” *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 369 (1956) (2011): 4751–4767.
- ¹⁷ Boris Bonev, Thorsten Kurth, Ankur Mahesh, et al., “FourCastNet 3: A Geometric Approach to Probabilistic Machine-Learning Weather Forecasting at Scale,” arXiv (2025), <https://arxiv.org/abs/2507.12144>.
- ¹⁸ Oliver Watt-Meyer, Brian Henn, Jeremy McGibbon, et al., “ACE2: Accurately Learning Subseasonal to Decadal Atmospheric Variability and Forced Responses,” *npj Climate and Atmospheric Science* 8 (1) (2025): 205; and Haiwen Guan, Troy Arcomano, Ashesh Chattopadhyay, and Romit Maulik, “LUCIE-3D: A Three-Dimensional Climate Emulator for Forced Responses,” arXiv (2025), <https://arxiv.org/abs/2509.02061>.
- ¹⁹ Boris Bonev, Thorsten Kurth, Christian Hundt, et al., “Spherical Fourier Neural Operators: Learning Stable Dynamics on the Sphere,” in *Proceedings of the 40th International Conference on Machine Learning* (PMLR, 2023), 2806–2823.
- ²⁰ Qiao, Nie, Vahdat, et al., “State-Specific Protein-Ligand Complex Structure Prediction With a Multiscale Deep Generative Model.”
- ²¹ Sun, Hassanzadeh, Zand, et al., “Can AI Weather Models Predict Out-of-Distribution Gray Swan Tropical Cyclones?”
- ²² Jacob B. Landsberg and Elizabeth A. Barnes, “Forecasting the Future with Yesterday’s Climate: Temperature Bias in AI Weather and Climate Models,” arXiv (2025), <https://www.arxiv.org/abs/2509.22359>.
- ²³ Bergonzo and Grishaev, “Critical Assessment of RNA and DNA Structure Predictions via Artificial Intelligence.”
- ²⁴ Tapio Schneider, João Teixeira, Christopher S. Bretherton, et al., “Climate Goals and Computing the Future of Clouds,” *Nature Climate Change* 7 (1) (2017): 3–5.
- ²⁵ Frank Jensen, *Introduction to Computational Chemistry* (John Wiley & Sons, 2017).
- ²⁶ Raissi, Perdikaris, and Karniadakis, “Physics-Informed Neural Networks.”
- ²⁷ Aditi Krishnapriyan, Amir Gholami, Shandian Zhe, et al., “Characterizing Possible Failure Modes in Physics-Informed Neural Networks,” *Advances in Neural Information Processing Systems* 34 (*NeurIPS 2021*), ed. Marc’Aurelio Ranzato, Alina Beygelzimer, Yann Dauphin, et al. (Curran Associates, Inc., 2021), 26548–26560; and Sifan Wang, Shyam Sankaran, Xiantao Fan, et al., “Simulating Three-Dimensional Turbulence with Physics-Informed Neural Networks,” arXiv (2025), <https://arxiv.org/abs/2507.08972>.
- ²⁸ Li, Kovachki, Azizzadenesheli, et al., “Fourier Neural Operator for Parametric Partial Differential Equations”; and Julius Berner, Miguel Liu-Schiaffini, Jean Kossaifi, et al., “Principled Approaches for Extending Neural Architectures to Function Spaces for Operator Learning,” arXiv (2025), <https://doi.org/10.48550/arXiv.2506.10973>.
- ²⁹ Nikola Kovachki, Zongyi Li, Burigede Liu, et al., “Neural Operator: Learning Maps between Function Spaces with Applications to PDEs,” *Journal of Machine Learning Research* 24 (89) (2023): 1–97.
- ³⁰ Zongyi Li, Hongkai Zheng, Nikola Kovachki, et al., “Physics-Informed Neural Operator for Learning Partial Differential Equations,” *ACM/IMS Journal of Data Science* 1 (3) (2024): 1–27, <https://doi.org/10.1145/3648506>.

- ³¹ Berner, Liu-Schiaffini, Kossaihi, et al., “Principled Approaches for Extending Neural Architectures to Function Spaces for Operator Learning.”
- ³² Vignesh Gopakumar, Stanislas Pamela, Lorenzo Zanisi, et al., “Plasma Surrogate Modelling Using Fourier Neural Operators,” *Nuclear Fusion* 64 (5) (2024), <https://doi.org/10.1088/1741-4326/ad313a>; Gege Wen, Zongyi Li, Qirui Long, et al., “Real-Time High-Resolution CO₂ Geological Storage Prediction Using Nested Fourier Neural Operators,” *Energy & Environmental Science* 16 (4) (2023): 1732–1741; and Miguel Liu-Schiaffini, Clare E. Singer, Nikola Kovachki, et al., “Tipping Point Forecasting in Non-Stationary Dynamics on Function Spaces,” arXiv (2023), <https://arxiv.org/abs/2308.08794>.
- ³³ Tingtao Zhou, Xuan Wan, Daniel Zhengyu Huang, et al., “AI-Aided Geometric Design of Anti-Infection Catheters,” *Science Advances* 10 (1) (2024), <https://doi.org/10.1126/sciadv.adj1741>.
- ³⁴ Haoyu Yang, Zongyi Li, Kumara Sastry, et al., “Large Scale Mask Optimization via Convolutional Fourier Neural Operator and Litho-Guided Self Training,” arXiv (2022), <https://doi.org/10.48550/arXiv.2207.04056>.
- ³⁵ Jiayun Wang, Oleksii Ostras, Masashi Sode, et al., “Ultrasound Lung Aeration Map via Physics-Aware Neural Operators,” arXiv (2025), <https://doi.org/10.48550/arXiv.2501.01157>.
- ³⁶ Jiachen Yao, Abbas Mammadov, Julius Berner, et al., “Guided Diffusion Sampling on Function Spaces with Applications to PDEs,” arXiv (2025), <https://doi.org/10.48550/arXiv.2505.17004>.
- ³⁷ Zelin Zhao, Zongyi Li, Kimia Hassibi, et al., “Physics-Informed Neural-Operator Predictive Control for Drag Reduction in Turbulent Flows,” arXiv (2025), <https://doi.org/10.48550/arXiv.2510.03360>.
- ³⁸ Yixuan Wang, Ziming Liu, Zongyi Li, et al., “High Precision PINNs in Unbounded Domains: Application to Singularity Formulation in PDEs,” arXiv (2025), <https://arxiv.org/abs/2506.19243>.
- ³⁹ Kaiyu Yang, Aidan M. Swope, Alex Gu, et al., “LeanDojo: Theorem Proving with Retrieval-Augmented Language Models,” in *NIPS ’23: Proceedings of the 37th International Conference on Neural Information Processing Systems* (Association for Computing Machinery, 2023).
- ⁴⁰ Dmitrii Kochkov, Janni Yuval, Ian Langmore, et al., “Neural General Circulation Models for Weather and Climate,” *Nature* 632 (8027) (2024): 1060–1066.
- ⁴¹ Chuwei Wang, Julius Berner, Zongyi Li, et al., “Coarse Graining with Neural Operators for Simulating Chaotic Systems,” arXiv (2025), <https://doi.org/10.48550/arXiv.2408.05177>.
- ⁴² Ibid.