# Transcript

Noah Kravitz (Host): Hello, and welcome to the NVIDIA AI Podcast. I'm your host, Noah Kravitz. Artificial intelligence is powering a new generation of physically accurate models and simulations that are changing the way research and planning are done.
In the past few months, we've explored industrial simulation on the podcast, talking about how companies like Siemens are designing and optimizing factories with the help of AI and technologies including digital twins. Today, we're exploring one of the most fascinating-and arguably important-uses of AI for simulation: climate simulation.
Mike Pritchard, Director of Climate Simulation Research at NVIDIA, is here to talk about the role of AI in climate science and how AI-driven climate modeling is helping us understand, predict, and prepare for extreme weather and climate change. Mike, welcome, and thank you so much for joining the AI Podcast.
Mike Pritchard (Guest): Thanks so much, Noah. It's great to meet you.
Noah: Maybe we can start by having you share a little bit about your own journey and how you got into working on climate science.
Mike: Sure. Well, it was a little windy road for me. I came out of high school not sure if I wanted to study jazz music or chemistry, and I landed in an astrophysics program. I learned about the cosmos, had my first summer job doing astrophysics research, and realized I didn't really care about how two galaxies are going to collide in 15 billion years enough to work on it all summer.
So I took a year off traveling halfway through my undergrad. I ended up visiting Bangladesh, met all these amazing people living so close to sea-level rise, and thought back to a course I'd taken on radiative transfer in planetary atmospheres. I realized there was a physics problem that really matters to the world. I haven't looked back since.
I did a master's, got a couple of ice-sheet simulators into San Diego to do a PhD studying the daily cycle of rainfall and how we can simulate rainfall systems in climate models better. And eventually, I realized there's a computer problem here: we know the equations we'd like to solve, but computers aren't up for it.
Any algorithm that could help me get more explicit cloud physics for realistic storm simulations into computer models became interesting. During 2017, AI became the most interesting algorithm to do that.
Noah: I have to ask-before we continue talking about climate science-are you a musician? Do you still play?
Mike: I play super casually. My jazz chops have dwindled, but there's a lovely sunset gig I play in San Diego with someone who lets me play not too many chords and improvise. I love music and communicating.
Noah: What's your instrument?
Mike: I play keyboards.
Noah: Great. I'm a drummer-we'll have to get together sometime.
Mike: Right on.
Noah: Let's talk climate science. Can you give a brief overview of what we mean when we talk about climate science? Then, maybe talk about how the technologies-and along with it, your perspective on climate science-have evolved over the years.

Mike: Sure. Cut me off if I nerd out too much, but we're all familiar with weather prediction, right? Climate prediction is similar, but with key differences.

Weather is like how the next few days unfold while the atmosphere still has memory of its initial conditions, which I've observed and can initialize a simulator on. Climate simulations solve very similar equations, but we're talking not about tomorrow's weather, but the climate-the average of weather, the statistics of weather. Those are predictable on different time scales for different reasons.

For example, here in San Diego, May is warmer than January because there's more solar energy-more watts per square meter-coming in. Add energy, and it heats up. Future climate change works the same way: add greenhouse gases, they back-radiate longwave radiation like little heat lamps in the sky, adding energy so the system heats up.

But the trillion-dollar question is: how much will it heat up? All these hazards-extreme storms, droughts, wildfires-scale with the overall warmth of the planet, like a fever. A lot of that depends on complicated processes we can't simulate well, like clouds. Will they dissipate like ice sheets do, revealing darker surfaces and amplifying warming? Or will they thicken and brighten the planet, reflecting more energy as the planet warms?

Simulating future climate brings supercomputers to their knees. We have to cover the whole planet at high resolution for 100 years, then run dozens of scenarios and hundreds of ensemble members to sample chaotic outcomes. It's a never-ending enterprise-like learning jazz.

I fixate on clouds because I think they're mission-critical.

Noah: You mentioned that around 2017 AI started becoming a tool of choice for simulation and for climate problems. How is AI transforming the way we simulate and understand how climate works on Earth?

Mike: AI can short-circuit Moore's law and bring resolution into simulators ahead of computational schedule. Take complicated cloud physics you can simulate convincingly for a small patch of atmosphere over a few weeks-don't scale it up to the whole planet for hundreds of years. Focus AI on that: learn the physics by training dense neural networks as emulators. Once trained, they run blazing fast.

In 2018, I had this wonderful experience where our hybrid physics-AI approach worked better than I thought, giving a 20x to 100x speedup in simulations. It totally changed my view of how much resolution and complexity we can afford today.

Fast-forward five years, and it's no longer just hybrid physics-AI simulations. We're now talking about full-AI simulators of the atmosphere. That brings huge ensembles for counterfactual weather scenarios, interactivity we never imagined, and the ability to "paint" resolution onto low-res predictions. AI is disrupting the entire Earth system modeling stack.

Noah: When you talk about resolution in climate modeling, what does that mean? Is it more precise timeframe predictions, or better predictions of what's going to happen?

Mike: In physics-based climate simulation, you write the governing equations, then divide the world up into a mesh to solve them numerically. The finer the mesh, the more computing power you need. To simulate 100 years, hundreds of times, the mesh is coarser than we'd like.

A typical government climate prediction 50 years out resolves scales of 25 kilometers or so. But if you look at satellite observations at that scale, you see complex interior structures-cloud processes, storms-that aren't explicitly represented in standard simulations. Instead, a human-crafted "parameterization" or "cartoon model" approximates that subgrid complexity. Resolution refers to the mesh size: the finer it is, the more processes you can resolve explicitly and the fewer guesses you need to make.

Noah: Can you talk about some breakthroughs you've been involved with over the years-using machine learning and AI to accelerate simulations?

Mike: Sure. One thread was my work as a university professor on hybrid physics-AI simulation. Embedding thousands of copies of an AI emulator inside a planetary model makes it hard to control-interactions between the AI subunits and fluid dynamics can push you out of sample. So we did a benchmarking activity at NeurIPS, teamed with NSF and NVIDIA, to create open code, APIs, and a Kaggle competition. That helped the community discover strategies to tame these brittle hybrid systems.

When I joined NVIDIA, I discovered a new breed of scientist training massively ambitious AI models covering the entire planet-pure AI, no physics solver, like video prediction but with a dozen atmospheric channels instead of RGB color. You train a network to predict the next time step of wind, temperature, humidity, etc., then roll it out autoregressively as a weather forecast. Three years ago, I was skeptical, but the world now agrees: these pure AI weather forecasts are the most skillful in the world, even if inscrutable.

We then asked: can generative AI do video-style super-resolution on climate data? Can we turn low-res, "blurry" predictions into super-resolved versions, synthesizing new high-res variables like radar reflectivity for rainfall? It worked far better than I expected. We now have a system called CorDiff, which does multivariate super-resolution and new-channel synthesis. It's being used to generate high-resolution state estimates of the atmosphere from coarse climate data.

Noah: When you talk about predicting weather and climate, is there a metric for accuracy-like a percentage number? How has it changed since advanced AI and generative AI?

Mike: For weather forecasting, it's very clear. The pace of improvement using AI is outpacing physics models. You can look at the "quiet revolution in weather forecasting" graph: physics-based skill improvements over decades have plateaued, while AI has leapfrogged that trajectory.

For climate prediction, it's harder: the answer is out of sample, 50 years from now. We predict a range of outcomes, and concepts like energy and mass conservation are crucial. The ecosystem effect is that researchers are adapting AI for weather prediction to climate prediction. For example, you can train an emulator not on weather observations but on climate projection outputs. Those trained AI models can be shared easily, letting people "run" climate predictions much faster than downloading huge data sets.

Some groups are building fully AI climate simulators with correct atmosphere-ocean coupling, mass and energy conservation, and proper responses to CO? forcing. It's promising but still work in progress.

Noah: What do we do with these predictions? How can advances in AI-driven climate modeling help us prepare and adapt to extreme weather and climate change?

Mike: I see two big impacts. First, we can study statistics of rare extremes-the low-likelihood, high-impact events. Observations are too limited to sample them well, and physical simulators are too slow to generate enough counterfactuals. AI models can generate huge ensembles at 1,000x speedup. For instance, Berkeley researchers used our AI model Forecast Net to generate 28,000 years' equivalent of summer 2023, giving 7,000 weather "realizations" per day instead of one. That massive dataset lets you study heat wave statistics, climate drivers, and future changes with satisfying sample sizes. Insurance companies are already using it.

Second, improved weather forecasting-more interactivity and steerability. My team released "Climate in a Model," an interactive climate sampler. You feed in three inputs-time of day, time of year, and ocean surface temperature-and it outputs thousands of samples of weather fields at 13 million pixels across 12 variables. You can even steer it by region: ask for tropical cyclones hitting Bermuda, for example, and it will generate plausible samples consistent with those boundary conditions. That kind of interactivity will empower planners, emergency managers, and policymakers.

Noah: Are governmental organizations and disaster response agencies using these interactive simulators to inform policy decisions and disaster planning?

Mike: It's an ecosystem in progress. We work with European science advisors like Bjorn Stevens at ECMWF. He envisions a workflow where you start with credible, best-in-class physics model output, then layer AI on top to broaden access-interact with the data while preserving provenance and credibility. That hybrid approach addresses liability and trust concerns.

Meanwhile, climate-tech startups are using AI for state estimation-turning proprietary sensor data into high-quality weather forecasting inputs. The broader ecosystem will decide how to use these capabilities-perhaps in digital twins of electrical grids, power stations, and ultimately the whole atmosphere and climate, enabling end-to-end optimization and resilience planning via differentiable simulation.

Noah: In the intro, I referenced digital twins in the industrial context. How do digital twins play into climate science?

Mike: We already think of climate simulators as digital twins of the Earth for scenario planning: "What if humans evolve the land surface this way?" or "What if we follow that mitigation pathway?" AI surrogates for physical climate models reduce latency in familiar digital twins. But imagine backpropagating through every layer-from the electrical grid to power stations to the atmosphere to the climate-and optimizing for resilience to low-likelihood extremes. That's the larger vision.

Noah: You gave a recent TED Talk about AI and climate simulation. Why give a TED Talk, and what message did you hope to leave the public with?

Mike: I wanted people to know something exciting is going on in atmospheric simulation-technology nobody would have predicted 10 years ago. These new AI simulators have unprecedented speed, interactivity, and multi-scale capabilities. I wanted to highlight that it's a beneficial use of AI for a problem that affects us all: weather and climate prediction. And there's a "cocoon" of collaboration between academia, government labs, and industry driving it forward.

One remarkable aspect I mentioned is AI's ability to incorporate far more observations than physics models-no human-written parameterizations to shock the model. We can use satellite and station observations directly, perhaps even build AI-only climate predictions from observations alone. The unification of AI for state estimation and prediction is happening now; we don't yet know the predictability limits of the Earth system in a world with powerful AI.

Noah: That reminds me of medical imaging and astronomy, where AI finds new markers in existing data. Do you see similar phenomena in weather observations-discovering unknown patterns in historical records that improve prediction?

Mike: Absolutely. There are papers from the University of Washington claiming to extend predictability limits-controversial but exciting. Phenomena like subseasonal forecasting (three to five weeks) involve ocean and land-surface memory that we hope AI can unlock. But we face a unique problem: only decades of satellite observations exist-tens of thousands of days, not billions of internet images. So some phenomena remain under-sampled. We'll always need a combination of physical simulators for deep-time projections and AI trained on observations for near-term improvements.

Noah: Does synthetic data play a significant role in climate science?

Mike: Yes. We can leverage large volumes of simulated data-thousands of years of AI forecasts or physical model output-for pre-training. Then we fine-tune on the highest-quality observations. That's an active area of research.

Noah: Looking ahead, what are the next big things you and the climate science community are working on, especially at NVIDIA?

Mike: Globally, the ocean simulation community is catching up to the atmosphere and weather community. The ocean absorbs the vast majority of excess heat we've put into the Earth system. Marine heat waves will largely control future interannual extremes. Full-ocean emulators have emerged this year, and people are beginning to couple them with full-atmosphere emulators. It's like the 1980s physics model coupling, but on an accelerated timeline-seasonal forecasting and El Niño predictions via AI-emulators.

At NVIDIA, I'm excited about interactivity. We published a paper called "Climate in a Bottle." It takes three inputs-time of day, time of year, and sea-surface temperature-and outputs 600 MB of sampled atmospheric fields across 13 million pixels and 12 variables. It preserves the correct seasonal and daily cycles, modes of variability, and even allows you to steer events, like drawing a map of where you want tropical cyclones to appear. It's the tip of the iceberg for new interactive, query-driven climate tools.

Noah: For concerned citizens or technologists listening, what can we do to get involved, support climate science, and support sustainability through AI?

Mike: If you're a citizen, contact your representatives about funding for sustained Earth-observing systems. Machine learning is nothing without data. If you value the observing satellites and modeling agencies that produce the data driving this AI revolution, let policymakers know.

If you're an ML researcher, try a climate-related Kaggle competition or contribute to open benchmarks-many are still unsolved. And have conversations: the planet is warming, our past emissions will continue to warm it for decades, and we need to think about what to do. AI will let us interact more directly with predictions of the future; awareness and dialogue are vital.

Noah: Mike, for listeners who want to learn more or follow your team's work, where should they go online?

Mike: Look up NVIDIA Earth2-that's the name of our initiative. You'll find software, open-source recipes for simulation technologies, APIs, and training materials. We hope the ecosystem builds on it.

Noah: Fantastic. Mike Pritchard, thank you so much for taking the time to join the podcast. It's a fascinating conversation, and despite the seriousness of climate change, you've left me optimistic about how these technologies can help us understand and adapt to the climate. What a time to be alive.

Mike: Thanks, Noah. It's a real privilege to represent a huge team of humble, earnest technologists, engineers, and scientists working on these problems. I'm just a mouthpiece-credit goes to them.