

Scaling Physics Intelligence for the Earth's Subsurface

Gege Wen

The progress of AI in the last decade has come with an enormous cost: energy. Modern AI development resembles a double Ouroboros, the Greek mythological symbol of two snakes biting each other's tails. To scale AI, we must reinvent the energy system, and to reinvent the energy system, we must harness the power of AI.

AI is a greedy consumer of power. Training a single large foundation model can consume as much electricity as thousands of U.S. households use in a year. What's more, AI doesn't just require "energy" broadly; it needs firm, stable, responsive power capable of meeting sudden compute spikes and massive parallel processing. Unfortunately, intermittent renewables like solar and wind alone cannot deliver that. The scalable and economically viable solution lies in Earth's subsurface: a mixture of natural gas-fired power plants with carbon sequestration, large-scale energy storage for renewables, and geothermal power.

But these subsurface energy systems are among the most complex and uncertain environments humans have ever tried to engineer. To design and operate them effectively, we need predictive tools that can understand how physics behaves in the subsurface, including how fluids travel within porous media, where carbon dioxide injection will cause high pressure, and whether heat can be produced sustainably in enhanced geothermal systems. All of these complex physics are occurring in the opaque, heterogeneous, and deep subsurface. That's a level of complexity that can't yet be properly captured with today's technology.

The digital toolkit that supports subsurface engineering today was built during the oil and gas boom of the 1990s and early 2000s. Having worked with these legacy systems early in my career, I've experienced firsthand their fundamental limitations. As an intern on a typical engineering team – comprising one manager, one geologist, one reservoir engineer, and one operation engineer on site – my main weekly task was to print out the production histories using Excel spreadsheets. On Mondays, we gathered in a small office, staring at columns of numbers and a large, printed map, debating where to send the drilling rig next. In this system, drilling

an exploratory well had a success rate of about 50 percent. This high failure rate shaped an ecosystem of tools designed around caution, not accuracy. Simulations could take weeks. Models were calibrated manually. Uncertainty wasn't something to eliminate; it was simply priced into the business.

These workflows were never meant to support a world in which we must quickly evaluate hundreds of potential sites or simulate multiphase carbon dioxide behavior across an entire sedimentary basin in days. Nor were they meant to provide answers to AI engineers asking whether their next data center can run on stable carbon-free power 24/7.

And so we arrive at the Ouroboros: AI needs subsurface energy to grow, and subsurface energy needs AI to be harnessed. To build the infrastructure that powers AI, we need to fundamentally rethink the way we model subsurface physical systems.

It amazes me how much harder it is to learn to simulate the physical world than to generate text. Language, for all its richness, is fundamentally one-dimensional. Words flow linearly; their meanings are flexible. There are many ways to say the same thing. In other words, natural language is low entropy and high redundancy.

The physical world is the opposite. Every molecule, every parcel of fluid, every pocket of pressure is governed by strict constraints across four dimensions: three in space, one in time. Every point in the field must simultaneously satisfy conservation laws and boundary conditions. This is a level of accuracy that was not required previously in generative language and image models.

This is why simulating physics with AI is so challenging – and so exciting. If we can train AI systems to capture the rules of the real world, not just the patterns of language, we would unlock an entirely new paradigm. By integrating physics into today's AI, we can explore thousands of geological configurations and storage scenarios and iterate quickly on the design of future energy systems, all without waiting days for an oil and gas legacy numerical solver to finish.

My team explores such models: subsurface physics intelligence that predicts how gases and fluids flow through porous rock using synthetic simulations rooted in physical laws and real-world experimental datasets. These models can forecast the long-term performance of geothermal wells and identify optimal well locations for carbon dioxide injectors, drastically shortening the engineering process for designing such projects.

One success story is the web application ccsnet.ai, which we built to make AI for subsurface carbon dioxide storage broadly accessible. The platform hosts AI models pretrained with large numerical simulation datasets, covering wide ranges of initial and boundary conditions for subsurface physical systems. Once trained, each prediction takes a fraction of a second. By democratizing access to fast and high-quality reservoir simulation, the tool enables groups that previously lacked these capabilities – including small operators, national labs, and even individuals inter-

ested in local projects – to independently evaluate storage sites. Today, the web app serves about two hundred users per month from every continent except Antarctica.

Of course, such work still needs to be scaled up to support industrial use. But it already offers a glimpse of what the energy industry could look like in the next five to ten years. Early in my career, a typical day meant going into the office, turning on my computer, setting up a numerical simulation case, and waiting – sometimes for days – for the results. I now envision the next generation of engineers and researchers beginning their day with instant, AI-driven physics predictions built directly into their workflow.

And the AI-enabled prediction will not stop at reservoir simulation alone. Fast physics prediction has the potential to accelerate decision-making across every stage of subsurface engineering. Seismic data could be integrated into subsurface models as they are acquired in the field; operational decisions could be adjusted autonomously by AI agents; and grid-level optimization systems could integrate subsurface assets in real time, producing energy in the lowest-carbon and most efficient way possible.

Only by simulating the physical world better can we build the energy systems needed to support AI's exponential growth. And only by using AI to solve subsurface physics problems can we ensure that the energy transition proceeds at the pace the world requires.

ABOUT THE AUTHOR

Gege Wen is an Assistant Professor at Imperial College London, co-appointed by Earth Science Engineering and I-X (Imperial AI). She is also the Cofounder and CEO of EarthFlow AI. She has published in *Energy & Environmental Science*, *Applied Energy*, and *Advances in Water Resources*.