

The Ocean Needs More AI, Not Less

The ocean is the most consequential system on Earth and one of the hardest to understand. AI can help, but not timidly: it will take foundation models built for the coast and the ocean, and a bigger ambition than the field has allowed itself.

Josh Humberston Founder & CEO, Coastal Measures

Kittery, Maine

Too little, not too much

In ocean and coastal circles, the AI skepticism I hear is that we should slow down: too much hype, too many confident models, too little measurement underneath them. I share a good deal of it; my own company exists because the data beneath most ocean models is thinner and messier than people assume.

But I have come to think the bigger problem runs the other way. The ocean is the most consequential system on the planet and one of the hardest to understand, and AI is the first tool we have had that might be equal to it. We have been using it far too timidly.

Where the money went

The four largest American technology companies plan to spend more than \$700 billion on capital projects in 2026, nearly double last year, almost all of it on AI infrastructure.¹ That is roughly \$2 billion a day, in service of search, advertising, social feeds and cloud computing.

I don't begrudge any of it. Well, maybe some. But the same chips and models are already spilling over into science. Hold the number next to this one: by UNESCO's count, less than 1% of national research budgets goes to ocean science.² Japan,

one of the largest national funders, spends roughly \$5 billion a year; the AI buildout spends that every three days. All the world's philanthropic giving to ocean science, about \$1.2 billion a year, is gone in under a day.




FIGURE 1 Four companies' planned 2026 capital spending, almost all of it AI infrastructure, against the largest national ocean-science budgets and all philanthropic funding for ocean science worldwide. Sources: Financial Times tally of company guidance via Sherwood News, April 2026; IOC-UNESCO Global Ocean Science Report, 3rd edition, 2026. Japan and Germany are derived from the report's statement that \$1.2 billion of philanthropy equals 22% of Japan's 2023 and 60% of Germany's 2022 ocean-science budgets.

The boom in language models came from a decade of pouring compute, data and capital at one problem until the models crossed a threshold. Nothing like that has been pointed at the ocean. Nobody can even say how much ocean-science money goes to AI, because nobody tracks it, but the order of magnitude is visible in the labs doing the work: a few dozen researchers here and there, not a few hundred billion dollars.

The largest system on Earth, barely measured

The ocean covers 71% of the planet and holds about 95% of its living space. It has absorbed roughly 90% of the excess heat in the climate system, sets the weather, feeds billions and carries most of world trade; as a country, the ocean economy would be the fifth largest on Earth.²



And we measure it thinly. As of this spring, 28.7% of the seafloor has been mapped to modern standards.³ The Argo array, the backbone of subsurface ocean observation, runs on close to 4,000 floats: one temperature and salinity profile every ten days for each patch of ocean the size of Maine.⁴

It gets patchier closer to shore, where the ports, towns and infrastructure are. Some 210 long-term NOAA water-level stations stand watch over 95,000 miles of U.S. tidal shoreline,⁵ and UNESCO notes that many observing platforms run with no redundancy at all.

This is the system we are asking engineers, insurers and town planners to make fifty-year decisions about.




FIGURE 2 How thinly the ocean is observed and funded. The Argo comparison divides the ocean's 361 million square kilometers by the array's roughly 4,000 floats: about 90,000 square kilometers per float, close to the land area of Maine. Sources: Nippon Foundation–GEBCO Seabed 2030 Project, April 2026; Argo Program; IOC-UNESCO Global Ocean Science Report, 3rd edition, 2026.

Hard is the point

People are tired of AI. Not the technology so much as the promise: that a model will fix everything and understand everything, as a black box that cannot say why the water rose. There is truth in both complaints, and they are the starting point of my argument, not an objection to it.

What AI is actually good at is finding structure in observed, nonlinear, tangled relationships that no analyst and no tractable set of equations can hold at once. The ocean is made of those relationships, across scales from millimeters to planetary and from seconds to centuries, in a system that is changing faster than our textbooks.

Held to account by the physics we already know, a learned model doesn't replace understanding. It produces it.



The gap a model finds between the equations we trust and the measurement is a clue about a process our working models have not yet captured, and that is how understanding of a rapidly evolving system advances.

Learned weather models already run alongside the physics-based ones at the world's leading forecast centers and beat them on many measures.⁶ The ocean is a harder, slower, less-observed version of the same problem, and it will not yield to a model borrowed from the atmosphere or scraped from the web. It needs its own: foundation models built for the coast and the ocean, models that hold an understanding of a place the way a brain does, many signals at once, across scales, with the physics as scaffolding. The first attempts, from small academic groups in Germany, the UK and China, are a start, not yet an ambition.⁷

That is the case for thinking bigger: not more AI because the ocean is underspent, but because it is the system we most need to understand and understand least, and because models of this kind could tell us what nothing else can: how to live with a changing coast more safely and more sustainably.

What thinking bigger has to mean

More AI in ocean science does not mean more chatbots pointed at the sea. It means scaling novel R&D in ocean-specific AI capabilities, starting with the least glamorous part: how we build the right training data.

More measurement. A learned model is only as good as the observations it is anchored to. AI's appetite for data is the best argument yet for more sensors in the water and on the shore, and for making the networks we already have work together; today coverage is partial and the pieces don't connect.

Physics first. We understand the ocean's physics far better than we understand social media feeds. A model should learn what the physics cannot cheaply resolve, not relearn what we already know. At Coastal Measures, our storm-surge models work this way: a physics-anchored baseline with a learned residual, promoted into production only when it outperforms the physics alone on data it has never seen. If it doesn't, it doesn't ship, which is what makes the claim falsifiable.






FIGURE 3 Physics first. The model never learns the tide; it learns only the residual the baseline leaves unexplained, from predictors with a physical meaning, and is promoted only if it beats the physics-only baseline on data it has never seen. The water-level traces are schematic. Source: Coastal Measures.

Open records. A model trained on closed, undocumented data cannot be checked. The record should be standardized, queryable and owned by the people who collected it, with every output traceable to its sources and honest about its uncertainty. This is why we built CUMULUS™, a data system made for coastal and ocean AI: it connects whatever sensors a port, a town or a lab already has, fills the gaps with partner hardware, standardizes everything into one open record the customer owns, and anchors every analytic to a physical baseline. The hardest part of ocean AI is not the model; it is the record the model learns from, and for most of the coast that record did not exist.

None of that is a smaller ambition than the hype. It is the only version of the ambition that survives contact with a nor'easter.

One percent of this year's AI buildout, about \$7 billion, would be the largest step-change in our understanding of the ocean in a generation. If there is a better place to point this technology than the largest, most consequential and least-understood system on the planet, I haven't found it. If you are working on it, or want to, I'd like to hear from you.

NOTES AND SOURCES

1. Alphabet, Amazon, Microsoft and Meta 2026 capital-spending guidance, compiled by Sherwood News from company earnings, April 30, 2026. [sherwood.news](https://www.sherwood.news)
2. IOC-UNESCO, *Global Ocean Science Report: Investing in Sustainable Ocean Solutions*, 3rd edition, Executive Summary, June 2026 (doi:10.71245/MSEK9116). Median share of national research budgets below 1%; philanthropy about \$1.2 billion a year in 2023–24, equal to 22% of Japan's 2023 ocean-science budget; 28% of Ocean Decade actions fully funded; ocean economy the fifth-largest if a country. unesco.org
3. Nippon Foundation–GEBCO Seabed 2030 Project, "Global seabed mapping reaches new milestone," April 20, 2026. seabed2030.org
4. Argo Program, "Argo's status." Ocean area 361 million km² ÷ 4,000 floats ≈ 90,000 km² per float; Maine's land area is about 80,000 km². argo.ucsd.edu
5. NOAA National Ocean Service: the National Water Level Observation Network's 210 continuously operating stations, and 95,471 miles of U.S. tidal shoreline. oceanservice.noaa.gov
6. ECMWF, "ECMWF's ensemble AI forecasts become operational," July 2025: the AI ensemble "outperforms state-of-the-art physics-based models for many measures." ecmwf.int
7. OceanRep (Alfred Wegener Institute, EGU 2025); IBM Research UK's Sentinel-3 ocean foundation model on Copernicus WvEEO (2026); OceanPile (Zhejiang University, arXiv 2605.00877, 2026). copernicus.org

Josh Humberston is the founder and CEO of Coastal Measures, a coastal ocean intelligence company in Kittery, Maine. He holds a PhD in oceanography from the University of New Hampshire and previously built marine data systems at Sandia National Laboratories, the U.S. Army Engineer Research and Development Center and NOAA.

Soundings is an occasional series from Coastal Measures on measurement, data and decisions at the coast.



Coastal Measures is a coastal ocean intelligence company in Kittery, Maine. We build and operate CUMULUS™, a platform that connects any sensor on any transport, standardizes the record so the customer owns it, and anchors analytics to physical baselines with stated uncertainty, for ports, coastal towns, utilities and research institutions.

coastalmeasures.com