



2025
AI
JOURNEYS

M | MIDAS MICHIGAN INSTITUTE
FOR DATA & AI IN SOCIETY
UNIVERSITY OF MICHIGAN



ABOUT

“AI for domain research” is a prominent topic that has caught the attention of many academic researchers. However, this is still just an idea for most domain researchers; their initial effort to explore this idea faces many challenges. At the same time, we also need to explore how this idea is transforming how we do research and the role of human researchers, so that this idea will lead us not only to better and faster research, but also more creativity and fulfillment of human researchers. This collection features “AI Journeys” presentations at the 2025 AI in Science and Engineering Symposium, organized by the Michigan Institute for Data and AI in Society (MIDAS). In these presentations, domain scientists as well as AI experts shared with the audience their explorations of using AI to advance domain research, including:

- How they selected a significant domain research question
- How they decided that they should leverage AI and a particular AI method
- How they adopted AI in the project(s)
- How they ensured sufficient expertise, resources and collaboration
- Results, successes and lessons learned.

We are grateful for the presenters and their candid stories, which will surely inspire others to use thoughtful and effective approaches to adopt AI to enable research breakthroughs.

TABLE OF CONTENTS

AI Model Design to Predict and Mitigate Financial Risk from Water Scarcity in Global Corporate Facilities

Peter Adriaens, University of Michigan — Page 3

Designing Trustworthy Drug Combinations with Mechanistic Neural Networks

Harkirat Singh Arora, University of Michigan — Page 4

Bridging Innovation and Equity: How the African Union is Driving AI for Sustainable Development in Africa

Barbara Glover, African Union Development Agency — Page 5

Navigating Environmental Data: Unsupervised Classification in Oceanic and Atmospheric Research

Dani Jones, University of Michigan — Page 6

Towards Interpretable Machine Learning Models Across the Geosciences

Fraser King, University of Michigan — Page 7

ChatATC: Large Language Model-Driven Agents for Strategic and Tactical Air Traffic Management and Control

Max Li, University of Michigan; Sinan Adulhak, University of Michigan; Wayne Hubbard, FAA; and Karthik Gopalakrishnan, Stanford — Page 8

Lessons from a Human-in-the-loop Machine Learning Approach for Identifying Vacant, Abandoned, and Deteriorated Properties in Savannah, Georgia

Xiaofan Liang, University of Michigan — Page 9

All you need is community

Abiodun Modupe, University of Pretoria — Page 10

Revolutionizing African Great Lakes Management: AI for Smart Monitoring and Data-Driven Decision Making

Grite Nelson Mwaijengo, Nelson Mandela African Institution of Science and Technology (NM-AIST) & African Center for Aquatic Research and Education (ACARE) — Page 11

Adaptive, Safe, and Efficient Generative AI for Scientific Applications

Qing Qu, University of Michigan — Page 12

Can LLMs and AI Agents Accelerate Real-World Adoption of Scientific Evidence?

Geoffrey H. Siwo, University of Michigan — Page 13

Thinking outside the black box with AI-assisted epidemiological models

Jon Zelner, University of Michigan — Page 14

Building AI for Bladder Cancer Survival Prediction: A Multidisciplinary Journey

Lubomir Hadjiyski, University of Michigan — Page 15



AI MODEL DESIGN TO PREDICT AND MITIGATE FINANCIAL RISK FROM WATER SCARCITY IN GLOBAL CORPORATE FACILITIES

PETER ADRIAENS, UNIVERSITY OF MICHIGAN

Peter Adriaens's AI journey didn't begin in a lab; it began with a question from a pension fund: How is water scarcity affecting the financial performance of the companies and holdings we invest in? That question has evolved into a cutting-edge AI project tackling one of the most underappreciated but urgent risks to global capital: water-related financial exposure across corporate facilities.

In a space where regressions still rule and risk-aversion dominates, Adriaens's team is developing a novel AI architecture to quantify water risk for investors. His work spans engineering, sustainability, and finance fields not often bridged, and aims to make climate- and water-related threats financially legible at the facility scale.

It all started when TIAA, a major pension fund, approached him with growing concerns about operational disruptions from water shortages, particularly in global semiconductor plants, manufacturing operations, and pharmaceutical supply chains. TIAA, like other asset managers and corporations, reports these risks under the Task Force for Climate-Related Financial Disclosures (TCFD), promulgated by the Financial Stability Board (FSB) in 2017 and incorporated in 2024 as a financial reporting requirement as part of International Financial Reporting Standards (IFRS) regulation. Adriaens's team began mapping how these risks propagate across 150,000+ facilities held by S&P 500 companies, a staggering \$7.3 trillion in insurance portfolio assets under management potentially affected.

At the heart of the work is an AI pipeline designed to simulate and analyze water purchase contracts, the legal agreements companies sign with utilities to guarantee water access for their operations. These contracts are rarely disclosed in financial reports, highly variable by region and sector, and central to operational continuity for industries like data centers, chemical manufacturing, and real estate.

To build a system that could provide actionable intelligence to CFOs, supply chain officers, and fund managers, Adriaens and their team had to integrate:

- Water risk data (e.g., drought, flood, and regulation data from WRI's Aqueduct platform),
- Facility-level corporate data (from the Bloomberg Terminal and the ORBIS private company database),
- Financial disclosures and earnings calls transcripts (parsed with transformer-based NLP models),
- And water procurement contractual language structures, enabling the assessment of counterparty risk and optimization of hedge strategies.

The result is an AI-assisted interface where a financial officer can ask, "What is our facility-level water exposure in Malaysia?" and receive a breakdown of risks, exposures, and contract clause recommendations, often pinpointing individual facilities as financially vulnerable or resilient based on water access.

Adriaens emphasized that progress didn't come from technical breakthroughs alone. His team held weekly meetings with asset managers, helping them build trust in the models and gradually move from regression-based comfort zones to more advanced rule-based and learning-based techniques.

One of the major learnings was around data asymmetry: most firms don't disclose any facility-level water use, and the models must infer risks using rule-based logic and financial proxies. Validation is being conducted through private partnerships, including with Kurita Water Industries and its companies on Japan's Nikkei 225, giving the team access to behind-the-firewall facility data and enabling side-by-side comparisons between predictions and internal reality.

Currently, the model is being scaled to support real-world applications — from informing water contract negotiations at AWS data centers to helping Japanese conglomerates like Marubeni rank and hedge their portfolio exposure. The next step? Embedding these insights into smart contracts using blockchain, in collaboration with Ripple's blockchain research initiative.

"Water risk isn't just about scarcity — it's about counterparty strategy, earnings resilience, and boardroom decisions."



DESIGNING TRUSTWORTHY DRUG COMBINATIONS WITH MECHANISTIC NEURAL NETWORKS

HARKIRAT SINGH ARORA, UNIVERSITY OF MICHIGAN

When Harkirat began their PhD in biomedical engineering, they weren't just looking to build a model, they were searching for a way to help solve one of the most pressing public health threats of our time: antibiotic resistance. Their journey led to the development of a novel mechanistic neural network model that aims to optimize combination therapies by balancing two competing needs: maximum potency against infections and minimum toxicity to human cells.

Harkirat's project centers on the insight that combining existing FDA-approved antibiotics could offer a workaround to the stalled pipeline of new antibiotic discovery. But with hundreds of potential drugs to combine, the search space is overwhelming. To tackle this, Harkirat turned to machine learning, but not just any black-box model. His work is grounded in domain knowledge, integrating omics data with established bacterial metabolic networks to featurize drug combinations in a biologically meaningful way.

Instead of using off-the-shelf models, Harkirat's neural network architecture was explicitly designed around known biological pathways. By segmenting input features according to the metabolic subsystems they affect, their mechanistic model enables interpretability, allowing researchers to trace which pathways are most influential in driving predicted potency or toxicity.

The tradeoff? The mechanistic model performed slightly worse than its more generic counterpart. But Harkirat embraces this compromise. "There's always a tension between accuracy and interpretability," they explain. "In healthcare, interpretability often matters more."

To validate the model's predictions, Harkirat partnered with experimental biologists at the Center for Chemical Genomics at the University of Michigan. They tested predicted drug combinations on kidney and liver cells, two of the most vulnerable organs to antibiotic toxicity.

One surprising finding: when vancomycin (a known kidney-toxic antibiotic) was combined with azithromycin, the toxic effect appeared to reverse. Analysis of electronic health records corroborated this result, showing that patients given the combination were significantly less likely to develop kidney complications than those given vancomycin alone.

The validation reinforced the model's clinical relevance, but it also highlighted how far machine learning still has to go in earning trust in biology and medicine. As Harkirat put it during their talk:

"Everyone trusts an experiment except the person who did it. No one trusts a model except the person who built it."

Their response is twofold: design models that reflect the mechanisms scientists already understand, and validate them rigorously with real-world experiments. Harkirat's model doesn't just output predictions; it reveals insights, such as the potential role of nucleotide salvage pathways in governing toxicity and the importance of alternate carbon metabolism and transport pathways in therapeutic effectiveness.

Looking ahead, Harkirat envisions expanding their work to new domains such as fungal infections and cancer, where drug resistance also poses a major challenge. He's also exploring new feature representations that combine protein interaction networks with metabolic data.

This journey is a model of how to do AI in biomedicine right: grounded in domain knowledge, transparent in design, and rigorously validated.

For Harkirat, AI isn't magic, it's a tool to co-create science that clinicians can trust and patients can benefit from.



BRIDGING INNOVATION AND EQUITY: HOW THE AFRICAN UNION IS DRIVING AI FOR SUSTAINABLE DEVELOPMENT IN AFRICA

**BARBARA GLOVER, AFRICAN UNION DEVELOPMENT
AGENCY**

Barbara Glover's AI journey is centered around building the policy scaffolding and regional alliances necessary for an entire continent to thrive in the age of AI. As a senior official at the African Union Development Agency (AUDA-NEPAD), Glover operates at the intersection of continental policy, socioeconomic development, and frontier technologies.

Her mission: ensure that Africa doesn't merely adopt artificial intelligence, but that it shapes and leads it, on its own terms.

For Glover, the starting point is scale. With 55 member states, the African Union represents a vast and diverse region with complex governance structures, enormous youth potential, and distinct technological needs. AI in this context isn't about beating benchmarks or publishing papers. It's about leveraging a young, rapidly growing population, 60% under the age of 25, to create inclusive, sovereign, and mission-driven technological solutions.

But Glover is clear-eyed about the challenges: fragmented regulations, limited cloud infrastructure, weak compute access, underinvestment in R&D, and the reliance on foreign data storage and models. The question they and their colleagues ask isn't whether Africa will adopt AI, but how to ensure that adoption happens in a way that aligns with local needs, protects data sovereignty, and generates equitable economic benefit.

The answer lies in deliberate collaboration.

Under Glover's leadership, AUDA-NEPAD has championed pilot-driven approaches that incubate AI tools in specific countries and scale them when successful. These pilots are supported by Africa's broader 50-year development blueprint — Agenda 2063 — and embedded in efforts like the African Continental Free Trade Area and the High-Level Panel on Emerging Technologies.

Glover points to a growing landscape of African-led innovation:

- In health, AI tools support cervical cancer diagnosis in under-resourced clinics.
- In agriculture, AI powers pest detection and soil health analysis.
- In finance and governance, tools like Flutterwave and GovChat are reshaping credit scoring and service delivery.
- In education, projects like Masakhane and Zindi are training the next generation of African AI researchers.

But Glover emphasizes that the policy landscape must evolve in tandem. She is a key advocate for harmonized, continent-wide AI regulation, pushing for shared data policies, ethical frameworks, and investment mechanisms that reflect Africa's priorities. Recent efforts include the AU's continental AI strategy, regional digital trade protocols, and the upcoming Smart Africa AI Council.

Importantly, Glover sees an opportunity for international researchers to join this momentum, not by imposing external models, but through demand-driven partnerships that reflect Africa's needs. For Glover, Africa's AI future isn't about catching up. It's about taking a different path, one rooted in context, equity, and co-creation.

"AI in Africa isn't about joining the race. It's about deciding where we're going — and designing systems that get us there together."



NAVIGATING ENVIRONMENTAL DATA: UNSUPERVISED CLASSIFICATION IN OCEANIC AND ATMOSPHERIC RESEARCH

DANI JONES, UNIVERSITY OF MICHIGAN

What happens when a physical oceanographer picks up machine learning? Dani Jones's AI journey began not with a grand research agenda but with a summer student and a simple idea: apply unsupervised learning to ocean profile data in a new region. That project, classifying **Argo float** profiles in the Southern Ocean using Gaussian Mixture Models, not only uncovered meaningful hydrographic structures, but it also launched a deeper exploration into how AI can support Earth system science.

Dani's work took off during their time at the British Antarctic Survey, where a culture of interdisciplinary exchange helped seed the **BAS AI Lab**. Through AI journal clubs, code reviews, and collaborative pilot projects, their team co-developed tools like **IceNet**, a deep learning model for Arctic sea ice forecasting that outperformed traditional baselines 2–6 months out. Another tool, **DeepSensor**, used convolutional Gaussian neural processes and active learning to guide where new observations could most improve environmental forecasts, a crucial need in sparse-data regions like Antarctica.

Since moving to the University of Michigan, Dani has shifted from oceans to the Great Lakes, launching the early stages of a new Great Lakes AI Lab. Their approach is to repeat the same grassroots process: start small, work with great students, build tools on open datasets, and grow organically toward broader scientific applications.

Dani's talk was also a reminder of the balance between interpretability and performance in environmental modeling. While deep learning models can offer powerful predictions, techniques like GMM and Bayesian regression trees provide interpretable outputs, critical when trying to inform science and policy. Their lab is now exploring both sides of this spectrum, building a "healthy diversity" of approaches.

"You don't always need a revolutionary new method. Sometimes, just trying a good method in a new place opens up a whole line of research."

- We'll examine a few of these different ML techniques in this talk

TOWARDS INTERPRETABLE MACHINE LEARNING MODELS ACROSS THE GEOSCIENCES

FRASER KING, UNIVERSITY OF MICHIGAN

Fraser King's AI journey started with clouds, literally. As a postdoc in the Department of Climate and Space Sciences and Engineering (CLASP), King develops AI models to better understand precipitation, cloud structure, and atmospheric dynamics. But unlike many machine learning researchers who chase performance metrics alone, King is focused on something deeper: trust.

"Just because a model gives a good result doesn't mean we understand it," they said. "For climate and atmospheric science, that's not good enough."

His work addresses a crucial question: Can we move beyond black-box predictions toward white-box understanding, especially when those predictions affect real-world decisions like flood warnings or climate projections?

King began by applying convolutional neural networks (CNNs) and U-Nets to inverse problems in atmospheric science, such as estimating precipitation rates from radar observations or reconstructing "blind zones" in satellite-based cloud profiling radars where surface clutter prevents accurate readings. These models performed well, sometimes even outperforming traditional extrapolation techniques. But performance was only part of the story.

To figure out why the models worked and how, King turned to explainability techniques. Using SHAP (Shapley Additive Explanations), saliency maps, and visual inspection of CNN feature maps, their team examined which input features mattered most. In one example, the model learned to prioritize cloud reflectivity just above the radar blind zone—a signal atmospheric scientists also use to infer surface precipitation. It also paid close attention to velocity gradients associated with melting layers, again, a known physical phenomenon.

"That gave us confidence," said King. "Not just that the model was right, but that it was right for the right reasons."

Still, explanation is not interpretation. To take the next step, King has begun experimenting with interpretable AI, building toy models with sparse autoencoders and analyzing neural activations at the level of individual neurons. He draws on recent work from vision models to investigate superposition, the idea that a single neuron can encode multiple unrelated features. In climate models, this makes it harder to trace model behavior to physical laws. But by intentionally designing simple architectures and isolating key neurons, their team is beginning to map which latent representations correspond to known atmospheric features.

King also explored unsupervised learning approaches like UMAP to uncover hidden patterns in precipitation particle data from ground-based distrometers. These methods revealed distinct clusters representing snow, rain, graupel, and virga, suggesting AI could help redefine how we classify precipitation, not just predict it.

What makes King's work stand out is its emphasis on trustworthiness as a first-order design principle. He doesn't see AI as a replacement for traditional physics-based models, but as a complement, an augmenting layer that can learn from data while still respecting scientific constraints.

As Earth system models grow more complex and essential to public decision-making, King's work is part of a broader shift in the geosciences, from modeling for accuracy to modeling for accountability.

"Explainable AI gives us clues. But to really trust these models, we need interpretability—models that let us trace predictions back to physical meaning."



CHATATC: LARGE LANGUAGE MODEL-DRIVEN AGENTS FOR STRATEGIC AND TACTICAL AIR TRAFFIC MANAGEMENT AND CONTROL

MAX LI, UNIVERSITY OF MICHIGAN; SINAN ADULHAK, UNIVERSITY OF MICHIGAN; WAYNE HUBBARD, FAA; AND KARTHIK GOPALAKRISHNAN, STANFORD

For Max Li, air traffic isn't just about planes in the sky; it's about stakeholder coordination, safety, and decision-making under uncertainty. At the heart of his research is a deceptively simple question: Can AI help air traffic flow managers better plan for capacity bottlenecks like weather delays without significantly overhauling the deeply manual, high-stakes system?

The answer, after many conversations with aviation stakeholders such as the Federal Aviation Administration (FAA), turned out to be yes, but with nuance. And with humility.

Air traffic control is a domain where introducing AI triggers two responses: excitement at the potential for AI to improve current workflows, or alarm regarding inappropriate use and AI replacing humans. So, Max's team avoided anything that suggested automation or replacement. Instead, they asked: how can LLMs best support the people already managing the system?

The key innovation lies in targeting the "repetitive half" of a human decision-making workflow, specifically, historical recall. When major U.S. airports face congestion (due to weather, wind shifts, or events like the Super Bowl), traffic managers at the FAA plan and evaluate the need for traffic management initiatives, such as Ground Delay Programs (GDPs). GDPs are strategic interventions that assign departure delays to inbound flights while they are still at their origin airport, so that they arrive during manageable time slots, ideally minimizing costly airborne holding.

Historically, GDPs are crafted by experts who consult past decisions and tailor them to the day's unique circumstances. Li's team saw an opportunity to train a large language model on more than 20 years of GDP data, including structured parameters and unstructured "comment" fields, to answer natural language questions like: "What was the traffic management strategy in place for Newark during adverse wind conditions?"

The resulting chatbot doesn't suggest plans. It retrieves and summarizes similar past cases, freeing the traffic manager to focus on today's distinctive constraints. This maintains human authority while boosting situational recall, offering a balance between familiarity and novelty.

Li's team conducted interviews with a variety of aviation stakeholders, from FAA to airlines, and iteratively improved the tool to meet their needs. One major request: show graphical outputs to complement text. Another: ensure outputs are grounded in historical data, not hallucinated. Now, the team is refining a retrieval-augmented generation (RAG) architecture and building a prototype mobile app for hands-on testing.

While quantifying time or cost savings remains a challenge, Li emphasizes that this is not about optimization; it's about operational support. By starting small and building trust, the project avoids the pitfalls of overpromising AI's capabilities in a highly safety-sensitive environment.

Looking ahead, Li envisions an ecosystem of advanced AI agents, offering context-aware insights, not only for strategic planning like GDPs, but also potentially for tactical roles in air traffic control. But for now, their team is focused on one big win: making traffic flow managers' jobs just a little easier, one query at a time.

"We're not replacing the human in the loop: We're giving them an easy-to-use AI sidekick trained on decades of prior data — and available on their phone."



LESSONS FROM A HUMAN-IN-THE-LOOP MACHINE LEARNING APPROACH FOR IDENTIFYING VACANT, ABANDONED, AND DETERIORATED PROPERTIES IN SAVANNAH, GEORGIA

XIAOFAN LIANG, UNIVERSITY OF MICHIGAN

Can artificial intelligence help shrinking cities reclaim abandoned and deteriorated properties for affordable housing? Faced with the city's ambitious goal of acquiring 1,000 properties for redevelopment, but with only a handful of members or a small code enforcement team and no reliable inventory of vacant, abandoned, and deteriorated (VAD) properties, Liang's team turned to machine learning, with a twist: humans in the loop.

But there was no ground truth to train a model. The city had only acquired about 100 properties over a decade, with inconsistent documentation. So, the team started from first principles, building an infographic-style social-technical map of how property decisions were made and what data (like tax delinquency, code violations, and crime reports) informed those decisions.

They proposed a human-in-the-loop active learning framework where local experts labeled a smart, diverse subset of properties selected by the algorithm. The result? A model that could identify likely VAD properties more comprehensively than either field surveys or simple proxies, while also exposing tensions in expert judgment, for example, human experts tended to overweight visual cues like broken windows, while the model relied heavily on tax data.

Yet technical success didn't lead to deployment. Savannah lacked the infrastructure and governance to operationalize the model, and some foundational datasets (like crime reports) turned out to be misleadingly geo-tagged. The biggest takeaway: AI tools are only as useful as the systems ready to receive them.

Liang's reflections go beyond Savannah. They challenge researchers to go beyond extractive "oversight labeling" and instead treat local expertise as a form of co-creation. They are now extending this approach to AI-assisted zoning interpretation, while also analyzing how cities are using AI, and teaching a new course on Urban AI to prepare future planners for this complex terrain.

"A better model can't fix bad data and weak governance. Sometimes, the real win is giving planners a new way to think, not a new tool to deploy."



ALL YOU NEED IS COMMUNITY

ABIODUN MODUPE, UNIVERSITY OF PRETORIA

Abiodun's AI journey reflects a vision rooted in community, culture, and collective empowerment. Trained originally in Nigeria and now based at the University of Pretoria in South Africa, Abiodun's path into machine learning began after a career in banking software, a pivot driven by a desire to explore deeper social questions through research.

Now a senior member of the Data Science for Social Good group at Pretoria, Abiodun helps lead a multidisciplinary master's program that bridges departments and faculties. But their work extends far beyond the university, into the broader African continent and its unique challenges.

Abiodun sees AI not as a commercial tool, but as a form of social infrastructure. "We know we have problems," they explained. "But we also know the best solutions will come from within, from us." To that end, Abiodun has championed grassroots AI initiatives across Africa, aiming to lower barriers to entry and unify African nations around shared tools and goals.

One major platform for this vision is **Deep Learning Indaba**, the largest annual gathering of African machine learning researchers and students. Abiodun first attended as a tutor in 2017. Since then, the event has grown exponentially, from 150 participants to nearly 1,000 recently, offering tutorials, mentorship, and a pan-African spirit of collaboration. "It's where we write our own story," they said.

Another initiative they hold dear is Masakhane, a grassroots natural language processing community dedicated to preserving and empowering African languages through AI. Rather than building new models from scratch, Masakhane evaluates and fine-tunes existing large language models (LLMs) using local data.

Through a participatory feedback system, users can flag inaccurate translations and directly improve the model, ensuring that AI reflects the linguistic and cultural realities of African societies.

These efforts are not merely academic. They are driven by lived realities: machine translation systems that confuse "modúpé" (Yoruba for "thank you") with unrelated meanings in other languages and multilingual societies where poor translations can erase nuance, identity, and value. "Once I lose my language," Abiodun noted, "I lose my identity."

Through Deep Learning Indaba, Masakhane, and new efforts like Hundzula, which brings together language practitioners, journalists, and sociologists for AI literacy training, Abiodun is working to democratize AI across disciplines and borders.

Their journey shows that meaningful AI innovation doesn't require vast compute or big-tech funding, it requires trust, cultural insight, and community. It requires networks of people who understand that building responsible AI starts by listening to those closest to the problems.

"Whatever we build, we must build it together, in our languages, for our people, with our values at the center."



REVOLUTIONIZING AFRICAN GREAT LAKES MANAGEMENT: AI FOR SMART MONITORING AND DATA-DRIVEN DECISION MAKING

GRITE NELSON MWAJENGO, NELSON MANDELA AFRICAN INSTITUTE OF SCIENCE AND TECHNOLOGY (NM-AIST) & AFRICAN CENTER FOR AQUATIC RESEARCH AND EDUCATION (ACARE)

Grite Nelson Mwaijengo is at the forefront of a growing pan-African effort to apply artificial intelligence to one of the continent's most pressing environmental challenges: the sustainable management of the African Great Lakes.

Stretching across East and Central Africa, the African Great Lakes, such as Lake Victoria, Lake Tanganyika, Lake Albert, and Lake Edward, contain 25% of the world's freshwater. They support the livelihoods of more than 60 million people and harbor approximately 10% of the world's freshwater fish species, many of which are endemic. However, these lakes face mounting threats from pollution, overfishing, habitat degradation, invasive species, and the impacts of climate change.

"Most of our lake monitoring still relies on conventional methods; manual sampling, slow laboratory analyses, and significant data gaps," Mwaijengo explained. "AI offers a way to change that."

Working through the African Center for Aquatic Research and Education (ACARE), Mwaijengo is part of a continent-wide network of scientists developing AI-powered systems to revolutionize the monitoring, modeling, and management of the African Great Lakes. Her vision is bold: deploy automated sensors, advanced satellite monitoring systems, predictive models, and smart governance tools to better protect these critical freshwater ecosystems and the millions of people who depend on them.

Key opportunities for AI include:

- **Automated Water Quality Monitoring:** AI-powered sensor buoys can detect key indicators of pollution and alert lake managers in real time. While a limited number of these buoys are currently deployed in Lake Victoria and Lake Tanganyika, their coverage remains limited. Mwaijengo envisions a much broader deployment, integrated with machine learning models capable of analyzing and interpreting incoming data to identify anomalies such as cyanobacteria blooms, oxygen depletion, and sudden changes in water quality.

- **Satellite Remote Sensing + Deep Learning:** Many parts of the lake are remote and difficult to access, making consistent monitoring a challenge. Mwaijengo's team is exploring the use of satellite imagery combined with deep learning algorithms to estimate water quality parameters, detect invasive aquatic weeds, and map large-scale environmental changes. This approach aims to fill critical data gaps and strengthen lake monitoring across vast and hard-to-reach regions.
- **Predictive Climate Models:** With rising lake levels and shifting temperatures, AI-driven climate models have the potential to forecast flood risks, thermal stratification events, and fish die-offs in cage aquaculture systems. "Right now, communities get no warning. AI could help change that," Mwaijengo said.
- **Smart Decision Support Tools:** Lake managers often struggle to interpret complex and diverse datasets essential for informed policy- and decision-making. Mwaijengo is advocating for AI-powered dashboards that integrate varied datasets and provide actionable insights to guide effective interventions.
- **Cross-Border Data Integration:** Many African Great Lakes span multiple countries, Lake Victoria alone borders Tanzania, Kenya, and Uganda, with two more nations sharing its watershed. AI-enabled data platforms could play a key role in harmonizing and sharing information across languages, borders, and governance systems. A prototype system is already in place through ACARE, but Mwaijengo hopes to make it more automated, scalable, and accessible to all stakeholders.

Challenges remain. Funding for deploying sensors and related infrastructure is limited. Technical AI expertise is scarce, and the long-term maintenance of systems remains uncertain. However, Mwaijengo sees opportunity in collaboration, with fellow African researchers and institutions, as well as with partners from the Global North.

"Harnessing AI for Africa's Great Lakes is not just about technology," she emphasized. "It's about securing a sustainable future for the lakes and the communities that depend on them. Together, we can turn data into action and innovation into impact."



ADAPTIVE, SAFE, AND EFFICIENT GENERATIVE AI FOR SCIENTIFIC APPLICATIONS

QING QU, UNIVERSITY OF MICHIGAN

Qing Qu's journey with generative AI is not just about building smarter models; it's about adapting these models for science and engineering challenges while ensuring they remain safe, efficient, and interpretable.

An assistant professor in Electrical and Computer Engineering, Qu's research has always focused on the mathematical foundations of AI: optimization, machine learning, and inverse problems. But in recent years, they and their team have been pushing into a new frontier: applying generative AI models, often associated with images and text, to complex scientific domains like weather prediction, fluid dynamics, and spectroscopic imaging.

At the core of Qu's approach is a simple but powerful insight: data and models in scientific problems often have **low-dimensional structures** hidden inside their complexity. If these structures can be captured, it becomes possible to make AI models **more interpretable, more controllable, and more efficient**, qualities often missing from large, black-box generative models.

Rather than discussing only theory, Qu shared real-world examples where their lab adapted generative AI methods to scientific problems:

Tackling Data Assimilation in Stochastic Systems

Predicting systems like weather patterns or ocean currents requires combining limited, noisy observations with uncertain models. Classical techniques like Kalman filters or particle filters either rely on unrealistic assumptions or struggle in high dimensions.

Qu's team developed **FlowDAS**, a novel AI method that learns the **state transition** dynamics between observations, rather than trying to model everything at once. Unlike diffusion models, which estimate joint distributions without structure, FlowDAS explicitly captures how systems evolve over time.

Their method showed striking improvements, outperforming particle filters and modern AI baselines in both low-dimensional (e.g., Lorenz systems) and high-dimensional (e.g., Navier-Stokes equations) settings, even when the underlying physics was unknown.

Advancing Generative AI Safety through Interpretable Attacks and Defenses

While generative models are powerful, they can also generate unsafe or copyrighted content, raising concerns about privacy, ethics, and robustness. Machine unlearning tackles the challenge often by fine-tuning models to forget harmful content, but are fragile and easily bypassed through adversarial prompts.

Qu's group tackled this by introducing an **interpretable attack framework**: instead of crafting opaque, hard-to-understand attacks, they analyzed the models' **token embedding spaces**. By learning structured, linear relationships in these embeddings, they could both create more robust attacks (to reveal weaknesses) and design **orthogonal defenses** that surgically remove unsafe behaviors without destabilizing the model.

Bridging Simulation and Reality in Scientific Imaging

In fields like spectroscopy, where training data is scarce and models must bridge simulated and real-world measurements, Qu's team explored **in-context learning (ICL)**.

By training transformer models on simulated examples and guiding them during real-world inference, they created systems that generalized across the "Sim2Real" gap, critical for advancing miniaturized sensing devices.

Through all these projects, a few themes emerged:

- Scientific AI demands careful modeling, not just bigger networks
- Interpretability and safety must be designed into systems from the start
- Low-dimensional thinking, and understanding the underlying structures, are key to making generative AI both more powerful and more responsible

Qu's work demonstrates that **generative AI** is not just about creating — it's about **understanding, adapting, and elevating** scientific discovery.

"Scientific AI isn't about generating more data — it's about uncovering the structures that truly matter."



CAN LLMs AND AI AGENTS ACCELERATE REAL-WORLD ADOPTION OF SCIENTIFIC EVIDENCE?

GEOFFREY H. SIWO, UNIVERSITY OF MICHIGAN

Two decades ago, in a cyber café in Kenya, a young undergraduate paid a dollar an hour to test a new biological hypothesis using dial-up internet. That idea, about how HIV might interact with ancient viral DNA sequences embedded in the human genome, eventually led him to present at a major conference in the U.S., supported by none other than Anthony Fauci. This early experience forged a lifelong belief: computing can radically accelerate how science is done, where it's done, and who gets to do it.

Now, as part of the **Ecosystems Finance and Health (EFH) initiative**, Siwo is applying large language models (LLMs) to bridge the longstanding gap between scientific discovery and real-world decision-making. Their vision: an AI co-scientist that can support biodiversity and public health policy, not by replacing researchers, but by making scientific insights more accessible to non-scientists.

At the heart of this effort is a new “intelligence platform”, a tool designed to help governments, funders, and conservation practitioners use AI to analyze research findings and local datasets. Whether advising on crop choice using soil data or evaluating proposals for biodiversity conservation, the goal is to empower non-technical users to benefit from scientific knowledge without having to parse technical literature.

To build this platform, the team launched a 3-day innovation sprint (or hackathon) in Nairobi, Kenya, in partnership with organizations like the Smithsonian, Kenya Institute of Primate Research, University of Michigan, AI Kenya, Science for Africa Foundation, Yale University, and Nature Finance. The sprint brought together 33 early-career African scientists, founders, and developers, many with hands-on experience integrating LLMs through APIs like OpenAI's, to explore how generative AI could analyze both scientific papers and structured datasets. With the support of facilitators drawn from organizations including IBM Research Africa, Microsoft Research Africa, Netflix, Michigan AI Lab and Elevance Health, the participants also gained training in technical areas such as retrieval augmented generation (RAG) and the development of AI agents.

The participants tackled two technical challenges:

- **Contextualizing scientific literature:** Could an LLM summarize and answer questions about domain-specific content like ecological or public health research papers, even incorporating citations or figures?
- **Automated interrogation of tabular datasets:** Could an LLM interpret CSV or database files, allowing a policymaker or NGO staff member to ask questions like “Which intervention is most cost-effective in this region?”

Rather than just prompting ChatGPT, participants built full-stack solutions using retrieval-augmented generation (RAG), semantic chunking, and external tool integrations like web search. One team even engineered the system to reference the precise figure or table from which its conclusions were drawn, an essential feature for transparency in scientific communication.

Beyond technical creativity, the sprint also tested a critical assumption: Can AI lower the barrier for scientific participation and decision-making? By co-designing tools with local scientists, the initiative reframes LLMs not as distant or general-purpose solutions, but as collaborative partners in regionally grounded, scientifically informed action.

Looking ahead, the team envisions expanding these tools into living labs, real-world ecosystems where AI-driven insights can inform environmental and health policies, paired with a financing hub to support sustainable, evidence-based projects.

Siwo closed by positioning this work not only within the growing field of AI for science but as part of a broader movement: AI for inclusive innovation. By starting with local expertise and problems, the initiative is helping to rewrite the story of who gets to shape the future of science, one model, one sprint, and one community at a time.

“We don't just want to accelerate discovery. We want to accelerate its impact — for everyone.”



THINKING OUTSIDE THE BLACK BOX WITH AI-ASSISTED EPIDEMIOLOGICAL MODELS

JON ZELNER, UNIVERSITY OF MICHIGAN

Jon Zelner's journey into AI-enhanced epidemiological modeling began long before the pandemic turned modelers into household names. A professor of epidemiology, Zelner has long focused on the spatial and social dynamics of infectious disease transmission, from MRSA to tuberculosis to COVID-19. But as data complexity has exploded and public trust in science has frayed, their mission has become more urgent: build models that not only predict and explain disease spread, but also earn the trust of those affected.

At its core, Zelner's work is about turning raw data into causal insight. His lab, EpiBayes, uses hierarchical Bayesian models to integrate disparate data streams, incidence rates, pathogen genomes, spatial data, and social determinants into a unified understanding of disease dynamics. For many years, this meant handcrafting interpretable models tailored to public health questions. But increasingly, Zelner's team is turning to AI to help scale, smooth, and synthesize.

That shift is less about replacing traditional epidemiological methods and more about extending them. With funding from a MIDAS PODS award, Zelner's team is exploring neural posterior estimation, a method that replaces brittle, hand-coded likelihood functions with neural networks trained on simulation outputs. The goal: enable faster, more flexible model fitting, especially when dealing with complex, multidimensional data.

The stakes are high. In public health, missteps in modeling can erode trust for years. Zelner recalls the damage done by oversimplified or poorly explained COVID-19 forecasts — black-box models that made strong claims with limited data and undermined public confidence in science. He sees trustworthy modeling as a two-part challenge: first, ensuring that scientists themselves believe their models are valid; and second, building public trust in the decisions that flow from those models.

One standout example of their team's work is a study of racial disparities in early COVID-19 deaths in Michigan. While narratives at the time emphasized comorbidities, Zelner's group used counterfactual modeling to show that infection timing, who was exposed earliest, during the most lethal phase of the outbreak, explained much of the disparity. Their models, grounded in plain-language results and real-world relevance, provided clarity without requiring black-box AI.

Still, as Zelner points out, some questions exceed the capacity of traditional models, especially when combining spatial, genomic, and social data in the face of emerging threats. His team is currently extending their neural inference methods to model MRSA transmission in Chicago, incorporating genomic sequences, incarceration data, and neighborhood-level social variables. The goal is to ask and answer mechanistic questions about inequality and disease, not just generate predictions.

But Zelner is clear: AI in their workflow is a **means**, not an end. It's a scaffolding tool, a way to simulate quickly, estimate robustly, and fit complex models, all in service of clarity, accuracy, and, above all, trust.

"We want AI to help us model better, faster, and more transparently — not replace the human judgment that public health depends on."



BUILDING AI FOR BLADDER CANCER SURVIVAL PREDICTION: A MULTIDISCIPLINARY JOURNEY

LUBOMIR HADJIYSKI, UNIVERSITY OF MICHIGAN

For Lubomir Hadjiyski, applying AI to medicine isn't just about algorithms; it's about building teams, defining the right clinical questions, and carefully navigating the realities of limited data. His group's work on bladder cancer survival prediction is a case study in how interdisciplinary collaboration fuels innovation.

Bladder cancer is a complex disease, ranking 10th globally and 4th among men in the U.S. While early-stage cases have a promising five-year survival rate near 90%, more advanced stages see survival drop sharply. Improving prediction models can help clinicians tailor treatments, counsel patients, and better allocate resources. But building accurate models in medicine, especially with small datasets, presents unique challenges.

The project began with a simple question from clinical collaborators: could AI help assess treatment response and predict survival from radiology images? Over time, the collaboration expanded, linking radiologists, oncologists, engineers, and data scientists into a cohesive team.

Their goal was ambitious: use information from CT urography images, clinical records, and advanced modeling techniques to predict five-year survival outcomes. The journey required integrating several modeling approaches:

- **Clinical Descriptors:** Traditional variables like tumor stage, lymphovascular invasion, and treatment history were modeled using nomograms and linear predictors.
- **Radiomics Features:** Image-based features capturing texture, morphology, and grayscale properties were extracted and modeled using neural networks.
- **Deep Learning Descriptors:** A hybrid method combined pre- and post-treatment imaging into a single input, allowing convolutional neural networks to learn subtle differences related to treatment response.
- **Large Language Models (LLMs):** More recently, Hadjiyski's team explored replacing manual clinical data extraction with LLMs (such as GPT-3.5/4.0), demonstrating that AI could retrieve structured information from medical reports almost as accurately as human experts.

Throughout the project, Hadjiyski emphasized practices critical for medical AI success:

- Careful **dataset partitioning** (training, validation, and testing) from the beginning to avoid data leakage.
- **Feature selection** to reduce overfitting, especially with small patient cohorts (163 patients in this case).
- **Combining modalities** — clinical, radiomic, and deep-learning features — into integrated models for better predictive power.
- Cautious **exploration of data augmentation** techniques, noting the limitations of GANs in generating realistic new samples for small medical datasets.

The combined models consistently outperformed models based on single data types alone, achieving an AUC (area under the curve) as high as 0.87 in survival prediction tasks.

Yet beyond the technical achievements, Hadjiyski highlighted an even deeper lesson: effective AI for medicine depends on **trust, collaboration, and shared vision** between technical experts and clinicians, from defining clinically meaningful tasks to validating models with real-world constraints.

Their team's work demonstrates that with careful design and interdisciplinary partnership, even relatively small datasets can drive advances that improve patient care, and that AI, when grounded in domain expertise, can extend the reach of medical science.

"AI models alone aren't enough — it's the human network behind them that brings meaningful medical impact."

2025

AI JOURNEYS

M | MIDAS MICHIGAN INSTITUTE
FOR DATA & AI IN SOCIETY
UNIVERSITY OF MICHIGAN