Is AI actually protecting at-risk communities from disasters? The honest answer is yes - but not where the hype is looking. The difference is being made not by clever models, but by the unglamorous plumbing that connects them to a village official's next move.

The real success of AI lies not in predicting disasters, but in enabling communities to act before they strike.

AI is transforming disaster management from reactive relief to proactive protection through hyperlocal, data-driven decision-making.

FROM PREDICTION TO PROTECTION: AI'S QUIET REVOLUTION IN DISASTER RESILIENCE

DR. SANJAY KUMAR SRIVASTAVA

Two villages, one revelation

In a village in Tamil Nadu's Virudhunagar district, a local administrator sat through a training workshop last year and stared at a map she had never seen before. It was generated by an AI platform, and it flagged the girls' school in her village as sitting inside a heat micro-hotspot - a pocket of elevated temperatures invisible to any district-level average. She already knew the children kept falling ill. She had never known why. According to a YourStory report on the programme, her response, once the map connected effect to cause, was immediate.

Fifteen hundred kilometres north, in Bahraich district of Uttar Pradesh, the same platform explained a different mystery: why stock at a small mandi kept drying up. The market sat in a high-risk zone. The community knew the symptom intimately; the data supplied the diagnosis.

> Communities have always known the symptoms of climate risks—AI is helping them understand the cause

These are small stories. That is precisely the point. The question of whether artificial intelligence is making a difference in protecting at-risk communities is usually answered with satellite constellations and supercomputers. But the real test of the technology is whether it changes what a panchayat officer, a health worker, or a grid operator does on a Tuesday afternoon before the water arrives. Increasingly - and quietly - it does.

> Hyperlocal AI risk maps are revealing hidden climate threats that traditional district-level assessments often miss.

The question behind the question

First, the uncomfortable arithmetic. The decades ahead will be riskier than the ones behind us. As a new conceptual

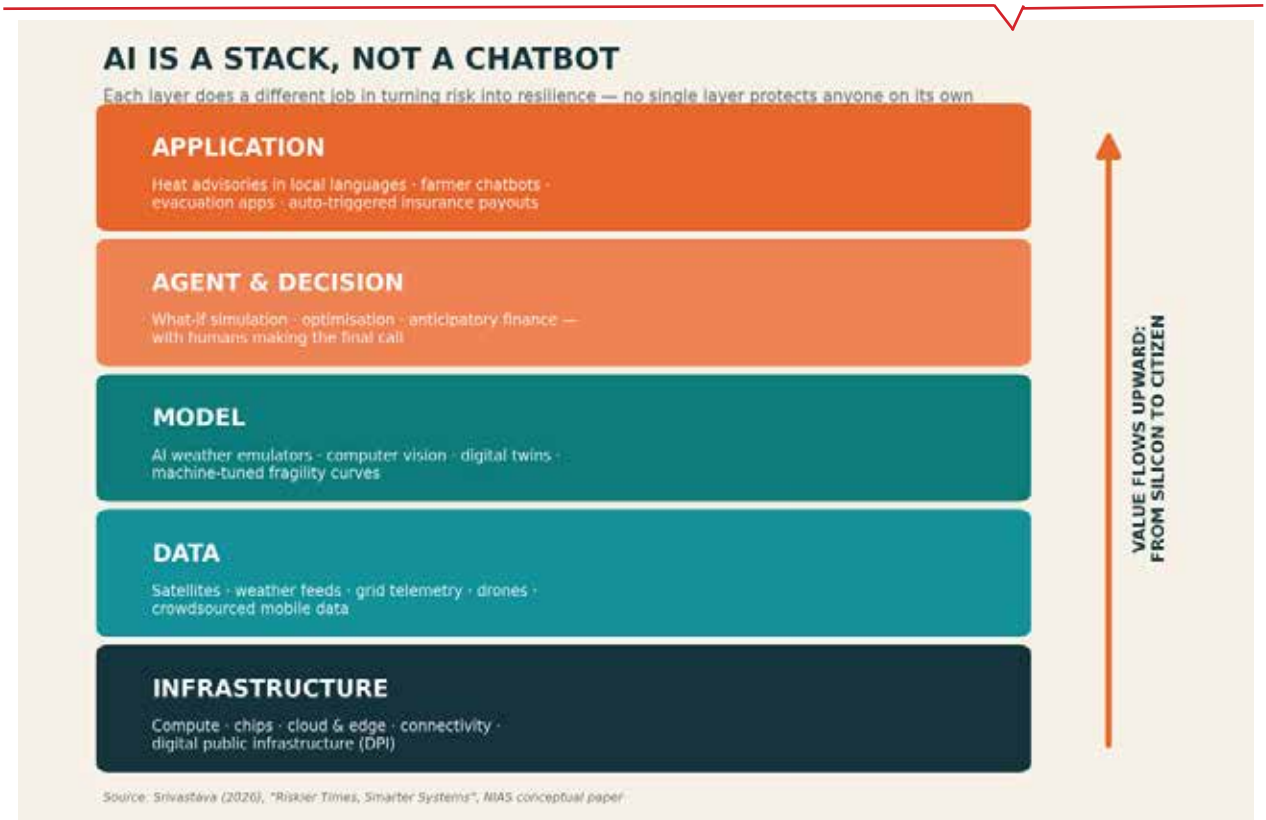

Figure 1 - The AI stack read as a resilience stack: value flows upward from silicon to citizen.

paper from the National Institute of Advanced Studies (NIAS) in Bengaluru argues, the excess heat trapped by greenhouse gases does not stay put: it lengthens heatwaves, destabilises the monsoon, energises cyclones, gnaws at Himalayan glaciers whose meltwater sustains nearly two billion people, and lifts the platform from which storm surges attack. These are not four problems but, in the paper's memorable phrase, one problem wearing four faces.

And the burden is not shared fairly. The countries that contributed least to emissions face the greatest exposure with the thinnest fiscal buffers. India sits at the sharpest edge: research cited in the paper suggests dangerous heat stress will become routine across South Asia even at 1.5°C of warming, and modelling points to a 30–40 per cent loss in outdoor work performance across India by century's end under high-emission pathways. For a nation of farmers, construction workers and street vendors, that is a direct tax on livelihoods. So the stakes of the AI question are not academic. If the technology cannot help here, its promise is decorative.

Stop thinking chatbots. Start thinking plumbing.

Public conversation has shrunk AI down to chatbots, chip wars and compute races. The NIAS paper offers a correction that should reframe the entire debate: AI is a stack of layers - infrastructure, data, models, decision agents, and finally the applications people actually touch - and each layer does a different job in building resilience. No single layer protects anyone. The gains come from stitching them into public systems.

India itself is the proof that such stitching works, even before AI enters the frame. The 1999 Super Cyclone killed close to ten thousand people in Odisha; storms of similar

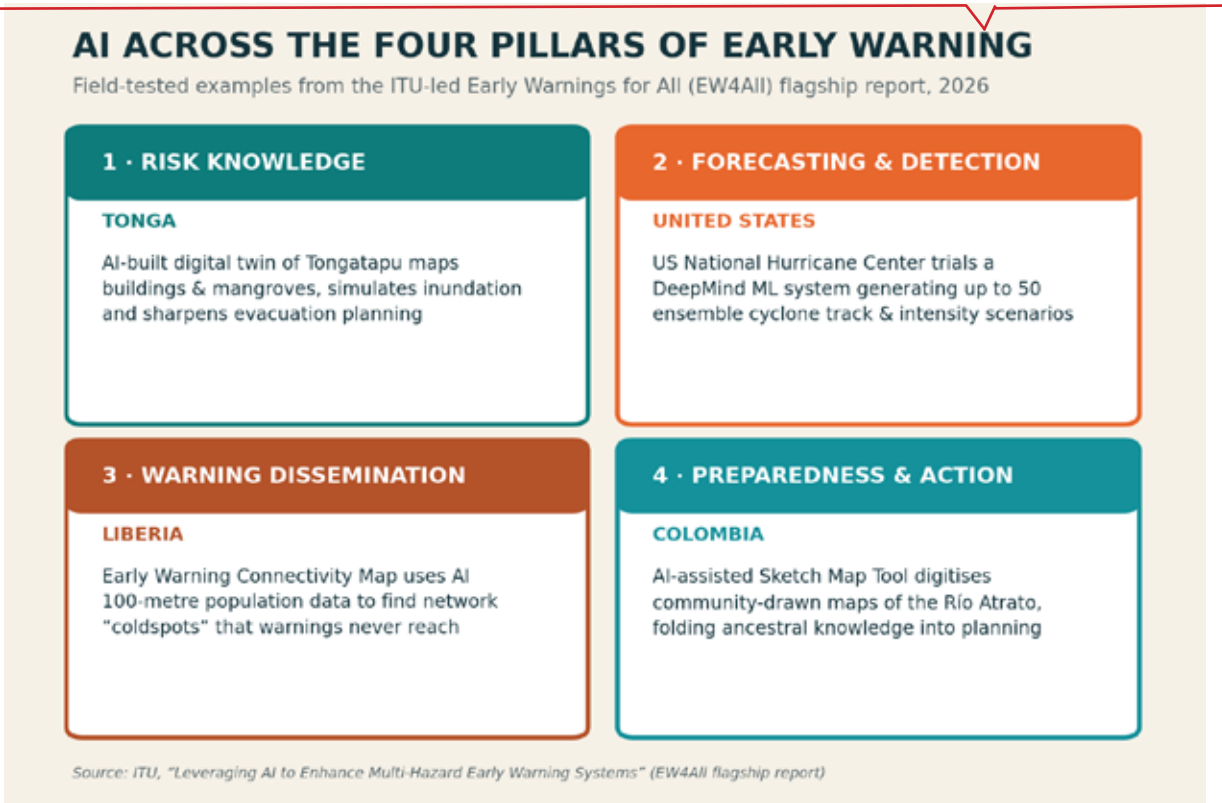

Figure 2 - What AI is already doing across the four pillars of multi-hazard early warning.

ferocity today kill a few dozen, because the state wired early warning into evacuation drills, shelters and social protection. Bangladesh's AI-supported flood forecasting, the paper notes, cut deaths not because its models were the world's best, but because forecasts were plugged directly into evacuation routines and safety nets that pay out before the water rises. AI delivers when embedded in preparedness. Treated as a stand-alone innovation, it decorates.

The global evidence is in

If you want field-tested receipts rather than conference-stage promises, the flagship report of the UN's Early Warnings for All initiative - led by the International Telecommunication Union and released this July - is the place to look. Its context alone is striking: 108 countries now report some multi-hazard early warning capacity, more than double the 2015 figure, yet many communities in the poorest countries remain outside the reach of a life-saving alert. The report maps, pillar by pillar, where AI is already closing those gaps.

The examples are refreshingly unglamorous. In Tonga, AI-driven asset mapping and a digital twin of Tongatapu island - down to counting coconut trees and tracking mangrove loss - now feed realistic inundation scenarios into local evacuation planning. The US National Hurricane Center is trialling a DeepMind machine-learning system that spins up as many as fifty ensemble scenarios for a cyclone's track and intensity. In Liberia, an AI-built, 100-metre-resolution population map was crossed with cell-tower data to expose connectivity "coldspots" - flood-prone places where a warning, however accurate, would simply never arrive. And in Colombia's Río Atrato basin, an AI-assisted sketch-mapping tool digitises maps drawn by hand in

Figure 3 - The difference, in numbers: what AI-assisted early warning and planning already deliver.

community workshops, folding ancestral knowledge into formal conservation planning rather than steamrolling it. Notice what unites these cases: none of them is about prediction alone. They are about knowing what is at risk, whom a warning can physically reach, and what a community does next. That is the four-pillar logic of early warning — and it is exactly where AI's pattern-hungry strengths meet the data deserts of the Global South.

India's last mile: 75 villages and counting

Which brings us back to Virudhunagar. The platform behind that heat map is Resilience360, built by the Indian startup Resilience AI, whose founder Samhita R. watched Cyclone Yas in 2021 and concluded that the deadliest element of a disaster is often the disorder and panic that follow it - including the conventional reflex of turning any available school into a relief centre, when many schools sit squarely in high-risk zones.

Backed by the United Nations and UNICEF, and working with the nonprofit RedR India, the company has now deployed AI-generated disaster management plans across 75 gram panchayats in Bahraich, Begusarai and Virudhunagar - flood country, quake country and heat country respectively. The numbers deserve attention. A hyperlocal, multi-hazard risk profile for any location in 24 hours. Village disaster plans that took government agencies three to eight months now emerge as working drafts in 72 hours, with human experts refining them. Every plan element converted into an action item with an owner and a deadline. One hundred and fifty grassroots operators trained to run it all - administrators, not climate scientists.

The deeper insight from the programme is humbling for technologists. These communities never needed to be told what was happening; they lived it daily. What they lacked was the scientific evidence explaining why - the link between a school's location and its sick children - without which the cycle simply continues. And

the work is already bending policy: the Ministry of Panchayati Raj has folded disaster management into its community-based risk reduction framework, and village-level databases are being designated as AI-enabled last-mile disaster records, so that villages hold their own risk data rather than depending on district averages.

Where the scepticism belongs

So, is AI making a difference? Yes - measurably, and in the places that matter most. But honesty requires marking where the job remains undone, and here all three sources converge on the same warnings. First, disasters cascade, and most risk tools still don't. Winter Storm Uri in Texas, Fukushima, the Odisha super cyclone - in each case the chain of failures, power to water to telecoms to hospitals, decided how bad things got. The NIAS paper's answer is the "GeoAI Guardian": a live model that, in one Odisha simulation, traces a single 220kV substation failure to six distribution zones within two hours, drained telecom batteries by hour eight, and four hospital ICUs dark within a day - and converts that foresight into a ranked list of orders. Move this repair crew. Fuel that generator. This should be a regulatory requirement, not a research project.

Second, governance is not garnish. The ITU report is emphatic that early warning is not really about technology but about the capacity to shape it - through rules guaranteeing transparency, human oversight and clear protocols for AI-supported decisions. Third, sovereignty matters: whether the Global South owns or merely rents the compute, data and models beneath its resilience will decide whether its adaptation intelligence is home-grown or imported - one more dependency in a world already full of them. Models trained on temperate baselines will misread the monsoon.

The verdict

The lazy answer to this essay's question is that AI predicts disasters better. True, and beside the point. The defensible answer is that AI, wired through the full stack into early warning, village planning, grid operations and social protection, is moving whole societies from prediction to protection - from paying for relief afterwards to acting before. Odisha proved a generation ago that systems, not luck, decide death tolls. The evidence from Tonga to Liberia to Bahraich suggests AI can power a leap of the same size, this time against the continuous, cascading risk of a hotter world.

The floods will keep coming; the physics guarantees it. But as the NIAS paper puts it, riskier times can no longer be avoided - unmanaged risk still can be. The stack exists. The frameworks exist. What remains, in every capital and every panchayat, is the decision to use them.

SOURCES

Srivastava, S. K. (2026). "Riskier Times, Smarter Systems: Leapfrogging Disaster Resilience and Climate Adaptation through the Full AI Stack." Conceptual paper, National Institute of Advanced Studies, Bengaluru.

ITU (2026). "Leveraging AI to Enhance Multi-Hazard Early Warning Systems: A practical resource to support Early Warnings for All." Flagship report of the EW4All initiative. itu.int

Balakrishnan, R. (2026). "How Resilience AI is using hyperlocal risk data to help villages prepare for climate disasters." YourStory SocialStory, 17 June 2026.

***

(Author is associated with National Institute of Advanced Studies, Indian Institute of Science Campus, Bengaluru as Chair Professor)