

AINA

AI and Analytics Magazine | Volume 7.0 | August 2026



Are We Ready?



From the Team

We live in a remarkable moment.

For perhaps the first time, millions of people are witnessing a technological revolution unfold in real time. Artificial intelligence is no longer confined to research laboratories. It writes, creates, predicts, reasons, and participates in decisions that shape our work, institutions, and everyday lives.

It is impossible to engage with AI without feeling a sense of excitement. New discoveries emerge every week. Problems once considered impossible suddenly appear solvable. Students, entrepreneurs, researchers, artists, and leaders now possess tools that were unimaginable only a few years ago. The pace of innovation can feel exhilarating.

Yet excitement is only part of the story.

AI has also introduced uncertainty. Questions surrounding employment, privacy, bias, accountability, and control have become central to public discourse. As these systems grow more capable, many of us find ourselves balancing optimism with anxiety. We wonder what AI will change, who it will benefit, and how society should guide its development.

This tension between possibility and responsibility lies at the heart of this issue of AINA.

Rather than viewing excitement and concern as opposing forces, we believe they can coexist. Anxiety often reflects the significance of what is unfolding before us. The future of AI matters because its potential is so immense. The conversations around governance, ethics, regulation, and accountability are signs that society is taking this technology seriously.

As students and AI enthusiasts, we recognise that how our generation builds, deploys, regulates, and governs these systems will shape the future.

This issue explores those questions. Through diverse perspectives and thoughtful analysis, we hope to encourage not only understanding but also responsibility.

The future of AI may be uncertain, but it is undeniably extraordinary. And despite the anxieties that accompany it, it remains one of the most exciting times to learn, to question, and to imagine what comes next.

Happy reading,
Team AINA



From the PGDBA Chairperson, IIM Calcutta

The Post Graduate Diploma in Business Analytics (PGDBA) is a unique full-time, residential programme jointly offered by three globally renowned institutions—the Indian Institute of Management Calcutta (IIM Calcutta), the Indian Statistical Institute (ISI), Kolkata, and the Indian Institute of Technology Kharagpur (IIT Kharagpur).

By bringing together the strengths of management, statistics, and technology, from the academic acumen of these three institutes, the programme has established itself as one of India's most distinctive offerings in business analytics.

The relevance of its curriculum is further heightened in the era of artificial intelligence (AI), as the courses provide comprehensive training in data analytics and equip students with data-driven leadership capabilities.

The programme employs a rigorous and competitive selection process to maintain a high-quality cohort of motivated learners, ensuring a well-balanced mix of experienced professionals and recent graduates. More importantly, the six month industry internship distinguishes PGDBA from most other programmes, enabling students to gain invaluable practical experience and emerge industry-ready.

As a result, graduates possess the critical thinking and analytical skills required to develop data-driven solutions to complex business problems using appropriate tools from data science.

The significance of this unique tri-institute programme continues to grow as AI-driven industrial transformations increase the global demand for professionals with expertise spanning business, analytics, and technology.

I am confident that our students will continue to excel in an AI-driven business world, drawing on the unique blend of expertise in management, statistics, and technology that defines the PGDBA programme. Just as our alumni have consistently made us proud by reaching new professional heights and creating impact across industries, I am certain that each new cohort will carry forward this legacy of excellence.

Dr. Abhipsa Pal

PGDBA Chairperson, IIM Calcutta



From the PGDBA Coordinator, IIT Kharagpur

We are living in an extraordinary epoch of human history where artificial intelligence is no longer a silent catalyst behind the scenes but the active architect of our global future.

From reshaping business models and transforming lifestyle habits to redefining the core operations of our legal, defence, and public policy frameworks, AI is fundamentally renovating the world market.

Yet, as we push the boundaries of what these technologies can achieve, we are confronted with an even more vital imperative: AI governance. Building intelligent systems is a technological triumph; ensuring they are ethical, transparent, accountable, and aligned with human values is a societal necessity.

Addressing a challenge of this magnitude demands a paradigm shift—moving away from siloed thinking toward a radically multidisciplinary approach.

This is where the unique essence of the Post Graduate Diploma in Business Analytics (PGDBA) truly shines. By amalgamating the technological grit of IIT Kharagpur, the statistical and mathematical rigour of the Indian Statistical Institute (ISI), and the strategic business acumen of the Indian Institute of Management (IIM) Calcutta, the tri-institute partnership creates a rare and powerful ecosystem.

It provides our students with a diverse opportunity to excel, nurturing them into leaders who do not just analyse data but contextualise it. They learn to view AI not merely as an engineering tool or a commercial asset but as a socio-technical construct that requires responsible stewardship.

As AINA 7.0 explores the critical theme of AI Governance, I congratulate the editorial team for curating a platform that addresses the most pressing conversation of our times. I wish our bright minds the very best as they step into the global arena, equipped to lead this digital renaissance with wisdom, integrity, and foresight.

Dr. Rudra Prakash Pradhan

PGDBA Coordinator, IIT Kharagpur



From the PGDBA Coordinator, ISI Kolkata

The Post Graduate Diploma in Business Analytics (PGDBA) is a tri-institutional, full-time, residential two-year programme conducted by three premier institutes (IIM Calcutta, ISI and IIT Kharagpur) in our country.

This programme presents a unique opportunity for students to receive world-class education and training in the field of business analytics. From these three institutes, students can benefit from learning interdisciplinary subjects, like statistical and machine learning theories for analytics from ISI, some technological training on analytics from IIT Kharagpur, and application of analytics from IIM Calcutta.

A major highlight of the PGDBA is the six-month industry internship, where students apply their knowledge to real analytics projects. This hands-on experience not only strengthens their technical and analytical skills but also provides critical industry exposure, preparing them to take on leadership roles in analytics-driven organisations.

At ISI, we take pride in our rich legacy of excellence in education and research.


Renowned globally for our contributions to statistics, mathematics, and computer science, we bring deep academic rigour to the PGDBA curriculum. Our faculty members are leading experts in mathematics, statistics, computer science and other interdisciplinary areas of research and students will get a high-quality learning experience from our faculty members.

We actively encourage students to engage in physical activities, social initiatives, and interactions with faculty and researchers, fostering both personal and intellectual growth. We are confident that the PGDBA will be a transformative journey for our students.

I look forward to witnessing their growth, contributions, and success in the dynamic world of business analytics.

Dr. Dibakar Ghosh

PGDBA Coordinator, ISI Kolkata



Copyright © 2026 PGDBA Publishing Team
All rights reserved.

No part of this publication may be reproduced, stored in a retrieval system, transmitted, or distributed in any form or by any electronic or mechanical means, including photocopying or recording, without prior written permission of the publisher, except for brief quotations used in reviews, research, or other non-commercial purposes as permitted under applicable copyright law.

The views and opinions expressed in the articles published in AINA are those of the respective authors and do not necessarily reflect the views of the Post Graduate Diploma in Business Analytics (PGDBA), the Indian Institute of Management Calcutta (IIM Calcutta), the Indian Institute of Technology Kharagpur (IIT Kharagpur), the Indian Statistical Institute Kolkata (ISI Kolkata), or the Editorial Board.

Every effort has been made to ensure the accuracy of the information contained in this publication. However, the publisher and editorial team accept no responsibility for any errors, omissions, or consequences arising from the use of the information presented herein. For any queries, please reach out to pgdba.aina@email.iimcal.ac.in

Published by
PGDBA Publishing Team
Post Graduate Diploma in Business Analytics (PGDBA)
Indian Institute of Management Calcutta
Indian Institute of Technology Kharagpur
Indian Statistical Institute Kolkata

Edition: AINA 7.0
Year of Publication: 2026

CONTENTS

- 10 AI... IN THIS ECONOMY?**
How AI Is Reshaping the global economy
- 15 CAUSAL ABSTRACTION**
Can we finally explain why AI makes the decisions it does?
- 20 BUILDING THE INFRASTRUCTURE FOR RESPONSIBLE AI**
Expert perspectives on governing AI before it governs us
- 30 THE CEO OF THE AI ERA**
The biggest management shift since the internet
- 33 CROSSWORD**
Test your AI vocabulary
- 34 SMARTER SIMULATIONS, BETTER ENGINEERING**
How AI is transforming computational engineering



AI... IN THIS ECONOMY?
History has repeatedly shown that technological revolutions create more prosperity than they destroy. Could AI be following the same path?

15

CAUSAL ABSTRACTION
As AI enters high-stakes decision-making, researchers are finally beginning to uncover how modern neural networks actually reason.



THE UNPROVABLE EQUATION
Intelligence can solve equations, but can it ever understand ethics, values, and human judgement?

40

YOUR LOCATION IS YOUR PASSWORD
Discover how banks are using geospatial intelligence to detect fraud before a transaction is even approved.



- 36 THE UNPROVABLE EQUATION**
Can intelligence ever understand morality?
- 39 INTRODUCING DASENSE**
Meet our new page for AI and analytics
- 40 YOUR LOCATION IS YOUR PASSWORD**
When your location becomes your strongest password
- 44 TEACHING NEURAL NETWORKS TO OBEY THE LAWS OF PHYSICS**
Can AI learn the laws of the physical world?
- 48 ALUMNUS SPEAKS**
Building a fintech startup from scratch
- 50 THE INVISIBLE FARMER**
Can AI finally improve farming productivity?
- 54 HOW INDIA CAN BECOME AN AI SUPERPOWER**
India's roadmap to AI leadership
- 56 THE BEST ANIMATOR IS STILL HUMAN**
Why human creativity still leads the way
- 58 TRILYTICS CONCLAVE**
Where India's brightest data scientists compete

AI... in This Economy?

RETHINKING ECONOMIC DEVELOPMENT IN THE AGE OF AI

Jaganatha Pandiyan



In the early days of commercial computing, when computers occupied entire rooms and memory was measured in kilobytes, Norbert Wiener, widely regarded as the father of cybernetics, saw something unsettling. “Let us remember,” he wrote, “that the automatic machine is the precise economic equivalent of slave labour. Any labour which competes with slave labour must accept the economic conditions of slave labour.” It was a bleak vision of technological progress, wherein machines would relentlessly undercut human workers, forcing wages downward until labour itself became economically expendable.

Seven decades later, Wiener’s warning has proved both prescient and incomplete. He correctly anticipated that automation would displace workers. What he failed to foresee was the extraordinary

expansion in the economy’s productive capacity. Computers did reduce the demand for entire professions such as clerks, telephone switchboard operators, bookkeepers, and even the human “computers” once employed to perform complex calculations. Yet they also created industries that had previously been unimaginable. Software, cybersecurity, e-commerce, cloud computing, digital media, and countless other sectors emerged from the computer revolution, generating trillions of dollars in economic value and employing hundreds of millions of people around the world. Automation did not simply replace labour. Instead, it expanded the frontier of what the economy was capable of producing.

That is the optimistic story we often tell ourselves. It is also an incomplete one.

Behind every technological revolution lies a quieter question that rarely appears in productivity statistics or GDP figures. Did the telephone switchboard operator become a semiconductor engineer? Did the bookkeeper become a software developer? In most cases, the answer was no.

This is the paradox at the heart of technological progress. Economies adapt far more easily than people do. While new industries eventually create more jobs than old ones destroy, the workers displaced by one revolution are rarely the ones who benefit most from the next. Aggregate prosperity can conceal individual hardship, and understanding that distinction is essential if we are to understand what AI might mean for the future of work.

History Repeats Itself

For the fourth time in modern history, society is convinced that technology is about to make human labour obsolete. The first was the Industrial Revolution. The second came with electrification and mass production. The third arrived with the computer revolution. Today, artificial intelligence has become the latest chapter in

that story. Each wave has displaced workers and sparked fears of permanent unemployment. Each time, those fears have been both justified and overstated.

Long before electronic computers existed, organisations such as NASA employed teams of “human computers” to perform complex calculations by hand. That profession has disappeared but that does not mean mathematics became less valuable. Rather, what it meant to be a mathematician changed. Once computation ceased to be scarce, the frontier shifted from arithmetic to abstraction. AI appears to be driving a similar transition. Consider software development, where the value of engineers is rapidly moving away from writing code towards designing systems and exercising judgement.

History also suggests that technological revolutions rarely translate into immediate productivity gains. Breakthrough technologies require complementary changes in organisational structure and management practices before their full economic benefits become visible. Economists refer to this as the implementation lag.


Robert Solow famously remarked

in 1987 that “you can see the computer age everywhere but in the productivity statistics.” Businesses took years to reorganise around computers before productivity accelerated. AI appears to be following a similar trajectory. While early deployments have produced mixed financial results, economists such as Daron Acemoglu still project meaningful long-term productivity gains once organisations fully adapt.

John Maynard Keynes famously believed that this transformation would eventually solve humanity’s oldest economic problem. In his 1930 essay *Economic Possibilities for our Grandchildren*, Keynes wrote that “for the first time since his creation, man will be faced with his real, his permanent problem, how to use his freedom from pressing economic cares.” Keynes believed that an eightfold increase in productivity over the next century would allow societies to reduce labour demand dramatically. A three-hour shift, he suggested, might become the norm.

In one sense, Keynes was right. Productivity has increased far beyond what he believed necessary to sustain such a society. Yet his prediction about leisure proved strikingly wrong. Rather than working dramatically fewer hours, people invented new forms of work and are working longer hours than ever before. Capitalism did not replace labour with leisure. It replaced old work with different work.

If history is any guide, the labour market will adapt with the AI revolution. New industries will emerge and the economy will continue to expand. The adjustment, however, will be painful for many workers, particularly those entering the workforce today. More importantly, history may not provide the right blueprint for what comes next. This time, AI is not merely changing jobs. It is changing the very economic model through which countries develop.



Why the AI Revolution Is Not Like the Previous Ones

Every technological revolution has expanded the economic pie and ultimately raised living standards. It is therefore tempting to dismiss concerns about artificial intelligence as simply the latest version of an old fear. That may yet prove to be correct. The reason for caution is not that AI will eliminate work altogether. What we must be concerned about is that it may disrupt the economic model through which much of the developing world has achieved prosperity.

Since the end of the Second World War, the path to development has followed a consistent pattern. As wages and land costs rose across advanced economies, firms began relocating labour-intensive activities to countries where production

was cheaper. Economists explain this process using the principles of comparative advantage. Companies improved their profits while developing countries gained employment, foreign exchange, and, perhaps most importantly, access to technology and know-how. Over time, these countries climbed the value chain, from low-value assembly to higher-value manufacturing, then to services and eventually to innovation.

This model transformed entire nations. China’s manufacturing rise and India’s IT boom were both products of integrating into global value chains. Similar strategies have guided development across much of Asia, Africa and Latin America. It is difficult to overstate how successful this model has been in transforming the lives of billions of people.

Artificial intelligence threatens to alter that equation. The concern

is often misunderstood as being that jobs will disappear overnight. Labour markets have repeatedly adapted to technological change and they are likely to do so again. The deeper question is whether AI weakens the economic incentives that made offshoring attractive in the first place.

If a software engineer assisted by AI can accomplish the work that previously required several developers, the labour-cost advantage of offshoring begins to shrink. If routine customer support can be handled by AI systems operating from company headquarters, maintaining large overseas call centres becomes less compelling. In many industries, the decision revolves around speed, coordination, intellectual property and strategic control. Those considerations naturally favour bringing work closer to home.

This shift is already visible and should concern developing economies far more than the prospect of individual jobs being automated. If advanced economies no longer need to offshore work at the scale they once did, the traditional ladder of development becomes much harder to climb. The challenge is not confined to workers acquiring new skills. It extends to entire nations rethinking how economic development itself will take place in a world where the cheapest worker may no longer be a worker at all.

The Inequality Problem

Technological revolutions have never benefited everyone equally. Every major innovation creates new winners and new losers before society eventually adjusts. The difference today is that AI is arriving at a moment when inequality is already historically high across much of the world.

For decades, economists have debated whether returns to capital systematically outpace returns



Before electronic computers, teams of human computers at JPL performed the calculations that helped lay the foundation for the U.S. space programme.
Source: NASA/JPL-Caltech

to labour. Whatever one's position in that debate, the fact that wealth has become increasingly concentrated in many advanced economies is difficult to ignore. Those who own productive assets are often better positioned to benefit from technological change than those whose primary asset is their labour. AI has the potential to accelerate that dynamic by increasing the value of existing capital while reducing the scarcity of certain forms of skilled labour.

Yet capital is no longer the only source of advantage. The AI economy is built upon a set of strategic assets that are even more difficult to acquire: vast computational infrastructure, specialised semiconductor supply chains, world-class engineering talent, proprietary datasets and the financial resources required to train frontier models. These advantages are highly concentrated. A small number of technology companies possess the compute, capital, talent, and research capabilities necessary to develop foundation models at the frontier, while the rest of the world builds applications on top of technologies they do not own.

There is another reason inequality may widen. AI exhibits powerful increasing returns to scale. The organisations with the best models attract the most users. More users generate more revenue, more feedback and more data, enabling further improvements that attract even more users. Success compounds upon itself. Over time, this creates winner-takes-most markets in which a small number of firms capture a disproportionate share of the economic value. Unlike previous industries, where many competitors could coexist, frontier AI naturally rewards scale, making concentration a feature of the market rather than an accident.

The implications extend beyond companies to countries themselves. Nations that possess frontier AI capabilities stand to accumulate more wealth, attract

more talent and shape global technological standards. Those without such capabilities risk becoming permanent consumers of intelligence produced elsewhere. They will innovate around AI, build products with AI and create value using AI, but the underlying platforms, and the greatest share of the economic rents, will remain concentrated in the countries that own the foundational technology.

Rethinking Development in the Age of AI

The challenge with AI is that it may be altering something much deeper than the labour market. It is reshaping the economics of production itself. If comparative advantage is no longer determined primarily by labour costs, then many of the assumptions that have guided development policy for the past seventy years should be reconsidered.

The first priority should be recognising AI infrastructure as essential economic infrastructure. In the twentieth century, countries invested in roads, ports, electricity grids and telecommunications because private markets alone could not build the foundations of long-term growth. The twenty-first century demands an equivalent commitment to compute infrastructure, research institutions, data, and reliable access to advanced semiconductors. Countries that lack the ability to participate in the AI ecosystem will find themselves dependent on those that do.

Second, governments must rethink how the gains from AI are distributed. The objective should not be to slow technological progress, nor to protect every existing job from inevitable change. History suggests that such efforts rarely succeed. The challenge is ensuring that productivity gains translate into broadly shared prosperity rather than becoming concentrated among a small number of firms

and capital owners. That requires serious conversations about competition policy, taxation, ownership of AI-enabled capital and new mechanisms through which citizens can share in the wealth generated by automated economies.

Finally, developing economies must recognise that the strategy which served them so well during the era of globalisation may not be sufficient in the age of artificial intelligence. Competing primarily on low labour costs becomes less attractive when labour itself becomes a smaller component of production. Future competitiveness will depend on research, education, entrepreneurship, digital infrastructure, institutional quality and the ability to build.

The more important question is who gets a larger slice. If we continue to measure success solely by productivity growth and GDP, we may celebrate an economy that is becoming steadily more efficient while overlooking one that is becoming progressively less inclusive. The challenge before policymakers is therefore not simply to accelerate AI adoption, but to ensure that the institutions governing this revolution evolve as rapidly as the technology itself.

The choices made over the coming decade will shape not only who benefits from AI, but also whether the next era of technological progress strengthens or weakens the social contract on which prosperous societies depend.

AI will undoubtedly create extraordinary prosperity. Whether it also creates broad-based prosperity depends not on the technology itself, but on the institutions we build around it. History has shown that innovation is remarkably good at expanding the pie. Our challenge now is to ensure that the recipe for sharing it evolves just as quickly. ■

Causal Abstraction

WHAT HAPPENS INSIDE AN AI'S MIND?

Pragathi Gudar



AI systems now influence real-world decisions. They help decide who gets a loan, which patients receive urgent care, and what legal advice someone might receive. They have become remarkably capable, until they aren't. And when they fail, they often fail with complete confidence. The problem with modern AI isn't that it fails. Traditional software fails all the time. The problem is that when AI fails, we often can't explain why.

In 2024, a Canadian tribunal ruled against Air Canada after its customer service chatbot stated that a bereavement discount could be claimed retroactively. The airline argued it shouldn't be liable for an autonomous system's mistakes. The tribunal disagreed: if a system speaks for you, you're responsible for what it says. But the deeper issue was technical.

Air Canada couldn't identify where the mistake came from. There was no buggy function to repair. The response emerged

from statistical patterns distributed across millions of parameters, learned behaviour with no obvious source. As AI systems draft legal briefs, approve loans, process immigration requests, and influence medical decisions, this inscrutability becomes dangerous. In 2023, two New York lawyers were sanctioned for submitting court filings with fabricated citations generated by ChatGPT. The citations looked legitimate but didn't exist. What's more troubling is that no one could trace how the model produced those specific references. The gap between what AI can do and what we can explain keeps widening.

The Illusion of Insight

Researchers have spent years building tools to peer inside neural networks. Attention maps show which words a language model focuses on. Probing classifiers test whether layers encode grammar or facts. Feature attribution estimates what influenced a prediction. These techniques have revealed genuine structure. Large language models aren't undifferentiated statistical engines; they

develop organised components that copy information, maintain grammatical agreement, and perform arithmetic. The notion that neural networks are pure chaos is no longer defensible. But visibility isn't understanding.

Most interpretability methods are observational. They show which parts of a model are active during a prediction, but they can't tell us whether those parts actually caused the decision. A well-known example came from a medical AI trained to detect pneumonia from chest X-rays. Early results looked impressive. But when researchers deployed it at another hospital, performance collapsed. The investigation revealed something surprising. The model wasn't relying only on lung abnormalities. It had also learned that portable X-ray machines were strongly associated with pneumonia in the training data. Patients healthy enough to walk received scans in the radiology department, while severely ill patients were often imaged at the bedside using portable machines. Because portable machines happened to correlate with severe illness, the model partially treated their presence as evidence of pneumonia. Attention


maps still highlighted the lungs because the model genuinely examined them. But attention alone couldn't reveal that the final prediction also depended on the portable machine. Only intervention, i.e. removing or altering that information and observing the prediction change, revealed the hidden shortcut.

In science, explanation requires intervention. To prove a variable matters, you must change it and observe whether outcomes shift predictably. If modifying a component leaves output unchanged, what you observed was incidental. Much of interpretability stops short of this standard. It shows what's present inside models, not what governs their behaviour. Causal abstraction begins where observation ends.

The Experiment

When physicians suspect a medication causes side effects, observation alone isn't proof. Controlled trials are required: one group receives the drug, another gets a placebo. If symptom rates change as predicted, causation is established. AI researchers are beginning to run similar experiments inside neural networks. Why? Because observing what's inside a model isn't the same as understanding why it decides what it decides. We need to know which internal parts actually drive the decisions, which ones matter, and which are just noise.

The technical name is activation patching. The idea is simple: replace internal representations and see what happens next. If the output changes predictably, you've found something important. Consider a straightforward task. GPT-2 correctly completes: "The capital of France is ____" with Paris, and "The capital of Germany is ____" with Berlin. Traditional tools might show the model attends to "France" when predicting "Paris," or that certain neurons activate for European countries.



Replace an internal activation and measure its effect on the output. A systematic change provides evidence that the patched component contributes causally to the prediction.


But this doesn't answer the crucial question: does the model actually use its representation of "which country" to determine "which capital"? Maybe not. It might rely on a statistical pattern. "The capital of F____" often precedes "Paris" in training data. That produces correct answers without implementing a genuine country-to-capital mechanism.

Activation patching tests this directly. Researchers run both sentences through the model, recording internal activations, the numerical patterns representing processing at each layer. Then

comes the intervention: while processing the France sentence, they replace the internal representation of "France" with the one from the Germany example. The visible text stays unchanged. The input still says "France." But the output changes. The model now predicts Berlin. Nothing was altered in the input. No retraining occurred. Only an internal representation was swapped, and the output followed exactly. That's evidence of causation. The manipulated activation behaves like a program variable: change it, and the result changes predictably. That might seem like a toy exam-

ple. But the same approach works on systems that matter: the ones deciding who gets a loan, which patients receive priority care, or why a chatbot gives certain legal or financial advice.

By repeating this experiment many times, researchers are beginning to piece together how information moves through the model. Rather than treating the network as one giant black box, they can identify groups of internal components that appear to pass information from one stage of reasoning to the next.



This is the key shift: from “the model appears to know geography” to “here’s proof that it actually uses country information to determine the answer.” Instead of watching and inferring, researchers test hypotheses by altering components and observing predicted effects. One approach hopes models improve with data. The other treats them as experimental objects. That shift, from observation to intervention, makes causal abstraction science.

The Search Problem


Activation patching only works if researchers know where to look. How did they identify which layers to test without guessing? The challenge is scale. Modern AI models contain thousands of interacting components and testing every component combination would exceed atoms in the universe. Therefore, we need a smarter solution.

Distributed Alignment Search, or DAS, solves this problem elegantly. It works like the process of elimination. Researchers start with a hypothesis: what internal step does the model perform? Then they generate candidate locations where this might happen: early layers, middle layers, late layers? They test each candidate by intervention. If swapping components in the early layers changes the prediction 70% of the time, those layers probably matter. If they don’t, they move on to the next layer. The best candidates get refined.

Each candidate undergoes activation patching tests. Does swapping these components between examples change predictions? How consistently? Candidates producing reliable changes, say 70% accuracy, are retained. Failures are discarded. The best refine further: if the middle layers score well, which specific layer? If layer 8 beats layer 7, focus there. If certain heads matter more, isolate them. This iterates. Initial guesses

might barely beat random chance. But each round narrows it down until eventually, precision sharpens and researchers can identify which internal components matter most and how strongly they control the output. When the search stops improving, they have their answer.

Researchers can trace where country information gets processed, which internal components transform it, and where it connects to the final answer. Swapping these components reliably changes the output. The primary computational pathway is visible. No pathway maps 100% of behaviour. Some decisions emerge from multiple routes, or involve noise. But the dominant mechanism is knowable, testable, and most importantly, fixable. Computational cost remains substantial. Mapping these mechanisms requires many rounds of testing. Scaling to much larger models is an active challenge, though researchers are developing more efficient methods. Still, the principle holds: mechanisms aren’t hidden in black boxes. Instead, they’re discoverable through systematic search and experimental verification.



From Lab to Ledger

In 2023, a European bank used activation patching to audit a credit model showing bias. Applicants from certain postal codes had approval rates eighteen per cent lower than others, even with identical finances.

Standard fixes failed. Removing postal codes didn’t help; the network inferred location from correlated features. Debiasing reduced gaps modestly but left mechanisms unclear. Regulators wanted to know whether a causal pathway from geography to decision existed. Auditors applied the same logic. They found two applications with identical financial profiles differing only in postal code, one from a majority-immigrant neighbourhood, and one not. Processed normally: 72 versus 83 scores.

Then they patched: reprocessing the first application with postal code activations from the second. The score jumped to 81. That eleven-point change from swapping twenty neural dimensions proved specific components caused the disparity.

The institution acted on this find-

ing by modifying weights along the problematic pathway while preserving legitimate risk assessment. The approval gap dropped from eighteen to three per cent but accuracy declined just over one per cent.

Germany’s BaFin accepted this as compliance evidence, a significant precedent. Accountability through mechanistic proof: identifying, testing, and correcting the causal pathway.

What This Cannot Do

While causal abstraction might be a leap towards more explainable AI, it faces real constraints. Computational expense is substantial; mapping mechanisms require tens of thousands of tests, and scaling to trillion-parameter models is a huge challenge. Not all capabilities localise neatly. Some behaviours emerge from diffuse interactions. One study found “sarcasm understanding” involved 47 components across 8 layers with only 63% accuracy, genuinely distributed, resisting clean decomposition. Mechanisms reliable in simple cases degrade with complexity. Models maintain multiple strategies, deploying different circuits for different inputs. Even successful interventions achieve perhaps 85% accuracy, leaving 15% unexplained through secondary mechanisms or errors. Causal abstraction maps dominant pathways without guaranteeing completeness. These limitations constrain application but don’t invalidate the approach. They clarify deliverables and where work must focus.

The Structural Turn

For more than a decade, AI progress was driven by scale. Larger models, more data and greater compute consistently produced better performance. The strategy succeeded remarkably. But capability and understanding have not advanced together. Models now

“

It doesn’t matter how beautiful your theory is, it doesn’t matter how smart you are. If it doesn’t agree with experiment, it’s wrong.

– Richard P. Feynman

generate code, reason through complex problems and assist with high-stakes decisions, yet the computations that produce these abilities often remain opaque. Bigger models are not inherently more interpretable. If anything, they demand even greater scrutiny.


That imbalance is beginning to reshape AI research. The next frontier is no longer simply building more capable systems but understanding the mechanisms that already exist inside them. Causal abstraction is part of this broader shift. Rather than treating neural networks as statistical black boxes, researchers are beginning to study them as engineered systems whose internal computations can be isolated, tested and verified. As models continue to grow in capability, our methods for understanding them must grow in rigour.

This shift also reflects a deeper realization. Correlations alone can produce remarkably capable systems, but they cannot tell us why a model reaches a particular conclusion or whether that reasoning will remain reliable when circumstances change. Understanding requires more than observing patterns. It requires uncovering the mechanisms that produce them. That philosophy

has guided scientific discovery for centuries, and it is now beginning to shape the next generation of AI research.

This evolution mirrors the trajectory of every mature engineering discipline. Early aviation advanced by building aircraft that flew farther and faster. Eventually, the emphasis shifted toward understanding fatigue, failure modes and safety systems. Civil engineering followed a similar path, moving from constructing larger structures to developing rigorous methods for testing materials, predicting failure and validating designs. AI is approaching a comparable inflection point. As models become foundational infrastructure rather than experimental software, reliability will depend not only on what they can accomplish, but on whether their internal reasoning can be systematically examined, challenged and improved.

The future of trustworthy AI will not belong to systems that merely produce the right answers. It will belong to systems whose reasoning can be tested, understood and improved. The next breakthrough in artificial intelligence may not be another leap in scale. It may be the moment when understanding becomes as fundamental as capability itself. ■



Building the Infrastructure for **Responsible AI**

Technology changes quickly but institutions do not.

That mismatch has defined every major technological revolution. New technologies can emerge within years; laws, norms, and organisations often take decades to catch up. Artificial intelligence is accelerating this cycle unlike anything before it.

The challenge before us is therefore larger than developing smarter models. It is about rethinking our institutions so that they remain effective in an AI-driven world.

In the conversations that follow, leaders from four different domains discuss not just how AI will transform their professions but also how the institutions they serve must evolve alongside it.

AI in Defence



Lt Col Kapil Singh Shekhawat (Retd)

Former Joint Director,
Indian Army

Q Let's begin with a simple thought experiment. Imagine two possible worlds side by side. In one, every country is racing to integrate AI into its military. In the other, all countries have collectively agreed to keep AI out of the armed forces entirely. Which world would you rather live in, and why?

We have once again reached that point where countries start debating on whether a new technology should be incorporated into the military or not. The last time this happened was with nuclear weapons and we have seen how countries that don't have a nuclear

arsenal are always on the receiving side when it comes to a war.

As a soldier, I would naturally prefer a safer world, one where militaries collectively decide to keep AI out of warfare. But as a realist, I do not believe that world is feasible. History teaches us an important lesson that whenever militaries fail to adapt to transformative technologies, be it gunpowder, tanks or nuclear weapons, they are at a strategic disadvantage. Those who adopt critical technologies early often gain a decisive advantage. Even if, hypothetically, countries agree to keep AI out of the military, how do we trust that? There have been instances where countries have se-

cretely developed nuclear weapons. There is nothing stopping them from doing the same with AI.

For India, it's clear that we cannot afford to sit out the AI revolution. We must understand it, develop capabilities responsibly, and build expertise before others do. At the same time, adopting AI should never mean outsourcing responsibility. Technology continuously evolves, but accountability should never be outsourced. Accountability should always remain human and from that worldview, I would go ahead not just with AI, but any new advancements in technology.

Q We are moving towards a world of autonomous weapon systems. Guns, jets, tanks and missiles are all increasingly becoming autonomous. Do you think that the advantages of autonomous systems outweigh the risks they come with? How do you feel about this change?

The advantages of autonomous systems are undeniable. Humans have limitations, be it physical, psychological, or emotional. Machines, on the other hand, do not tire, become emotionally overwhelmed, or lose concentration. They can process enormous volumes of information at speeds no human can ever match. However, military operations are rarely black and white. During my years in service, I learned that battlefields are shaped by uncertainty, incomplete information, and rapidly evolving circumstances. Soldiers often operate in ambiguous situations where judgement, and not just data, determines outcomes. That is the fundamental distinction: machines operate

on data, while soldiers operate on judgement. Every new piece of information changes the battlefield. Context evolves continuously, and experienced soldiers constantly reassess situations as fresh intelligence arrives. That ability to interpret ambiguity, rather than simply process inputs, remains uniquely human and is difficult to encode into any autonomous system.

Autonomous systems can identify threats and support soldiers. But when it comes to lethal force, meaningful human judgement and control are required. A machine may be technically capable of pulling the trigger, but that does not mean it should be granted the authority to do so.

Q Defence is often viewed through the lens of battlefield strategy, but logistics and timely resource allocation are equally critical. How do you see AI and data analytics reshaping sourcing and procurement?

Be it aviation or ammunition, data analytics plays a crucial role in sourcing and procurement. I have spent many years handling procurement at various echelons of the Indian army, including at Army Headquarters. One lesson became very clear: battles may be fought on the battlefield, but they are sustained through logistics.

AI can be transformative here. Predictive analytics can help anticipate equipment failures before they happen, identify inventory shortages, and optimise procurement cycles. Imagine knowing six months in advance which helicopter component is likely to fail, where ammunition will need replenishing, or whether a supply-chain disruption could affect operations.

Even if predictive systems are not perfect, anticipating 70 to 90 per cent of future contingencies can give a major strategic advantage. Wars are not necessarily won by

“

A machine may be technically capable of pulling the trigger but that does not mean it should be granted the authority to do so.

the stronger side; they are often won by the side that is the best prepared for contingencies. Military success often comes down to delivering the right resource, at the right place, at the right time. Traditionally, logistics has been reactive. AI offers the possibility of shifting toward predictive logistics, where military supply disruptions are anticipated and mitigated before they escalate.

Q As AI becomes embedded across defence functions, the skills required from a military officer will inevitably change. How should military training programmes evolve to prepare future leaders?

The military leaders of tomorrow must be equally comfortable with technology and human leadership. Future officers will need fluency in areas such as data analysis, AI systems, cybersecurity, and autonomous technologies. But equally important are timeless leadership values like courage and accountability. Training programmes should not simply teach officers to ask, "What does AI recommend?" They should also teach them to ask, "When should I challenge or override AI?" The ability to question AI recommendations, particularly in high-stakes situations, may become as important as understanding the technology itself.

When I joined the Army more

than two decades ago, computer literacy and cybersecurity were already included in the training programmes. Today, I'm sure that AI training is already a part of military curricula. But the officers who succeed on future battlefields will be those who combine technological competence with sound soldiering instincts.

Q If you could choose one area where AI could meaningfully transform how defence organisations function, what would it be?

Decision support across logistics, maintenance and resource allocation. Wars are not lost because soldiers lack courage. Sometimes wars are lost because supplies arrive late, spare parts are not available, or commanders do not have the required information. AI can help improve situational awareness and can offer predictive insights which can be game-changing.

I believe that in the battlefield of the future, victory does not belong to the one that has the most data. Instead, victory belongs to the side that can transform data into insights faster than its adversaries. AI can become a force multiplier for national security and I would love to see AI more in logistics and maintenance rather than replacing soldiers on the front. ■

AI in Judiciary



Manik Goyal

Advocate, Supreme Court of India

Q Indian courts have a backlog of over 50 million cases. Many believe that AI could be the solution to faster justice. Do you agree that AI could be the solution to faster access to justice?

AI is not the solution to judicial backlog. The biggest causes of delay in India are structural and not technological. There are chronic judicial vacancies, and procedural inefficiencies such as the culture of repeated adjournments. For example, over 25,000 judicial posts remain vacant across subordinate courts, reducing capacity of the judicial system. The limited establishment of Gram Nyayalayas also

illustrates the continuing gaps in grassroots access to justice. AI cannot solve these institutional gaps. Faster justice requires major institutional reforms. The problem is therefore not the absence of technology but the absence of institutional capacity. AI can certainly make existing processes more efficient, but it cannot compensate for missing judges, inadequate court infrastructure, or procedural delays built into the system.

That said, AI can support the system by improving accessibility. The Supreme Court has adopted tools that translate judgements into multiple Indian languages, helping litigants better under-

stand proceedings. Real-time transcription tools are also improving courtroom accessibility for people who may struggle with legal language. These applications are meaningful because they expand access to justice without interfering in judicial reasoning. The current approach of Indian courts appears cautious but sensible: use AI to improve accessibility and efficiency, while remaining careful about where human judgement must remain central.

Q With deepfakes and AI-generated media becoming more sophisticated, documents, photos, audio and video recordings that used to serve as reliable evidence can easily be fabricated. How should courts evolve in determining what qualifies as evidence and what doesn't in the AI era?

Forensic tools are evolving alongside these threats. Courts are getting equipped to authenticate digital evidence and identify manipulated material. There is a registrar general in the courts who checks if the evidence submitted is AI-generated.

The law is also evolving. Recently, the Bharatiya Sakshya Adhiniyam came into effect, replacing the Indian Evidence Act. The new act retains and updates the legal framework governing electronic and digital evidence.

In 2023, the Delhi High Court ruled in favour of a victim whose digital likeness and voice were misused. This was one of the earliest Indian judgements on deepfake personality rights. Therefore, the institutions are evolving along

with deepfake technology. There are multiple layers of scrutiny, making AI-generated evidence difficult to pass as credible.

Q One major concern with using AI systems in the judiciary is bias. If algorithms learn from historical legal data, they could reinforce existing inequalities in gender, religion, caste, and economic status. How real is this problem for India and how can we make AI-assisted decisions more responsibly?

That risk is real, especially considering the US' experience with COMPAS and the worsening of racial biases in bail decision-making. AI systems learn from historical patterns, and if those patterns contain bias, the system can reproduce the same bias. A High Court recently flagged that certain subordinate courts were taking the help of AI to write judgements. There are strict regulations in place now to prevent such practices but this is an example of how AI should not be used in the judiciary. The issue is not merely that AI can be biased, but that it can scale those biases across thousands of cases with an appearance of objectivity. That makes human oversight even more important when liberty, rights, or constitutional protections are involved.

In India, the Supreme Court has taken a cautious approach by prohibiting the use of AI in making sensitive decisions such as determining bail eligibility, or criminal sentencing. It is also a mandate in India that only those AI tools that have been specifically designed to prevent biases can be used in the judiciary. These regulations are necessary to ensure that existing biases do not perpetuate into the AI era.

Q People are now turning to AI for legal advice. For example, one of the students at IIT Kharagpur built an AI exten-

Editorial Insight

A legal problem is rarely just a legal problem. It involves social consequences and emotional distress. AI should expand access to justice – not replace lawyers.

sion that can read through the terms and conditions of any online service and can flag any concerning clauses before a user clicks on "I agree". What is stopping people from building similar tools for more sophisticated legal contracts? Do you think AI could reduce the scope of the traditional functions of legal consultants?

No, at least not entirely. AI tends to miss nuances that a legal consultant can see immediately. AI suggestions without human oversight also run into the problem of citing non-existent cases to strengthen one's argument, which can lead to major embarrassment. Until AI systems become consistently reliable, every AI-generated legal argument must be independently verified against authoritative legal databases before it can be relied upon. Currently in India, using AI for legal consultancy is not allowed. However, based on the experiences of other countries, I'm sure that the role of a legal consultant might evolve but it is very unlikely that AI can perform the role of a legal consultant on its own.

Legal disputes are rarely one-dimensional. They involve strategy, negotiation, client counselling, and an understanding of facts that often extend well beyond what is written in a contract. AI can

make legal professionals more productive, but it cannot replicate the experience and judgement required to navigate complex legal disputes.

Q Lawyers are now using AI for researching and drafting legal documents. However, in the process, they miss out on reading and researching the old-fashioned way. Could this create a generation of lawyers overly dependent on technology and not as well-read as the current generation of lawyers?

That is a legitimate concern. While AI is genuinely helpful in the legal profession, overdependence comes with risks, particularly because AI systems still hallucinate and generate inaccurate case citations. Lawyers and AI should complement each other. Therefore, legal training and education need to keep up with the advancements in technology. These days, graduates are using AI through prompting without fully understanding its implications. Therefore, there is a need to incorporate AI literacy in legal education and training so that we can use AI more responsibly. Ultimately, the future of legal practice is not lawyers versus AI, but lawyers who understand AI working alongside it to deliver better outcomes. ■

AI in Business



Manpreet Singh

President, Indian Chamber of International Business

Q There are two competing narratives around AI: one says it empowers small businesses by giving them access to tools and skills, while another suggests large firms will become even more dominant because they control the data and compute. Which direction do you think reality is moving towards, especially in India?

For small businesses, AI is already levelling the playing field. Tasks such as creating brochures, flyers, marketing materials, and digital campaigns, which previously required external agencies, can now be completed in minutes. That dramatically improves speed

and access to market. For more advanced users, the value of AI comes from decision support. One thing that would be a game-changer for me, and I don't know if it already exists, is an AI tool that can analyse fifteen years of emails and instantly identify your best customers, unproductive partnerships, and potential new partners. An insight like this can be a complete game-changer for business owners. I cannot outsource this task to an external company as they could provide these insights to my competitors. I cannot delegate this task to my own employees either because my competitors could gain access to this information when the employee moves to a different company. Using an

AI tool is ideal here, and although it comes with privacy challenges, the implications are different. AI tools are generally meant for mass deployment and it's very hard for my competitors to get access to my information among millions of other customers that the AI company is serving.

Q Different countries are regulating AI differently. For example, EU has the AI Act. The US has taken a decentralised approach. How can we create common standards for AI regulation without slowing innovation?

My opinion is that the moment governments get involved, the innovation stops. Regulation is necessary and the Government of India has already mandated that AI content needs to be labelled as AI content, which is an important regulation as it helps customers make informed decisions without slowing down innovation.

What we need is collaborative regulation and cooperation between the industry and government. For example, industry associations can develop a regulatory framework and submit it to the government. Technology will stabilise over time, but trying to regulate too aggressively at an early stage risks slowing meaningful progress. AI is evolving at extraordinary speed, far faster than traditional policy-making speed. Regulation must therefore evolve continuously instead of remaining static in the face of rapid technological change. That means governments and industry cannot work in silos and must work together to innovate responsibly.

Q As AI automates work, we are seeing developed countries reshoring operations, making the outsourcing model obsolete. India has a huge IT services industry that relies on outsourcing. How should India adapt to this change?

India has adapted before, and it will adapt again. When computers first emerged, people feared that typists and data-entry professionals would disappear. Instead, technology transformed work and created new jobs that didn't exist before.

AI should be viewed similarly. Certain coding or repetitive IT tasks may become automated, but human capability and judgement will remain essential for solving problems. In fact, development cycles that once took months may eventually take days or even hours, which is good as it dramatically increases our productive capacity.

India's strength lies in its scale and talent. While we may not create the next Deepseek or Claude, we will be massive users of these tools to create use cases for the entire world because at the end of the day, we have the manpower, the youth and the intellect to use these systems as building blocks. Jobs as we know them may be affected but in the end you still need manpower to use new technologies and tools to get work done. The goal is to accomplish more in less time. As technology frees people from certain tasks, they can redirect their efforts towards new and more productive forms of work. This shift can create new roles, industries, and economic opportunities.

Q Previously, countries used to build manufacturing around cost-effective labour. But in the AI era, it's difficult to imagine a manufacturing powerhouse with just cost-effective labour and no AI capabilities. What is one

“

While we may not create the next DeepSeek or Claude, we will be massive users of these tools to build use cases for the entire world.

AI capability India must build to compete globally when it comes to manufacturing?

The answer depends on the kind of manufacturing we are talking about. A factory doing injection moulding of plastic products doesn't need AI to operate. However, AI becomes incredibly valuable in designing a new product which is cheaper, faster to implement and has a better application. On the design side, AI will help reduce the time-to-market for any product. However, if we talk about an industry like pharmaceuticals, then AI is a game-changer. Drug discovery often takes years and AI can dramatically shorten that cycle by identifying molecular interactions and narrowing down treatment paths.

Then there is robotics. AI can definitely help with building autonomous factories. However, automation at that level can be problematic because we could have robots building weapons by themselves in an autonomous weapons manufacturing facility, for example and that is concerning.

Q AI is evolving so rapidly that predicting skills that would be valuable even five years from now is difficult. In that sense, how should education evolve? Should we retain the existing model

with a focus on certain skills or should we fundamentally rethink education for the next generation?

I think our education system is completely outdated. I work with the Ministry of Education and I have made some recommendations to them. In some countries, they have books in coffee shops and play areas on critical thinking, radio signals, and AI dynamics meant for toddlers. The books have pictures meant for visual learning so that they can follow. That is the kind of education we need in the AI era. We need to expose students to concepts like systems thinking, technology, AI, and analytical reasoning far earlier than we currently do. Learning itself must also become more experiential. Students today increasingly acquire practical skills through digital platforms, short-form videos, and interactive content rather than traditional textbooks alone. Our education system should embrace these new modes of learning while strengthening critical thinking and problem-solving abilities, ensuring students know not only how to use AI tools, but also how to think independently alongside them. Education must become more future-facing. Most importantly, we should listen more closely to what younger generations want to learn and design future-facing curricula accordingly. ■

AI in Public Policy



Dr Rakesh Varma

ex-IAS (VR), Uttar Pradesh
Cadre

Q In the AI era, students are increasingly using AI to outsource thinking right from school, and as a result, they might not develop the ability to think critically. Given this shift, do you think that our democratic institutions will weaken, since critical thought and intellectual debate are crucial for a well-functioning democracy?

AI should be understood for what it fundamentally is: a system built on existing data and correlations. The danger arises when people begin treating AI as a substitute for thinking rather than a tool for thinking better. In reality, AI

can only respond based on what has been fed into it. That makes human judgement even more valuable, as without strong reasoning skills and critical thinking, even asking the right questions to an AI agent becomes difficult. In that sense, the age of AI may increase the value of natural intelligence. The real differentiator will no longer be access to information but the ability to question, interpret, and apply that information responsibly. AI can generate answers instantly, but it cannot replace intellectual curiosity or independent reasoning. Those human capabilities will become even more valuable in an AI-driven society.

The real policy challenge begins in education. Everywhere in the world, education is heavily influenced by the socioeconomic factors in that country and we need to move away from that if we truly want the next generation to develop critical thinking skills. While adding AI to the curriculum might be necessary, we also need to emphasise depth and rigour rather than surface-level training. The need of the hour is to expose students from all disciplines, be it engineering, business, medicine, or social sciences, to critical thinking. More importantly, the education system should actively reward curiosity and analytical ability rather than rote learning. Without aligning incentives, policy reforms often remain well-intentioned ideas that fail to change classroom behaviour. One cannot make a policy change without aligning incentives accordingly to achieve the intended objective. Technology alone cannot build thoughtful citizens. A healthy democracy ultimately depends on citizens who can question evidence, evaluate competing viewpoints, and reason independently. Thoughtfulness must come from education, starting from a very young age.

Q We are moving towards a future of AI-powered facial recognition and surveillance in public spaces. While it can help law enforcement maintain law and order, it can also compromise citizen privacy and can easily be used to suppress dissent. Where should we draw the line between public safety and individual freedom and privacy?

Governments around the world

want to create digital identities for their citizens so that they are easier to track. But I believe that privacy is supreme and should never be compromised.

What is even more concerning is that these platforms or applications that track citizens are sometimes built by foreign entities who are interested in the local data. Even if there are rules that ensure that the data is stored within the borders of the country, it barely makes a difference as it does not prevent foreign entities from accessing it.

I don't think AI-powered facial recognition can be implemented in any country that cares about privacy. Since the younger generation has grown up in the digital age, they understand the concept of privacy better than anyone else and therefore, they must be at the centre of this movement for the right to privacy. However, social media culture is worrying because the younger generation seems to be comfortable sharing their private lives on the internet.

Nevertheless, privacy need not be compromised to achieve better law and order. We must ensure that citizens' right to privacy is protected.

Q Analytics has become an important part of policymaking. It helps governments identify issues and target interventions for maximum impact. However, there are nuances and practical realities that analytics alone cannot reveal and that become apparent only through community engagement. How should policymakers strike the balance between data-driven decision-making and grassroots-level understanding of the public to ensure that policies are not just effective on paper but also meaningful in practice?

Before relying on analytics, policymakers must ask fundamental

Editorial Insight

The purpose of analytics is to uncover the story hidden in the data – not to use data to justify a predetermined narrative.

questions: How authentic is the data? How was it collected? Who collected it? What assumptions were made during data collection? Poor-quality data inevitably produces unreliable insights, which can lead to bad policies and poor outcomes. Good Analytics cannot make up for the poor quality of data.

It is also very easy to twist the narrative a certain way using data to support it. There have been instances where analysts frame a narrative first and then find patterns in data to support their narrative rather than trying to uncover the story that the data is telling.

In medicine, there is a practice of asking "clinically, how does the patient look?" because even when all the tests indicate that a patient is fine, they might still be sick. The same applies to any data-driven decision making. While data can support and help policymakers, human judgement remains critical as data analytics alone cannot solve our problems.

At the same time, AI regulation is also becoming an issue. Currently, some AI models are built to not be too critical of the government. If that is going to be the path of AI, then using AI for research can become challenging as without asking the right questions and without critically analysing facts, a

country cannot prosper. These are the major issues with using AI and data analytics for policymaking.

Q India has a huge IT industry that provides high-paying jobs to graduates. However, with AI automating entry-level work, people fear that these opportunities might shrink. How should we adapt to this shift?

Most jobs cannot be completely automated. One area where the industry is investing a lot is customer service. Most apps now connect their customers to chatbots or AI agents first, routing them to a human only when required. However, people still prefer human interaction, and for that reason, even if the AI is able to do the job, a customer might still prefer to talk to a human at a call centre.

Jobs are not at risk. What we need to do is give employees more discretion in solving problems. They also need to be trained so that they have the required skills to solve the problems independently. However, the future is no longer repetitive work and will mostly revolve around human judgement and critical thinking. The workforce of the future will therefore be defined less by routine execution and more by sound decision-making and adaptability. ■

The CEO of the AI Era

Shobhna Raj, IIM Kozhikode



I

n 2023, hundreds of employees of a technology company received the news that they had lost their jobs not in a meeting, not from a manager, but from an automated email triggered by an AI-powered workforce optimisation tool. The algorithm had analysed productivity metrics, cost curves, and future demand forecasts. From a technical perspective, this choice was “optimal.” From a human perspective, it was disastrous.

When backlash followed, the CEO issued an apology. The words were familiar: the process lacked human oversight; lessons had been learned. But the damage had already been done. The episode raised an uncomfortable question for modern leaders: if an algorithm makes a decision on your behalf, does responsibility still sit with you? In the age of AI, the answer is in-

creasingly yes, and that answer is fundamentally redefining what it means to lead at the highest level.

The Accountability Problem

Artificial intelligence is transforming the nature of leadership in a subtle but fundamental way. The onus is shifting from individual decisions to the systems that support those decisions, often repeatedly and independently of direct human intervention. In the past, a CEO could trace a problem back to a specific meeting, manager, or policy. But with AI, that becomes impossible. The algorithms are self-contained, self-acting, and self-exponentiating, scaling decisions to levels far beyond individual human oversight. But when things go wrong, it's not the

model or the data pipe that gets scrutinised. It's leadership.

This growing autonomy has introduced deep uncertainty around accountability. Research on agentic AI shows that as systems act with greater independence, responsibility becomes harder to trace. What matters now is not whether harm was intended, but whether it was anticipated. This is why transparency and explainability have emerged as central concerns in organisations that make AI-driven decisions. People want to know why a system acted the way it did, and who authorised it to act that way. Leadership, as a result, is being redefined less by control and more by foresight, by setting boundaries around what AI can decide alone, where human judgement must intervene, and how decisions can be defended

when challenged. In the age of AI, the CEO's role is no longer just to make the right call in the moment, but to take responsibility for the intelligence that operates on their behalf. Increasingly, leadership is judged not by whether AI makes mistakes, but by whether those mistakes could have been foreseen, questioned, or prevented. The quality of governance, rather than the sophistication of the technology, becomes the true measure of responsible leadership in the AI era.

This change is most apparent in the area of ethics. When an AI system causes harm, the first question is not which model failed, but who allowed it to happen? As companies that used AI tools began to experience the backlash of the public in terms of privacy violations, biased treatment, and

lack of transparency, a pattern was clear. No engineer appeared on television. No policy statement was referenced. The blame rested squarely on the CEO.

The Need for Ethical Leadership

AI has quietly bridged the gap between ethics and leadership. Decisions that were once hidden in code are now manifesting as moral failures at the leadership level. Research on corporate digital responsibility indicates that as AI permeates the organisation, ethical standards will follow a predictable trajectory, from the compliance department to the leadership level, as stakeholders view digital harm as a failure of leadership, not technology. The recent spate of reported AI-re-

lated incidents only cements this truth. Trust erodes fastest when leaders cannot explain how or why a system acted the way it did. Ethical leadership in the age of AI is no longer about making statements of value or responding to crises after they happen. It is about foresight: determining where AI should never operate alone, what safeguards must be in place before deployment, and what level of risk an organisation is willing to accept in the name of efficiency.

This moral dilemma immediately translates into strategic decision-making. When Satya Nadella says that AI should enhance human judgement, not replace it, he is actually highlighting a strategic weakness. AI systems make decisions quicker than a leader can challenge them, and this is very tempting. Autonomous systems provide recommendations with such confidence that they inspire leaders to follow them, rather than question them. Studies have found that as AI systems become less transparent and more confident, leaders tend to rely on them because they are precise. The result is a dangerous illusion of certainty.

This shift also changes the culture of executive decision-making in ways that are easy to overlook. In many boardrooms, the most valuable conversations no longer revolve around what the model predicts, but around the assumptions embedded within it. What assumptions were made? Which objectives was the system optimising? What trade-offs were accepted during training? Which groups stand to benefit, and which may be disadvantaged? These questions rarely have numerical answers, yet they increasingly determine whether an AI-driven decision is perceived as legitimate. The future CEO must therefore cultivate an organisation where questioning an algorithm is viewed not as resistance to innovation, but as an essential exercise of responsible leadership.



“You all deserve a personal pat on the back.”

Strategy, in such an environment, can easily become optimisation without wisdom. Research has shown that it is precisely when leaders are struggling to explain the results produced by machines, as opposed to human reasoning, that transparency and explainability become most important. In the age of AI, the CEO’s edge no longer lies in speedy decision-making. Instead, it is the willingness to slow down, to question the model, to be uncertain, even when the machine is certain it is correct.


The Expanding Role of the Future CEO

Ultimately, artificial intelligence has not reduced the CEO’s role; it has honed it. Leadership is no longer about information or analytical expertise. It is about judgement. While AI can suggest, forecast, classify, and optimise, it cannot bear responsibility or explain ethical trade-offs. These are irreducibly human functions, and they reside unequivocally at the top.

The leaders who will endure in the age of AI are not those who deploy the most sophisticated systems first, but those who understand that technology amplifies whatever leadership already exists. Where ethics are performative, AI reveals the gap between values and behaviour. And where strategy is based solely on speed, AI hastens organisations towards decisions they will find difficult to defend.

Ultimately, AI is forcing CEOs to lead in a different way, by placing ethics before profits and understanding before speedy decisions. The defining question for tomorrow’s CEOs will therefore not be whether they embraced AI quickly enough, but whether they built organisations capable of using it responsibly, transparently, and with the courage to challenge it when necessary. In a world that is being increasingly driven by AI, leadership is a function of character and no longer just a function of intelligence. ■

Crossword



ACROSS

- 1 You are happy that the residuals of your predictions are not a big fan
- 6 It might need some guidance but its definitely a star at path finding
- 8 A not so massive, stellar explosion
- 9 A gold fish might benefit from this model
- 12 An office supply store might not sell this, but you'd hope to have this series
- 13 Your estimates are BLUE and that's the best you can do
- 14 My partner can be innocent but they are quick at catching fakes

DOWN

- 2 One might say it doesn't dwell on the past, which could be a good quality
- 3 One might say that a risk analyst and an architect might spend their time designing this
- 4 A test that the British would ace
- 5 A test that Indians would ace
- 6 This is all you need
- 7 Something fishy, one might use, to investigate temporal occurrences of events
- 10 An analyst might not need a cowboy's help, but they might need his _____
- 11 The only wish that it can grant is to cure your 'splitting headache'



Scan to play with hints (or just click on the QR).

Smarter Simulations, Better Engineering



Prof. Suvanjan Bhattacharyya is an Associate Professor in the Department of Mechanical Engineering at BITS Pilani. He was included in Stanford University's list of the world's top 2% of scientists.

Ask anyone who has spent real time running a CFD or heat transfer simulation what the hardest part of the job is, and very few will say “the equations.” The equations are well known. What wears people down is the waiting. A single high-fidelity run involving turbulence, multiphase flow, or conjugate heat transfer can tie up an HPC cluster for hours, sometimes days, and that is before anyone has actually learned anything useful about the design.

I have lived with this problem for

most of my career. My own doctoral work was on swirl and turbulence promoters for convective heat transfer in duct flow, and the years since have only added more variables to track — nanofluids, magnetic fields, vibration, phase-change materials. Every new parameter is another simulation, and every design question an engineer actually cares about — should I optimise the cooling channel of an electric vehicle battery, the rear wing of a race car, or the airflow inside a data centre — demands not one simulation but hundreds, sometimes thousands. That is the gap artificial intelligence is now stepping into, and it is worth being precise about what it is, and is not, doing.

A co-pilot, not a replacement

Let me be direct about something many people get wrong: AI is not replacing the Navier-Stokes solver. Engineers still need it, and for good reason — it remains the only tool we fully trust for validation. What AI does is sit alongside the solver and take over the repetitive parts of the job.

In practice, this has settled into three approaches that most of us in the field now recognise. The first is the data-driven surrogate model, trained on a library of existing CFD runs, which can predict the behaviour of a similar new design almost instantly. The second is the physics-informed neural network, or PINN, which bakes the governing equations

directly into the training process so the model still respects mass, momentum and energy conservation even when experimental data is thin — something I find particularly relevant to the kind of correlation-starved problems we deal with in heat exchanger design. The third is the machine-learning-assisted solver, where the underlying CFD code is untouched but AI is used to refine the turbulence model, guide adaptive meshing, or steer an optimisation loop.

I tend to describe this to my own research students using a comparison most of them already understand: it is not unlike GPS navigation. You are still the one driving the car, and you still need to know how to drive, but a good navigation system saves you from wasting an afternoon stuck in traffic you did not need to sit through. AI in CFD plays a similar role. It does not remove the driver. It removes the unnecessary detours.

Where this matters most in heat transfer

Heat transfer is, in some ways, an unusually good fit for this shift, and the reason is almost historical. Many of the correlations we still teach and use were developed decades ago, often from a fairly narrow set of experiments, for equipment and operating conditions that look nothing like what we run today. Stretch those correlations far enough beyond their original assumptions and they simply stop being reliable.

This is precisely where some of the more interesting work is happening. PINNs have shown they can predict thermal behaviour inside complex porous heat exchangers from surprisingly small datasets. Neural operators are being used to accelerate topology optimisation in ways that transfer across different mesh resolutions, which used to be a real headache. Wide neural networks are improving boiling heat transfer predictions in refrigeration systems, and physics-informed LSTMs are being used to estimate long-term fouling-related performance loss — the kind of slow degradation that is notoriously hard to capture with a single snapshot simulation.

I have worked on a version of this problem myself, using artificial neural networks to predict the thermal performance of solar air-heater tubes fitted with turbulence promoters, and more recently exploring machine-learning models for thermohydraulic prediction in hybrid grooved-tape inserts. The appeal is not subtle: once trained, these models give you in seconds what would otherwise take a full parametric CFD study to establish.

Battery thermal management in electric vehicles is a good illustration of why this matters beyond the lab. An engineer there is juggling cooling performance against packaging space, weight, and cost, all at once. Running a few thousand full CFD cases to explore that trade-off space properly is simply not realistic on any commercial timeline. What AI offers is the ability to screen a much wider set of design alternatives quickly, and then send only the strongest candidates forward for the high-fidelity simulation that should always have the final word.

Beyond faster simulations

The more I look at where this field is heading, the clearer it becomes that the benefit is not only speed.

Diffusion models can now reconstruct reasonably complete flow fields from sparse simulation data. Differentiable CFD frameworks allow turbulence models to be optimised end-to-end rather than tuned by hand. Deep reinforcement learning is being used to discover flow-control strategies for drag reduction and combustion stabilisation that a purely human-driven design process might never have stumbled onto. Geometry-aware neural operators are finally starting to cope with the large, irregular industrial shapes that simple benchmark geometries never prepared them for.

What this really changes is the kind of question an engineer can afford to ask. Instead of running more simulations to chase the same question, you can compare competing objectives, explore a wider design space, and still leave room for the engineering judgement that no algorithm should be allowed to replace.

Has this actually reached industry?

It has reached industry, and reasonably fast. NVIDIA's Physics-NeMo framework and its collaboration with SimScale on turbomachinery evaluation are among the clearer signals, with reported speed-ups described as orders of magnitude faster than conventional workflows for suitable cases. Several major CAE vendors, including Ansys, Siemens, Altair and Cadence, now support training AI surrogates directly on a company's own proprietary simulation history, which means decades of internal engineering knowledge can be folded into a model rather than left sitting in old project archives. Engineers are increasingly able to interact with a digital twin, change a geometric parameter, and see the aerodynamic or thermal consequence almost immediately, rather than queuing another overnight solver run.

What still limits it


None of this should be mistaken for a solved problem. These models can fail badly once pushed outside the range of data they were trained on, particularly across multiple spatial scales — exactly the kind of situation that turbulent and multiphase flows tend to create. For anything safety-critical, rigorous validation against trusted numerical and experimental results is not optional, and I would be uneasy with anyone who suggested otherwise.

The harder challenge from here is not simply making simulations faster. It is building AI methods that respect physical laws consistently, not just on the cases they happened to be trained on, while still being capable of surfacing genuinely new engineering insight from data we already have.

What I would tell engineers entering this space

If there is one thing I would want a younger engineer to take from all this, it is that AI is not replacing CFD. It is changing how we use it. AI is good at repetitive exploration, at screening thousands of candidate designs quickly, and at telling you where your limited computational budget will do the most good. Traditional CFD is still the tool you reach for when you need to verify a result, certify a design, or genuinely understand a new physical phenomenon, and that is not going to change soon.

The future, as I see it, belongs to hybrid workflows where physics-based simulation and AI are used together rather than treated as competitors. Engineers who are comfortable with both will be the ones who design faster, test more ideas, and take on problems that would otherwise have been too expensive or too time-consuming to even attempt. ■



The Unprovable Equation

WHY AI CAN NEVER MASTER HUMAN ETHICS

Mainak Jana



Imagine a courtroom in the future. The person deciding a criminal's sentence isn't a judge. It's a computer. This computer has been trained on every law book, every case file, and every ethical rule. Depending on the previous training, it can just 'calculate' a decision, because it is nothing but a long matrix multiplication for the machine.

The idea to build a perfect, logical machine to handle the moral dilemmas in which we are often stuck, is incredibly tempting. We're messy, we're inconsistent, and we let our feelings get in the way. Wouldn't it be great to hand over the hard stuff to a machine that doesn't feel for itself, but just mimics the output given some particular input information?

But there's a catch. A big one. And to see it, we have to look back about a hundred years to a discovery made by a quiet mathematician named Kurt Gödel. What he proved doesn't just apply to numbers, but philosophy as a whole. It builds a fence around what any logical system, including the smartest AI we can imagine, can ever truly know. And it shows us that the deepest questions of right and wrong will always have to be answered by us, humans.

Our love affair with perfect rules

We humans love a good set of instructions. Think about a recipe. You follow the steps, and you get a cake. In school, many of us learned Euclidean geometry, which can be beautifully constructed by 5 basic axioms using a finite number of logical steps. We've tried to do the same thing with morality. Different philosophies are like different recipe books. Utilitarianism tries to minimise the total 'harm' and maximise the total 'good'. Whereas deontology tries to establish a universal moral framework or duties like "never lie" or "never steal". Each one starts with its

own basic rule and tries to build a system of perfect behaviour from the starting points.

The dream for artificial intelligence is to take one of these "moral recipe books," program it in, and let the AI do the cooking. Give it a goal like "maximise fairness" and with enough brainpower, it could work out some clean, flawless moral math. It sounds like a solution to many of our problems: the AI could work through millions of complex social situations to find the optimal, fairest outcome. No human error. No bad day. Just clean, flawless moral math.

The man who broke logic

In the early 1900s, mathematicians had a similar dream. They wanted to build one perfect, complete set of rules for all of mathematics. One of the main proponents of this movement was David Hilbert. They wanted to make a complete and consistent system to axiomatise all of mathematics. A system where every true statement could be proven (completeness), and no rule would ever contradict another (consistency). It would be the ultimate instruction manual for numbers.

Then, in 1931, Kurt Gödel published a paper that shattered that dream. What Gödel proved is hard to wrap your head around, but its gist is earth-shattering. He showed that any set of rules complex enough to do basic math (like adding and multiplying) is fundamentally flawed in one of two ways:


- I. **It's Incomplete:** There will be true statements that you can never prove using the rules of the system. True things that are forever out of reach.
- II. **It's Inconsistent:** The rules will eventually lead

to a contradiction, like proving that something is both true and false at the same time. That makes the whole system useless.

Gödel's genius was in creating a kind of logical self-trap. Let's have a dummy example for our understanding. Write the following statement on a sheet of paper:

"This statement cannot be proven"

- I. If you can prove the statement, then it's a lie (because it says it can't be proven). So, your system just proved a lie. By definition, lies can't be proven. Only true statements can be proven. That's a contradiction.
- II. If you can't prove it, then the statement is telling the truth!




So, you have a true statement that you cannot prove.

Although this example uses deliberately self-referential language, Gödel showed that any consistent, effectively defined formal system capable of expressing elementary arithmetic contains statements that cannot be proved within that system. In other words, some truths may lie beyond what the system’s own rules can establish.

What this means for our AI Judge

Now, let’s bring this back to our dream of a moral AI. To build an AI that makes ethical decisions, we have to give it a formal system, i.e. a set of rules and goals. LLMs use statistics, mathematics and probability theory to find patterns in language. Using these patterns, they generate text and they can only answer ethical questions based on this framework. Statistics and probability are formal systems that are subjected to Gödel’s incompleteness. So, no matter how advanced we make our LLMs, their scope is always bound by limitations stated by Gödel. So, an AI judge will always have moral blind spots. According to Gödel’s incompleteness theorem, there will always be true statements which the AI system can’t conclude using the set of rules on which the AI is trained. The answer exists, but the AI’s own brain cannot find it.

This is why the dream of the flawless AI judge is an illusion. Think about the hardest decisions we make in life. Should a terminally ill patient be allowed to choose how they die? Should someone who stole food to feed their family be punished in the same way as someone who stole out of greed? None of these questions have answers waiting to be discovered like a hidden equation. They require interpretation, compassion, and the willingness to accept that two reasonable people can arrive at different yet defensible



<< Please don’t make me choose!
Just tell me what to do. >>

conclusions. Unlike a mathematical proof, moral judgement also depends on context, lived experience, cultural values and consequences that may remain uncertain even after careful deliberation. When faced with a truly novel moral dilemma which is not in its code, the best AI can do is spin its gears. It can’t invent new understanding. It can’t have a moral breakthrough. It can only follow the recipes it was given, and Gödel proved that no recipe book is complete enough.

The part we can’t outsource

This is basically a question about the fundamental difference between humans and machines. However convincing it may seem, machines can’t think. They can only recalculate weights and biases to generate outputs that make it seem like they are thinking.

Gödel’s unprovable truths are the domain of human wisdom. Our ability isn’t just to follow rules, but to know when the rules are wrong, inadequate, or need to change. We can look at a system from the outside. We use empathy, intuition, shared stories, and our sense of mercy. This makes us human. And at the end of the

day, moral judgements must be humane.

So, the real danger isn’t a robot uprising. It’s that we will be so impressed by the AI’s speed and consistency that we mistake its calculations for real moral judgement. We might stop seeing the blind spots. That may be the greatest risk of all. Not that AI will become too intelligent, but that we will gradually stop exercising the judgement that makes us human in the first place.

The future isn’t about building a machine to replace our conscience. It’s about building tools that know their limits. An ethical AI should be like a powerful calculator for morality: brilliant at crunching numbers within the rules we set, but always flashing a light that says, “Final decision requires human review.”

Gödel’s math, in the end, is a gift. It’s a reminder that some things can’t be automated. The exhausting yet beautiful work of figuring out how to be good to each other is our job. No algorithm will ever take that burden from us. That responsibility has always belonged to us, and perhaps it always will, no matter how intelligent our machines become. And maybe that’s exactly how it should be. ■

INTRODUCING
DATA
SENSE.

For the **Unfiltered, Unbiased, Unsupervised.**

Follow us for fun facts and insights on all things data, analytics and AI.

<https://www.linkedin.com/company/datasense-by-aina/>



Scan to explore the page
(or just click on the QR).

Your **Location** is Your Password

HOW GEOSPATIAL DATA IS REWRITING BANKING SECURITY

Harsha Boddu

Imagine yourself buying a cup of coffee at a café in Kolkata. You tap your card, expecting the familiar beep of approval. Instead, the transaction gets declined. Not because your balance is insufficient. Not because the PIN was entered incorrectly. Not because the OTP failed. Your phone was detected moving at highway speed several kilometres away, triggering a security flag. To the bank, something about the transaction did not align with physical reality. You could not realistically exist in two places at once.

Security is no longer based only on proving identity. Increasingly, financial systems are trying to determine whether a transaction makes sense in the real world. Geography, movement, and behavioural context have begun to function as invisible layers of authentication.

From Identity Verification to Contextual Intelligence

For decades, financial security revolved around identity verification. If the password matched, access was granted. If the OTP was correct, the transaction proceeded. The assumption was that correct credentials implied legitimacy.

However, fraud evolved faster than traditional authentication systems. SIM swap attacks intercept OTPs. Mobile devices can be remotely cloned. Organised crime networks move stolen funds through mule accounts spread across different regions. Automated bots attempt thousands of transactions within seconds. In this environment, modern authentication systems ask, “Does this



behaviour make sense in the physical world?” Authentication is no longer treated as a single checkpoint. It is becoming increasingly dependent on space and time.

Your Movement Has a Pattern

Although most people assume their movements are random, they rarely are. Daily life tends to follow predictable geographic rhythms: home, workplace, commute routes, favourite cafés, shopping areas, or gym. Over time, these recurring movements create what can be described as a mobility signature.

Modern fraud detection systems analyse timestamps, location, and behavioural consistency to understand this pattern. Every new transaction is evaluated against an established baseline. If a purchase naturally fits into a customer’s normal behavioural geography, the system sees little reason for concern. But when a transaction sharply deviates from established movement patterns, suspicion ris-

es. Imagine a customer withdrawing cash from an ATM in Kolkata at 6:00 PM while, minutes earlier, an online purchase appears to originate from Delhi. For both activities to be genuine, the customer would need to travel faster than is realistically possible.

Fraud systems calculate what is known as geovelocity, that is, the distance between transaction points divided by elapsed time. If the resulting speed exceeds realistic human mobility thresholds, the transaction is flagged for risk. In effect, the system evaluates physical plausibility alongside digital identity.

The Rise of the Geospatial Risk Score

Behind nearly every approval or rejection lies an invisible calculation: the geospatial risk score. Modern fraud engines simultaneously process numerous spatial and contextual signals. They examine the distance from a customer’s usual geographic cluster, assess whether move-

ment between transactions appears realistic, compare IP routing regions with device GPS signals, and evaluate whether transactions occur in locations statistically associated with higher fraud activity. Machine learning systems analyse these variables in real time to estimate behavioural deviation. Fraud detection is no longer based on rigid rules alone. It has evolved into a probabilistic inference problem solved within milliseconds.

This intelligence already powers payment systems used across the world. Organisations such as Visa evaluate hundreds of signals. Mastercard uses advanced analytics and device-location correlation to identify coordinated fraud patterns. In India’s rapidly expanding digital payments ecosystem, major banks integrate AI-driven geospatial monitoring into card and UPI transaction pipelines to detect unusual behavioural movement before funds are released. These systems operate silently in the background, invisible to customers but involved in nearly every transaction.



“I’ll pay cash. I’ve decided to eat healthy but my bank is not used to it yet.”

The Smartphone as a Spatial Witness

The smartphone has evolved into far more than a communication device. It now functions as a dynamic authentication layer. Mobile banking applications no longer generate only OTPs. They continuously contribute GPS telemetry, device fingerprints, network routing information, and behavioural signals. When a transaction occurs, fraud systems may compare the physical location of the device with the merchant location, detect suspicious jumps in movement, or identify spoofing attempts. In this way, the smartphone becomes a spatial witness validating not only possession of the device but also whether the transaction occurs within a believable physical context.

Networks, Not Individuals

Modern fraud rarely operates through isolated individuals. Increasingly, it functions through coordinated networks of accounts, merchants, devices, and locations. Graph-based intelligence systems examine hidden relationships between geographic regions and transactional behaviour to identify suspicious clusters. Multiple accounts operating across cities may appear unrelated on the surface,

yet behavioural patterns can reveal coordinated activity.

When Legitimate Transactions Get Declined

At times, genuine customers experience declined transactions especially while travelling rapidly between cities, using VPN services, or making purchases from unfamiliar locations. What feels legitimate to the customer may appear statistically abnormal to an algorithm trained on routine behavioural patterns.

Rather than assuming malicious intent, the system attempts to reduce uncertainty in real time. Although this cautious approach can occasionally create inconvenience, it also prevents much larger financial losses across payment ecosystems.

The Limits of Spatial Intelligence


Despite its growing effectiveness, geospatial AI is not without limitations. Legitimate users may occasionally be flagged due to unusual travel behaviour, changing routines, inaccurate GPS signals, or virtual private networks that alter geographical IP visibility.

Concerns regarding privacy and data dependency also remain significant. Since these systems rely heavily on behavioural and location data, maintaining transparency and safeguarding user trust become essential challenges. Furthermore, building sophisticated geospatial fraud infrastructure requires substantial investments in real-time analytics, AI systems, and computational resources, making widespread adoption difficult for smaller institutions. In other words, geospatial intelligence remains a probabilistic system designed to reduce fraud rather than eliminate it entirely.

Geography as a Silent Password

Traditional identity verification is not disappearing. Instead, it is being strengthened by contextual intelligence in which movement, behaviour, and geography quietly support decisions in the background. In this invisible game, banks are no longer verifying only who we are. They are quietly asking, does this transaction fit the story of your everyday life?


In the coming years, contextual intelligence will likely extend beyond banking into insurance, cybersecurity, digital identity, and smart cities, where decisions increasingly depend on understanding behaviour rather than simply verifying credentials. As artificial intelligence becomes more adept at interpreting human movement and context, geography may evolve from a passive backdrop into an active participant in digital trust. The future of authentication may be less about secrets we remember and more about patterns we naturally create. In a world where algorithms increasingly learn our routines, trust is becoming something we continuously demonstrate rather than occasionally prove. The strongest password may no longer be something we know or possess, but the ordinary patterns of life that quietly make us who we are. ■



An illustrative heatmap showing how a bank might assess transaction risk using geospatial and contextual signals.

Teaching
Neural
Networks to
Obey the
Laws of Physics

Aditya Arvind



H

erbert Simon, winner of the Nobel Prize in Economics in 1978, beautifully summarised the importance of differential equations for scientific enquiry in the following quote: “The general paradigm is: Given a blueprint, to find the corresponding recipe. Much of the activity of science is an application of that paradigm: Given the description of some natural phenomena, to find the differential equations for processes that will produce the phenomena.” Differential equations are used extensively in every field of study you can imagine. In physics they play an important role in modelling physical systems, from modelling the motion of celestial bodies to describing the behaviour of smallest particles. In engineering they is used to model the behaviour of circuits in electrical engineering and structural vibrations in civil engineering. In biology they are used to model the spread of diseases in epi-

miology and population cycles such as those of lynx and hares. Furthermore they are also used in economics to determine the fair price of financial derivatives and social sciences to model social dynamics. They quietly underpin our daily technologies, from weather forecasting and GPS navigation to medical imaging and modern aerospace systems.

Ultimately differential equations serve as a universal mathematical language to represent the change in some phenomena with respect to some observed variable like time, displacement and provide a framework for understanding complex systems. Therefore, to extract meaningful insights about the system the equations describe, we need to solve them. Often, analytical solutions are impossible, requiring numerical methods and computational simulations to approximate real-world behaviour accurately.

Limitations of
Traditional and Deep
Learning Approaches

Differential equations are generally solved using numerical methods like finite difference or the finite volume methods. These methods approximate derivatives by discretizing the domain into a grid or mesh and then solving the resulting system of algebraic equations. Many traditional methods require carefully constructed meshes that conform to the geometry of the domain. However, these methods suffer from limitations like mesh generation. For complex geometries or moving boundaries, creating and maintaining a good mesh can be difficult and time-consuming. Traditional methods also require large computation costs to solve these complex problems and it grows with the number of dimensions. Whereas vanilla deep-learning approaches require huge amounts of data to train them which is expensive to get in many cases. Furthermore they suffer from poor generalisation outside the range of training data in the unseen ranges of the domain. Despite these flaws thanks

to neural networks’ universal approximation and great flexibility they can be used to model highly complex and nonlinear systems. Therefore to bridge this gap we need a method to impart the knowledge of the physical systems into the neural-network architecture to use the strengths of deep learning while also ensuring that the model remains data efficient.

Overview of PINNs

This is where PINNs come in. PINNs approximate solutions to the partial differential equations by training a neural network to minimise the loss function. The loss function for PINNs consists of two parts: the boundary loss and the physics loss. The boundary loss reflects the initial and the boundary conditions across the space-time domain’s boundary and the physics loss consists of the PDE equations calculated at selected points in the domain which are known as the collocation points. Thus during training the neural network learns to satisfy both data and physics. This makes the predictions more reliable and physically consistent.

At a conceptual level, PINNs represent a fundamental shift in machine learning. Traditional neural networks are often described as “black boxes” because they learn purely from examples, with little regard for the mechanisms that generated the data. PINNs, by contrast, treat the laws of physics as part of the learning process itself. By embedding prior scientific knowledge directly into optimisation, PINNs achieve meaningful predictions with far less data, making them particularly attractive where experiments are expensive or labelled data is scarce.

Key Concepts in
PINNs

1 Collocation Points

Think of collocation points like an exam set by a teacher. The teacher cannot ask the student to solve every single problem ever created in the exam; that would be impossible. Instead, they select

a representative set of problems and verify that the student’s solutions are correct and consistent with the underlying rules. In a PINN, the “problems” are coordinate locations in the domain, and the “rules” are the differential equations. The more carefully the collocation points are spread, the more confident we can be that the network has genuinely learned to satisfy the physics throughout the entire space-time domain.

2 How are collocation points chosen?

In many cases, collocation points are sampled randomly and uniformly across the domain using a technique called Monte Carlo sampling. If your problem lives in a two-dimensional space with time ranging from 0 to 1 second, you might scatter 10,000 random (x, y, t) triples throughout that volume. The sheer number and random coverage ensure that no region of the domain is consistently ignored. However, uniform random sampling has a subtle weakness. Some regions of the domain might be physically “boring” i.e. the solution changes slowly and the physics is easily satisfied there. Other regions might be “hard” with sharp gradients, shocks, or rapid transitions where the physics constraint is most difficult to enforce. A network trained on uniform collocation might waste most of its effort on the easy regions and underfit the hard ones. This insight has motivated a growing body of work

on adaptive collocation. These methods periodically re-sample collocation points, choosing more of them from regions where the residual is currently large, essentially telling the network “you are struggling here; let us focus our physics checks right where they are most needed.”

3 Automatic Differentiation

Knowing where to check the physics (collocation points) is only half the story. To actually enforce a differential equation, you need to compute derivatives of the network’s output with respect to its inputs. If the governing equation is the heat equation, for example, you need the second spatial derivative and the first time derivative of the temperature field predicted by the network. How does one compute these derivatives of a neural network? The answer is automatic differentiation, often called autodiff or AD for short. Modern deep learning frameworks like PyTorch and TensorFlow all implement automatic differentiation natively.

4 How Collocation Points and Autodiff Work Together

The real magic of a PINN emerges when collocation points and automatic differentiation work in tandem. Below is a step-by-step breakdown of a single training iteration. First, a batch of collo-

cation points is selected from the domain of, say, 1,000 random (x, t) pairs. These are passed through the neural network in a forward pass, producing predicted field values at each point. Then, automatic differentiation is used to compute the required derivatives of those predictions with respect to the input coordinates. The derivatives are assembled into the governing differential equation, and the residual: how far the equation is from being satisfied is computed at every collocation point. Finally, this residual is combined with any data loss and boundary condition loss into the total training loss, which is minimised by backpropagation to update the network weights.

Where PINNs Are Making an Impact

PINNs have found applications across an impressive range of scientific and engineering domains. In fluid dynamics, they have been used to infer full velocity and pressure fields from sparse sensor measurements, essentially reconstructing an entire flow from only a handful of data points, because the Navier-Stokes equations provide the missing information. In biomedical engineering, PINNs have been used to model blood flow in patient-specific arteries and to infer material properties of soft tissues from imaging data. In geophysics, they have tackled wave propagation through complex heterogeneous media. In climate science and materials


science, they are being explored as mesh-free solvers for equations that are notoriously expensive to simulate on traditional numerical grids.

Perhaps most exciting is the inverse problem capability. Traditional numerical solvers take known parameters (like the thermal conductivity of a material) and compute the solution (the temperature field). PINNs can run this process in reverse: given partial observations of the temperature field, they can simultaneously fit the solution and infer the unknown thermal conductivity. This ability to solve forward and inverse problems within a unified framework is one of the most distinctive strengths of the PINN networks.


Limitations and Challenges of PINNs

Despite the transformative potential of PINNs, they are not a cure all solution. Training can be slow and difficult, especially for problems with multiple competing loss terms that are difficult to balance. They also tend to underperform traditional numerical methods for high-dimensional problems or when very high accuracy is required. The choice of network architecture, number of collocation points, loss weights, and optimiser all have significant effects on performance, and there is currently no universal recipe for getting these right. Research is actively addressing these challenges. Improved loss balancing strategies, better network architectures, and more intelligent adaptive collocation schemes are being developed and tested. The field is new and evolving rapidly, and many experts believe that the limitations of today’s PINNs will be substantially mitigated within the next few years. ■

Anatomy of a PINN



FORWARD PASS →



← Backprop updates weights using these exact derivatives

Alumnus Speaks



Rachit Tripathi

Co-founder, Blacklitt
Alumnus, PGDBA '17
Alumnus, UC Berkeley '22

Q After working at large banks like Goldman Sachs and HSBC, what made you start BlackLitt? What gap in the market are you trying to solve?

The core problem we observed was a lack of trust in wealth management services. Much of the traditional wealth ecosystem in India is built around selling products with hidden commissions, be it mutual funds, insurance, or investment plans. The commissions and charges are usually not openly told to the customers at the time of signing up. That creates a structural issue: advisors

are often incentivised to maximise commissions rather than optimise outcomes for clients.

At BlackLitt, we wanted to build a more transparent model. Instead of hidden incentives, we believe in charging upfront for advisory. The benefit of such a model is that the incentives are aligned. Advice becomes genuinely client-first as the advisor's only goal is to create good outcomes for the client as he/she is no longer thinking about his/her own commissions.

We began with mutual funds because we saw growing demand for direct investment products. There

are two types of mutual funds: direct plans and regular plans. Banks and wealth managers usually sell regular plans and they come with hidden fees. People are starting to realise that and are now moving to direct plans. But the problem is that they might not get good advisory services when investing in direct plans and that is the gap we are trying to bridge.

Q You've seen investing from both institutional and retail perspectives. How different are these worlds?

Surprisingly, they are far more similar than people imagine. At the end of the day, everyone wants better returns. Institutions may have liabilities such as pension obligations or certain liabilities, while individuals think about children's education, marriage, buying a home, or retirement. But fundamentally, both are solving the same problem: how to manage assets against future obligations.

Q In your career, you must have been in a situation where the model you are testing tells you something that goes completely against your intuition. Whom do you trust in such a situation? The model's findings or your own intuition?

I generally trust human judgement. Models are ultimately built by humans, and investment decisions still involve context and behavioural dynamics that machines struggle to capture. For example, just recently, NASDAQ fell by 4% after the published US job growth numbers were better than expected. But why does

NASDAQ fall when the economy is doing well? It's unintuitive but these things happen in the stock market. At the end of the day, we are not betting against machines. We are betting against humans and that's the reason why human judgement is crucial.

Quantitative systems are extremely useful for high-frequency trading but when it comes to long-term investing, humans are still much better. AI can offer valuable insights but human judgement is still supreme when it comes to long-term. Sometimes assets appear irrationally priced from a model's perspective, but people still invest because they believe in a company's long-term story. Markets are ultimately driven not just by numbers, but by expectations, narratives, and human behaviour, which remain difficult for AI to model consistently.

To answer your question, I go back to the model to understand the unintuitive finding and if there is an error, I fix it. I am not comfortable trusting a model completely with my money or my client's money.

Q Financial markets are supposed to be efficient but they are far from that, at least currently. As more investors make AI-assisted decisions, do you think markets will become more rational or simply develop new kinds of biases?

By design, the financial market can never be fully rational. There will always be players with more information than the others and this leads to information asymmetry. AI, at the end of the day, is just a tool and if people start acting on AI-based recommendations, then there will be another set of people who short those positions because they know that people will act upon those recommendations. In that sense, human emotions and biases will remain central, even in the AI era, at least as far as the financial markets are concerned.

“

Machines are much better at very fast execution, spotting arbitrage and high-frequency trading... For long-term investing, it's still human-driven.

Q What is one genuinely valuable use case of AI in finance that you have encountered - one that you believe has the potential to fundamentally transform the industry?

The biggest value of AI today is productivity. Tasks that once took hours can now be completed in a few minutes. AI also helps reduce operational risk by catching errors, improving checks, and identifying inconsistencies before you go into live trade. For professionals, this translates into a significant productivity advantage.

AI does not replace expertise, but it enables experienced analysts and portfolio managers to spend less time on repetitive tasks and more time on making high-quality investment decisions.

At the same time, overreliance can be dangerous. AI still hallucinates and makes mistakes. That is why every AI-generated output should be treated as a starting point for analysis rather than a final answer. In finance, even small errors can have serious consequences, which means human oversight remains essential.

Q India has no shortage of technical talent, yet relatively few global companies emerge from here. What needs to change for us to build global

companies in India?

The biggest challenge is mindset. India still has a relatively low appetite for risk. Many people prioritise stability from government jobs or big companies, which makes entrepreneurship harder.

The second challenge is access to capital. Building innovative companies, especially in AI or deep technology, is expensive and there isn't enough capital in India to fund these initiatives. But there is also a cultural factor: we often expect immediate success but deep innovation takes time. Companies usually struggle for years before a breakthrough and it requires immense patience and strong support systems.

We are beginning to see progress in sectors such as deep tech, robotics, and space technology, but scaling globally requires patience, capital, and a willingness to think long term. The pace of innovation is currently slow compared with that of our global peers, but the potential is there. If we can cultivate greater risk-taking, improve access to capital, and remain patient with deep technology ventures, I believe India can build globally competitive companies over the coming decade. ■

The Invisible Farmer

Sahit Gunda &
Karthik Ashtekar

India's growth story is often associated with the service industry. Yet agriculture remains one of the country's most important economic pillars. Agriculture and allied sectors contribute roughly 18% of India's GDP while supporting the livelihoods of more than 42% of the workforce. This suggests that a large share of workers remain engaged in relatively low-productivity activities, limiting incomes in rural areas and slowing broader economic growth. As India seeks to improve productivity and build climate-resilient rural economies, access to affordable credit becomes increasingly important.

Farmers require capital not only for seeds and fertilisers, but also for building productive infrastructure such as irrigation systems, and storage facilities, things needed to improve agricultural productivity and incomes. Despite this, many farmers remain outside the formal lending ecosystem. In India, nearly 61% of farmers

depend on rainfed agriculture, making their yields and incomes highly vulnerable to changing weather patterns. If access to credit remains restricted, farmers will have a hard time escaping the poverty cycle. Small and marginal farmers account for nearly 86% of the country's operational land-holdings, yet only about 41% of them are covered by commercial banking channels. As a result, many agricultural households still depend on high-interest loans from moneylenders, friends, relatives, and other informal sources.

Why Farmers Remain Outside the Formal Banking System

Traditional lending institutions assess borrowers using documents such as salary slips, income tax returns, audited financial statements, credit bureau scores, and established banking histories. Most farmers do not possess these documents. Also, since their income is seasonal and weather-dependent, past income might not even be a good indicator of their credit-worthiness. A farmer may own productive land, consistently generate agricultural output, and possess strong repayment capacity, yet still appear "invisible" to a commercial lender.

Historically, institutions such as the Grameen Bank, cooperative societies, regional rural banks, and microfinance institutions helped bridge this gap through trust-based lending models built on community reputation. While these played an important role in expanding financial inclusion, they face limitations with scale.

Building a Digital Financial Footprint

In the AI era, the question for lenders is slowly shifting. Instead of asking, "Does this farmer have a salary slip?", they are beginning to ask a more relevant question: What does the farmer's economic

activity reveal? A farmer may not have formal documents, but they often leave behind many smaller signals. Purchases of seeds and fertilisers, crop sales, utility payments, mobile payments, Farmer-Producer Organisation participation, all tell part of the story.

Individually, these signals may look ordinary, but together, they form something powerful: a digital financial footprint. This footprint enables lenders to see patterns that were previously hidden. Platforms such as Jai Kisan's Bharat Khata are important because they digitise rural transactions that earlier lived only in verbal agreements. Once these records become structured and traceable, they can help convert informal economic activity into financial evidence. India's Account Aggregator framework further strengthens this shift. By allowing individuals to securely share financial information across institutions through a consent-based system, it provides borrowers with a portable, verifiable financial identity. For farmers who were previously invisible to formal lenders, this means credibility no longer depends solely on conventional paperwork. It can also be built from the economic activity they were already generating.

Predicting Income Instead of Looking at the Past

Agricultural lending differs from most consumer or business lending. For a salaried employee, past income is often a reasonable indicator of future repayment capacity. For farmers, the situation is far more uncertain. Their income depends on the next crop, the next rainfall pattern, the next market price, and the condition of the land during the growing season. A farmer may have repaid past loans responsibly, but a failed monsoon can quickly change their ability to repay. This is why agricultural credit requires a different approach. Lenders cannot rely only on historical records. They also need to estimate future income. Modern credit models are therefore beginning to combine data from different sources to understand not just what a farmer earned in the past, but what they are likely to earn in the coming season.

1 Geospatial Intelligence

Geospatial intelligence plays an important role in this shift. Satellite imagery from sources such as Sentinel, Landsat, and MODIS

enables lenders and agri fintech platforms to monitor agricultural land without relying solely on manual field visits. These images can help identify whether land is being actively cultivated, how crops are developing, and whether vegetation conditions appear healthy during the season.

One widely used measure is the Normalised Difference Vegetation Index, or NDVI. It helps estimate crop health by measuring the difference in the reflectance of visible and near-infrared light from vegetation. In simple terms, healthier crops usually reflect light differently from stressed or damaged crops. By tracking these changes over time, lenders can get a clearer view of crop progress, land use, and potential yield.

2 Soil and Agronomic Data

Crop output also depends heavily on the condition of the soil. Two farmers may cultivate the same crop in the same district, but their expected yields can differ significantly depending on soil moisture, nutrient levels, salinity, etc. When lenders use agronomic data and understand the quality and condition of the land, they can make better estimates of the upcoming harvest.

3 Weather and Climate Intelligence

Weather remains one of the biggest sources of uncertainty in Indian agriculture, especially because a large share of farmers still depend on rainfed farming. If rainfall is expected to be below normal, or if a region is facing a higher risk of drought or extreme heat, the model can adjust its estimate of future income. Similarly, favourable weather conditions during key stages of the crop cycle can improve the expected repayment outlook. This is especially important in the context of climate change as farming incomes



How NDVI helps map crop health based on reflectance of visible and near infra-red light

are becoming more exposed to sudden and uneven weather patterns.

4 Market and Transaction Data

A good harvest does not automatically guarantee strong income. The farmer must also be able to sell the crop at a reasonable price. Commodity prices, local market demand, historical sales records, buyer relationships, and supply chain transactions help lenders understand the commercial side of farming. They show whether the farmer is connected to reliable buyers, and whether the crop has a strong market at the expected time of harvest.

Explainability

Financial institutions and agri-businesses cannot rely entirely on “black-box” decisions. Lenders, regulators, risk teams, and farmers need to understand why an AI model made a specific prediction. This need for transparency becomes especially important because of a common weakness in machine learning known as shortcut learning, or the “Clever Hans” effect. In such cases, an AI model may appear highly accurate during training, but instead of learning

the real logic behind a problem, it may be relying on irrelevant or misleading patterns in the data.

A model using satellite, weather, and alternative data may classify a farm as “high default risk” because regional agroclimatic data shows severe rainfall shortages. However, explainability frameworks such as SHAP, can break down the decision at the individual level. The analysis may reveal that although the region is facing low rainfall, the specific farmer has invested in micro-irrigation, improved soil health, or adopted practices that reduce the actual risk. Explainable AI, therefore, plays a critical role in responsible agricultural lending. It helps lenders separate genuine risk from accidental bias. In a sector where the core problem is already financial invisibility, explainability ensures that AI does not create a new form of invisibility behind a credit score.

Industry in Action

One of the most advanced examples of data-driven lending comes from L&T Finance. Its proprietary AI ecosystem, including Project Cyclops, performs multiple automated validations before making a lending decision. Instead of relying on a single score, the system evaluates multiple layers

of information before approving a loan.

Samunnati is another example. The company has adopted a different approach. Rather than assessing individual farmers one by one, it focuses on underwriting Farmer Producer Organisations (FPOs). Using data-driven assessment and underwriting, Samunnati evaluates operational capability, governance quality, and value-chain activity rather than relying solely on historical financial statements.

The Promise of Credit Access

The larger promise of AI in agricultural finance is not limited to faster loan approvals or more accurate credit scores. Its real potential lies in changing how agriculture itself is understood. As agricultural education, research institutions, agri-tech platforms, and digital public infrastructure become more connected, farming can gradually move from inherited guesswork to informed decision-making. The future farmer may not only depend on experience passed down through generations, but also on dashboards, satellite insights, predictive models, and scientific advisory systems. ■



Keeping farmers away

How India Can Become an AI Superpower

Prabal Biswas, IISER Pune

Saty Nadella, CEO of Microsoft aptly said that “India has the opportunity to do AI for its own people first, but in doing so, it will create the playbook for the rest of the world.” This suggests that solutions developed in India for its linguistic diversity and healthcare scale can be exported to the Global South.

The Hardware

India has recently emerged as the leader in IT and software exports and is the third largest exporter of smartphones and second largest by manufacturing volume. So, we can export AI IP and Edge AI, directly embedding it in chips of millions of phones we make. Now, we are more reliant on foreign-owned compute (Azure) that puts forth a dependency risk. This poses a risk if there is a variation in global geopolitical winds. So, it is essential for India to build its own compute. The government has launched the IndiaAI Mission (\$1.2 billion investment). This in-

volves making foundational model compute capacity of about 38,000 high-end GPUs (India AI Compute Capacity). The policy focuses on implementing India AI Innovation Centre and India AI Startup financing where it is proposed to develop and host indigenous LLM and data centers (in Maharashtra, Tamil Nadu, Karnataka) and to fund and support serious startups by providing access to the GPUs at a subsidised rate. One can take the analogy of ISRO launching satellites for other countries, likewise we can provide industries in other countries with access to our GPUs to solve their problems.

The Data

As NVIDIA CEO Jensen Huang famously said at the 2024 NVIDIA India Summit, “India should not just export data to import intelligence”. India has vast amounts of well-structured data through systems such as Aadhaar, UPI, ONDC which is present nowhere in the world. India should share some of the government datasets for use by AI companies like UPI. Wait! still confused how a government can share citizen’s data without their consent? Well, it’s non-personal data. It’s not saying “Mohan bought something from Sohan for 500 rupees in Mumbai” but “1 billion dollar transaction

in Mumbai in Q1”. The model has been adopted in Singapore and the EU.

However, access alone is not enough. India must establish robust governance frameworks that define who can use these datasets, for what purpose, and under what safeguards. Transparent standards for privacy, security, and accountability will be essential to build public trust in AI.

The Models

We must also focus on building ‘context-aware models’ trained with the data of healthcare, agriculture, and finance. Conventional models are trained on US-based financial data and may therefore be unsuitable for underwriting loans for rural farmers in India. However, if we can train a model on our UPI dataset, it can. This will be a great leap towards our progress. Data shows 10% of Indians speak English fluently. Most AI models are trained primarily on English-language data and are therefore unable to meet the needs of many Indians. So, the government should more prominently fund projects like Bhashini under the Natural Language Translation Mission. It can provide real time translation and language accessibility. It is essential to break the barriers of language to establish a developed nation.

“

We are not in the race to beat LLM in poetry! We are in the race to build AI that helps billions get better healthcare, better education, faster justice.

– Nandan Nilekani

The Talent

India lacks skilled labour particularly in the field of AI. We must encourage more research in quantum computing, applied mathematics and physics required to design these LLMs. Colleges like IISERs, IITs must have courses on ‘AI’ and ‘ML’ so that students can directly be a part of the workforce after their graduation. It is necessary to realise that the government cannot do all this alone, so public-private partnerships must be done in order to achieve goals in a shorter span of time. Startups and industries must give opportunities to young college students to work as an intern and gain hands-on experience in these fields.

The Compute Infrastructure

The government must set up ‘AI green corridors’ in states with renewable energy to power data centers like Gujarat and Karnataka. ‘Essential Infrastructure Status’ must be granted to AI data centers. Tariffs and import duties on GPUs must be removed if the big companies agree to lend a part of it to small startups at a subsidised rate. The government must open up departments in existing universities like IISER, IITs and other research institutes and offer M.Tech, PhD, PostDoc opportunities with higher stipends to attract graduates.

The Future

As Nandan Nilekani, co-founder of Infosys, rightly said, “We are not in the race to beat LLM in poetry! We are in the race to build AI that helps billions get better healthcare, better education, faster justice.” We must focus on ‘Applied AI’ that solves problems of billions of people. We must ensure that the next era of technology is not just ‘Made in India’ but ‘Thought in India’. In doing so, India can move beyond being the world’s technology back office to becoming an architect of AI solutions designed for economies across the Global South. ■



The Best Animator is Still Human



Himanshu Sharma is an award-winning animator and filmmaker. He is the Creative Director of GenAI at Pocket FM.

Hi everyone.

I'm Himanshu Sharma. I'm an animator, filmmaker, and currently the Creative Director at Pocket FM, where I lead the GenAI creative team. Before joining Pocket FM, I worked with Stride Ventures, and over the last nine years I've been fortunate to work across animation, filmmaking, visual storytelling, and creative direction. Along the way I've directed several short films that have been recog-

nised at different film festivals, and every project has taught me something new about storytelling and the creative process.

One thing I've realised is that animation has never stood still. It has always evolved with technology. We started with hand-drawn animation, then came digital tools, 3D software, motion capture, real-time rendering, and today we are witnessing another major shift with artificial intelligence.

Whenever a new technology enters the industry, people ask the same question. Will it replace artists? Personally, I don't think that's the right question. The better question is: How can artists use this technology to become even better storytellers?

AI is Changing the Way We Create

Today, AI can generate concept art, animate characters, create environments, synchronise lip movements, upscale footage, clean audio, edit videos, and even produce animated sequences from text prompts. Tasks that once required large teams can now be completed by much smaller teams in a fraction of the time.

As someone leading a GenAI team, I get to see these capabilities every day. AI is incredibly powerful. It

has completely changed the speed at which we can experiment with ideas. Earlier, if a director wanted to test five different visual styles, it could take days. Today we can generate multiple directions within hours and quickly decide what works best.

That speed changes the creative process. Instead of spending weeks finding one idea, we can explore many ideas and spend more time refining the strongest one.

Another area where AI has made a huge impact is post-production. Editing, colour correction, background removal, object clean-up, scene detection, subtitles, dubbing, and voice generation have all become much faster. Independent filmmakers who once struggled with limited budgets now have access to tools that allow them to produce content at a much higher quality than before.

AI has also improved pre-production. Directors can generate storyboards, concept art, mood boards, and shot references before production even begins. Script analysis tools can identify pacing issues and estimate production requirements. These tools don't make creative decisions for us, but they help us make better decisions much faster.

These advances also make creative iteration far more accessible.



Small studios, independent artists, and students can now prototype ambitious ideas, experiment fearlessly, and bring compelling stories to life with fewer barriers.

Where AI Still Falls Short

That said, after working closely with AI, I can confidently say that it still has very real limitations. One of the biggest challenges is character consistency. This is something every animator understands immediately. A character is not just a single image. That character has to exist across hundreds, or sometimes thousands, of frames. Every expression, every costume detail, every proportion, and every movement needs to remain consistent. Even small inconsistencies become very noticeable once the animation starts. AI is improving rapidly, but maintaining that consistency across an entire project is still one of its biggest weaknesses. Sometimes facial features change. Sometimes hairstyles change. Sometimes body proportions shift from one shot to another. In professional animation, these things are simply not acceptable.

When artists build characters manually, every detail is intentional. We know why a character looks the way they do. We understand their personality, body language, and emotional expressions. That level of control is still difficult to achieve with AI alone.

Another challenge is emotion. Animation is not only about making beautiful visuals. It's about making audiences feel something. It's about timing, performance, silence, expressions, and subtle movements. Those decisions come from observation, empathy, and human experience. AI can generate impressive images because it has learned patterns, but it doesn't actually understand why a scene feels emotional.

Sometimes I see AI-generated

work that looks technically perfect but emotionally flat. Everything is polished, yet something feels missing. That's usually the human touch.

The Human Touch in Animation

I experienced this very clearly while working on my film *The Fierce Tiger*. For this project, we deliberately chose to create the characters manually rather than relying completely on AI. It definitely took much longer. Every design, every expression, and every movement required careful attention. There were moments when AI could have helped us move faster, but we felt the manual process gave us something more valuable.

The characters felt alive. As artists, we were making thousands of tiny creative decisions that AI simply wouldn't make. Those decisions added personality, emotion, and authenticity to every frame. It reminded all of us that animation is ultimately about connecting with people, not just producing images.

Another topic people often discuss is employment. There's no denying that AI is automating many repetitive tasks. Clean-up work, background generation, basic editing, and several production processes can now be completed much faster than before. Naturally, this changes the type of work available in the industry.

But every technological shift creates new opportunities as well. Today, we need artists who understand prompting, AI workflows, creative supervision, visual development, and quality control. We need people who know when to trust AI and when to ignore it. The role of the artist is evolving rather than disappearing.

I believe the artists who continue learning will always stay relevant because creativity has never been

about using a particular software. Creativity is about solving problems and telling stories. That's something no technology can replace.

The Future Lies in Collaboration

People sometimes think AI and manual animation are competing with each other. I don't see it that way. I think they're partners. Let AI handle repetitive work. Let it speed up production, organise assets, generate variations, and automate technical processes. But let artists focus on storytelling, performances, world building, character design, and emotional moments. That's where human creativity shines. Some of the greatest animated films in history were created long before AI existed. They continue to inspire audiences today because they were made with genuine human emotion.

Technology will continue to evolve. AI will become more capable every year. Character consistency will improve. Animation quality will improve. Entire production pipelines will become faster. But audiences don't fall in love with technology. They fall in love with stories. As filmmakers and animators, our responsibility is not simply to use the newest tools. Our responsibility is to create work that people remember long after the credits roll.

So my belief is very simple. Don't fear AI. Don't depend on it either. Use it wisely. Use it where it genuinely adds value, but never stop developing your own creativity because that is the one thing technology cannot automate. If we combine the efficiency of AI with the imagination, experience, and emotional understanding of human artists, I truly believe the future of animation and filmmaking will be more exciting than ever before. ■

TRILYTICS CONCLAVE



Dr. Debarag Banerjee and Dr. Kamlesh Kumar with the PGDBA batch during the inauguration ceremony of Trilytics 2025

T

rilytics 2025, the fourth edition of the PGDBA programme's flagship analytics conclave, once again brought together the best of academia, industry and student talent under one roof. Organised by the students of the Post Graduate Diploma in Business Analytics (PGDBA), the event celebrated analytical excellence through a national case competition, keynote addresses, expert panels and hands-on workshops, showcasing the growing influence of data science in solving real-world business problems.

The centrepiece of Trilytics 2025 was its national analytics case competition, **sponsored by L&T Finance**, which attracted over **7,700 registrations** from students representing premier institutions

across the country, including the IITs, IIMs, NITs, IIITs and several leading universities. Participants competed for a prize pool of INR 2 lakh.

After multiple rounds of rigorous evaluation, the winners of the competition were:

Winner - Team Dheerajpreet_s, IIT Roorkee

1st Runner-up - Team Analytica Trio, IIM Calcutta

2nd Runner-up - Team PlanB, IIM Sirmaur

Beyond the competition, Trilytics 2025 featured an exceptional lineup of speakers who offered participants a front-row view of the rapidly evolving AI landscape. **Dr. Debarag Banerjee**, Chief AI & Data Officer at L&T Finance,

discussed how artificial intelligence is transforming digital lending and financial services at scale.

Dr. Srinivas Telukunta, Global Head of Generative AI & AI Agent Programmes at PepsiCo, explored the emergence of agentic AI and its growing role in enterprise decision-making. Academic perspectives came from **Prof. Aritra Hazra**, IIT Kharagpur,

who examined AI's capabilities in strategic gameplay, and **Prof. Ankit Kumar**, IIM Calcutta, who conducted an interactive workshop on the latest developments in artificial intelligence.

A key highlight of the conclave was the panel discussion titled "AI and ML in Industry 5.0: Disruptions, Challenges and the Road Ahead," featuring **Prof. Samit Paul** (IIM Calcutta), **Dr. Kamlesh Kumar** and **Mr. Sandip Kumar** (L&T Finance), and **Ms. Urvashi Ghosh** (Kotak Mahindra Bank). The discussion examined how organisations can responsibly

deploy AI while balancing innovation, governance and measurable business value.

The event also fostered networking, mentorship, and collaboration opportunities, enabling participants to engage with peers, professors, experts, and emerging industry trends.

Trilytics 2025 was more than a showcase of technical excellence. It was a celebration of the people and ideas shaping the future of data science. From students presenting innovative solutions to industry leaders sharing lessons from large-scale AI deployments, every session reflected the growing importance of analytics in solving meaningful real-world problems. As the field continues to evolve, platforms like Trilytics play a vital role in bringing together curious minds, fostering collaboration, and inspiring the next generation of data scientists. ■



Dr. Srinivas Telukunta explaining the growing role of agentic AI in decision-making



Team Dheerajpreet_s from IIT Roorkee, winners of Trilytics 2025, receiving the cash prize from Dr. Kamlesh Kumar



**The AINA 7.0
Editorial Team**

The future of AI will be shaped
not only by what we build but by
how **responsibly** we build it.

Thanks for reading. Scan the QR code and let
us know what you think of the magazine
(or just click on the QR).



AINA has a new home.
Scan or click to explore.

