



PEOPLE . PLANET . PROGRESS



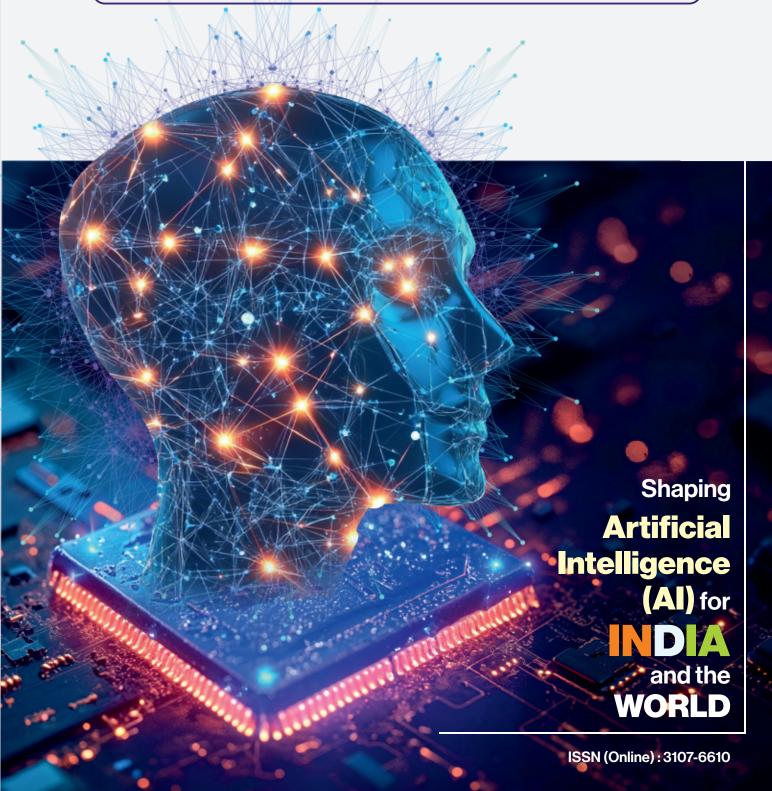




TABLE OF CONTENTS





roreword	O1
The Evolving Landscape of Artificial Intelligence: From Foundations to Future Frontiers	03
Evolution of India's Al Policy and Strategies	27
Understanding India's AI Ecosystem: Featured interview with CEO of IndiaAl Mission	31
OPSA's Role in Empowering India's Al Revolution	41
Al Policy Priorities for India	47
Al Across Sectors: Mapping the Impact of Artificial Intelligence in India	53
Al for Empowering Science and Engineering	83
Digital Public Infrastructure (DPI) for AI	91
Al for Arts and Heritage	97
Building Al for a Sustainable Tomorrow	105
Al Impact Summit 2026	113

Foreword

Prof. Ajay Kumar Sood

Principal Scientific Adviser to the Government of India



rtificial Intelligence (AI), once a theoretical concept introduced in the 1950s, has evolved into a transformative force that now touches nearly every aspect of human life. From its early roots in symbolic reasoning and machine learning to today's advanced systems capable of perception, decision-making, and adaptation, AI has steadily moved from laboratories into the heart of our daily experiences powering healthcare diagnostics, enabling smart agriculture, personalising education, and streamlining governance.

As its cross-sectoral influence continues to grow, the need for preparedness is no longer optional; it is essential. Strengthening our technological foundations is not just about meeting today's demands, but about equipping future generations to thrive in an Al-driven world. In this pursuit, technological sovereignty, the ability to independently develop, deploy, and regulate Al systems has emerged as a critical national priority. For India, this means securing its digital infrastructure, protecting citizen data, and reducing reliance on foreign technologies to ensure long-term resilience and strategic autonomy.

In line with this, the Government of India has launched the IndiaAl Mission, an initiative to ensure that the country not only keeps pace with global advancements, but leads with inclusive, ethical, and context-aware Al solutions. Recognising the vast linguistic and cultural diversity of the nation, the mission prioritises the development of Al tools that are responsive to native tongues and respectful of local norms.



This spirit of inclusive innovation finds its most promising expression in rural empowerment. Whether through voice-enabled agri-advisory platforms that support farmers in their native dialects or Al-powered healthcare diagnostics that reach remote areas, the mission strives to democratise the future. IndiaAl is redefining what it means to build technology "for the people" by promoting accessibility, affordability, and relevance, while lowering barriers to entry and transforming innovation into a shared opportunity.

India's leadership in building transformative Digital Public Infrastructure (DPI) through platforms like Aadhaar, UPI, and DigiLocker has laid a strong foundation for scalable, interoperable, and inclusive tech ecosystems. Moreover, DPI is simplifying global connectivity and business processes, and the infusion of AI adds intelligence, and adaptability to these efforts.

As countries around the world race to harness AI for strategic and economic advantage, India's collaborative model is bringing together government, academia, startups, and industry to ensure that innovation is not only cutting-edge but also locally grounded. These partnerships foster cost-effectiveness by pooling resources, reducing duplication, and enabling shared infrastructure like data repositories and compute platforms. Such joint efforts democratise access to AI capabilities, allowing emerging economies and smaller players to participate meaningfully in the global AI landscape.

Beyond commercial success, Al must be a force for social good. It must address pressing challenges such as healthcare access, education quality, climate resilience, and financial inclusion. In this regard, the Ministry of Education's establishment of Centers of Excellence (CoEs) in Al represents a vital step forward. The CoEs in Al, focusing on healthcare, agriculture, education and sustainable cities, established in collaboration with leading institutions aim to drive interdisciplinary research, nurture industry-academia partnerships, and empower students, researchers, and innovators from diverse backgrounds. Aligned with the vision of "Make Al in India and Make Al Work for India," the initiative is a strategic step toward building an inclusive, innovation-driven Al ecosystem that addresses national challenges and empowers citizens across sectors.

The Office of the Principal Scientific Adviser (OPSA) has also played a thoughtful role in guiding India's Al journey, helping to shape regulatory and techno-legal frameworks that balance innovation with societal safeguards. OPSA prepared a concept note, "Al Regulation in India", that highlights transparency, accountability, and fairness, and calls on the government to develop regulations, standards, and ethical guidelines. It proposes three key approaches: a Whole of Government strategy to ensure cohesive policy, a balance between innovation and risk management, and a robust Data Governance including techno-legal framework. The multistakeholder Advisory Group's report on "Al Governance and Guidelines Development" further reinforces this commitment, introducing forward-looking concepts such as Digital by Design Governance and Do No Harm to ensure accountability.

As we shape AI for India and for the world, our collective responsibility is clear: to build technologies that serve humanity, respect diversity, and empower every individual. With inclusive design, thoughtful governance, and a spirit of collaboration, India is poised to lead the way in crafting AI that is not only globally competitive, but deeply humane.

The Evolving Landscape of Artificial Intelligence:

From Foundations to Future Frontiers

rtificial Intelligence (AI) stands at the forefront of technological evolution, offering new opportunities to enhance daily life and contribute to evolving social systems around the world. As a field of computer science, AI is fundamentally concerned with enabling machines to exhibit cognitive functions traditionally associated with human intelligence, such as perceiving, learning, reasoning, problem-solving, and decision-making. The ambition is to create intelligent machines or computer programs that can mimic, and in some cases enhance, human intellectual capabilities. This capacity to learn and solve novel problems, act rationally, and engage with the world like humans defines the essence of AI.

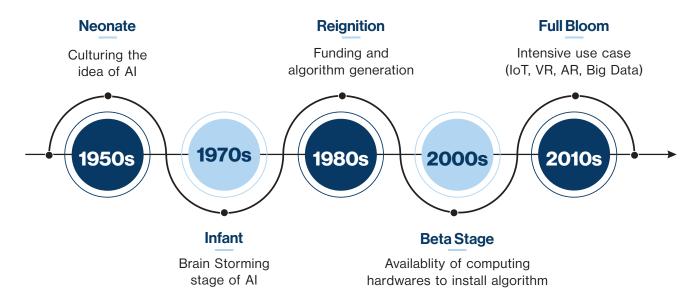
As the challenge of transforming raw data into actionable insights grows, especially under the pressure of rapid decision-making, AI is emerging as a powerful ally. It processes vast amounts of information with speed and precision, uncovers patterns that might escape human notice, and supports more informed, data-driven choices. Across industries, AI is driving productivity and innovation: improving healthcare diagnostics, streamlining supply chains, enabling scientific breakthroughs, and optimising energy use. By automating routine tasks such as document processing, content categorisation, and query resolution, AI frees people to focus on areas where human intelligence excels: problem-solving, critical thinking, and creative exploration. Today's challenges like climate change, resource management, and public health require timely, coordinated responses; AI stands as a vital partner in building sustainable and resilient futures.





Shakun Shan Specialist (IT and Digital Services) Office of PSA

Evolution of Al



Source: Press Information Bureau

The journey of AI is far from a recent phenomenon. Its theoretical and technological roots stretch back more than 70 years, shaped by the vision of pioneers such as Alan Turing, Marvin Minsky, and John McCarthy. In 1950, Alan Turing's seminal paper, Computing Machinery and Intelligence, posed the groundbreaking question, "Can machines think?", laying the conceptual groundwork for decades of exploration. Six years later, John McCarthy coined the term "artificial intelligence," marking the start of a new era in computing research.

Since then, Al's evolution has unfolded in distinct phases. The 1950s were its "neonate" stage—a fresh, intriguing branch of technology that captured the imagination of scientists worldwide. The 1970s became a time of experimentation and bold ideas, as researchers probed the boundaries of what machines might one day achieve. By the 1980s, interest surged again, and the first early applications began to surface. In the 2000s, leaps in computing power transformed Al from a research concept into a practical tool, opening the door to real-world adoption. Then came the 2010s, a decade of rapid expansion, when Al became central to innovations in the Internet of Things, virtual and augmented reality, and the vast, data-driven realm of Big Data.



In India, progress unfolded alongside these global changes. In the early years, institutions such as IIT Kanpur and IISc Bangalore began building expertise in computer science, setting the stage for future breakthroughs. A key milestone came in 1986 with the launch of the Knowledge Based Computer Systems (KBCS) project the country's first large-scale Al research effort.

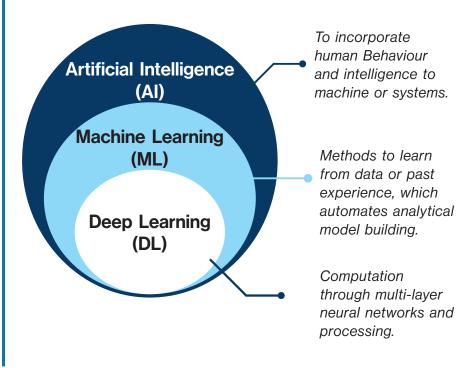
The late 1980s and 1990s brought further momentum. The creation of C-DAC in 1988 greatly boosted India's computing capabilities and opened new opportunities for AI research. Software companies began to explore AI for automating business operations, and by the 2000s, major players like TCS, Infosys, and Wipro were investing in AI projects. Universities expanded their programs in machine learning and AI, training the next generation of specialists.

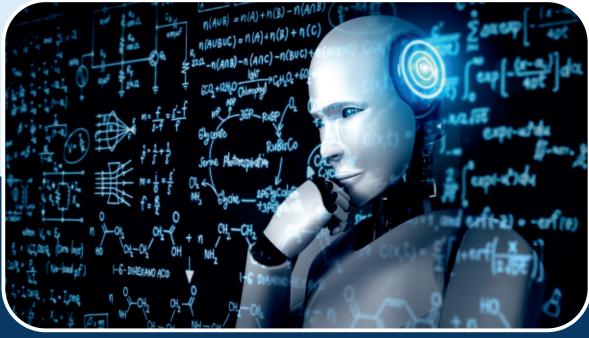
In the 2010s, developments accelerated. The Digital India initiative (2014–15) placed emerging technologies at the heart of the country's growth plans. In 2018, NITI Aayog introduced the National Strategy for AI, outlining how AI could support both economic development and social progress. Around the same time, AI-focused startups began to thrive, drawing in funding and global attention.

Now, in the 2020s, AI has become a central force in India's technological vision. It is being applied to strengthen healthcare, improve agricultural efficiency, develop smart cities, modernise education, and enhance governance. Programs such as "AI for AII" reflect a commitment to making AI accessible and beneficial to every part of society transforming it from a specialist's tool into a key driver of an inclusive, forward-looking nation.

Understanding the Landscape of Al Models From Machine Learning to Generative Al

Al is a bit like an ever-growing city, full of bustling districts, each with its own character, yet all connected by invisible roads of shared ideas. Within this city are the neighbourhoods of Machine Learning, Deep Learning, Neural Networks, Computer Vision, Generative Al, Large Language Models, Edge Al, Explainable Al, and Agentic Al. Each plays its own role, yet together they create something greater than the sum of their parts.





Machine Learning:

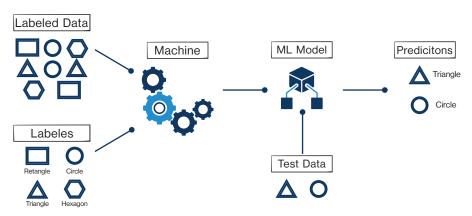
The Adaptive Core of Al

At the heart of AI is Machine Learning (ML) the adaptable model that can learn not by remembering the dataset, but by figuring out the patterens by themselves. Give ML enough experiences, and it will start to make predictions about what comes next. At its core, ML involves the use of algorithms to process data and derive insights or "learn" from it, rather than being explicitly programmed for every task.

Supervised Learning:

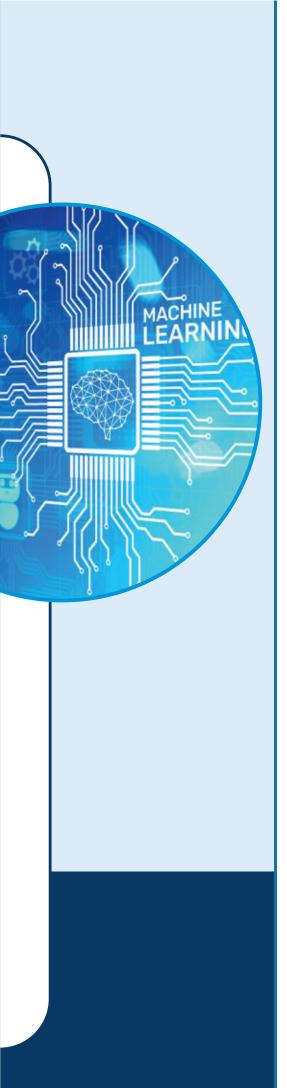
A Prediction Rule that Performs

A central objective in ML, particularly in supervised learning, is to create a prediction rule that performs well on new, unseen data, given a sufficient number of labelled examples. For instance, predicting house prices based on square footage in a specific zip code involves learning a prediction rule (e.g., a linear predictor) from existing size-price examples. The underlying assumption is that a well-fitting and parsimonious rule learned from observed data will maintain its predictive accuracy on new data. Other tasks, such as classifying spam emails, also benefit from extensive labelled examples to ensure robust predictive performance. ML algorithms essentially function by identifying patterns and making predictions, often without necessarily delving into causality or underlying latent structures.



How supervised learning uses labeled datasets to train models that can then predict labels on new, unseen data

There's something reassuring about a car that seems to understand its surroundings. Imagine it rolling through a busy street, where people cross at odd intervals and cars never quite stick to their lanes. What's remarkable is how it behaves, almost as if it remembers every lesson taught to it. It's seen thousands of examples each tagged by carefully with the right labels and action-"for crosswalk sign car needs to slow down and pause for pedestrians to cross". Each time it spots someone waiting to crosswalk sign, these actions kick in. The car eases off the accelerator, stops where it's meant to, and you notice it keeps you passenger or pedestrian protected. Its judgment isn't magic. It's the result of hours, maybe years, spent studying what makes a street safe, all stitched together into responses that feel natural and attentive.

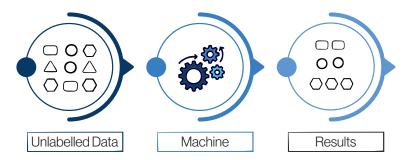


Unsupervised Learning:

The Detective in the Data

Another fascinating approach used in ML is also known as unsupervised learning, where, unlike methods that rely on human-labelled data, unsupervised learning sets the algorithm on a journey of its own: tasked with uncovering hidden patterns within vast collections of information, entirely without human guidance. Imagine stacks upon stacks of complex, unorganised data. Somewhere in that overwhelming heap lies a subtle clue, perhaps a suspicious financial transaction or an unusual user behaviour on a network. To us, it might be like searching for a needle in a haystack, but this is where the unsupervised learning approach of ML becomes extremely useful.

One of its most powerful use cases is in fraud detection whether in cybersecurity or finance where it notices deviations in patterns that the human eye might easily miss. Equally impressive is its role in data analytics, where it quietly maps out unexpected relationships between data points, revealing connections that were never obvious before. All of this happens simply through careful exploration: no instructions, no labels, just the algorithm's own ability to learn and organise the chaos into something meaningful.



The machine groups unlabeled data based on similarities, revealing hidden patterns and clusters without human-provided labels

Not every scene is predictable. Imagine driving past a market people crowd the curb, someone rolls a cart across, bikes are parked in odd spots, and there's an impromptu drum circle you never expected. No one sat down and labeled all these things for a computer's benefit. Instead, the car needs to figure it out on the go.

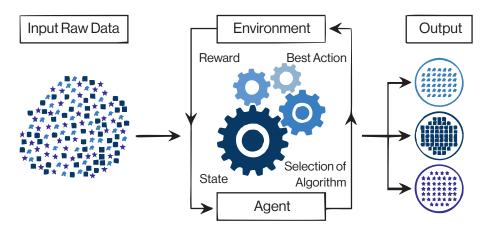
It starts seeing patterns groups of people here, odd shapes there trying to make sense of things it hasn't been told about. When something doesn't fit the familiar patterns, the car plays it safe, slowing down and paying extra attention. Honestly, it's not too different from how we might pause and take a second look at unfamiliar commotion. This ability to find structure in chaos lets the car handle the unexpected, making sure no surprises lead to careless mistakes.

Reinforcement Learning:

The Apprentice Learning by Trial and Error

Another approach generally used in ML is an amalgamation of both supervised learning and unsupervised learning called reinforcement learning. The algorithm for reinforcement learning is much like an eager apprentice, not given detailed instructions or predefined results but instead encouraged to experiment and learn through consequences.

In reinforcement learning, the machine learning model acts like an agent placed in an environment. It learns by trying actions, getting rewards for good moves, and penalties for bad ones. Through many cycles of trial and error, the ML model refines its approach, mastering tasks from navigating digital mazes to optimising logistics in the real world. This approach is especially powerful in fields where outcomes unfold over time, such as robotics, automated trading, or energy management systems.



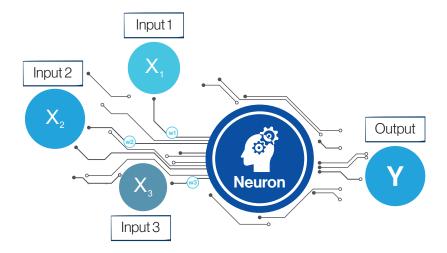
The agent learns optimal actions by interacting with the environment, receiving rewards, and refining its strategy to turn raw input data into desirable outputs.

Getting through a busy junction can feel like a dance. The car moves forward, holds back, and occasionally makes a bold move to join the flow. It didn't learn this confidence overnight. Instead, its actions are shaped by trial after trial, some ending well with smooth merges and others punished with awkward stalling. Every good move gets stored away as a "do this again"; every poor choice becomes a lesson. Over time, the system develops its rhythm, knowing when to step up and when to fall back. Watching it, you might think of a new driver slowly becoming confident, learning unwritten road rules and the difference between theory and improvisation. The car's ongoing education reminds us why patience and practice are so important in mastering street savvy.

Artificial Neural Networks:

The Brain-Inspired Engines of Al

Artificial Neural Networks (ANNs) are computational models inspired by the structure and function of the human brain's neural networks. A biological neuron, a fundamental unit of the brain, transmits and receives signals to and from thousands of other neurons via connected synapses. Similarly, in an ANN, artificial "neurons" process inputs. Inputs are approximately summed, and if this sum exceeds a certain threshold, the neuron "fires" or sends a signal to subsequent neurons. This interconnected structure allows ANNs to learn and adapt by adjusting the "strength" of connections between neurons based on training data. ANNs are important because they can search through large amounts of possibilities, use shortcuts to find solutions, handle incomplete or conflicting information, and learn and adapt over time.



A simple Artificial Neural Network with three inputs (X, X_i, X_i) feeding into a single neuron, which processes the signals and produces an output (Y).

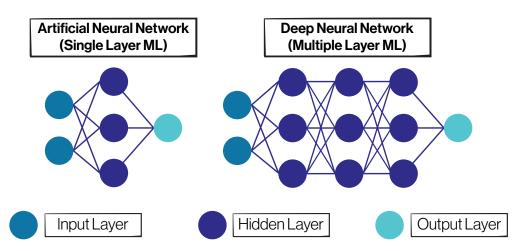
At a bustling intersection, so much happens at once kids wave from the sidewalk, dogs dart past, and a delivery van pulls a sudden halt. The car's neural network kicks in, sorting all this cluttered information into neat categories almost instantly. You can imagine someone glancing around, tuning out background noise to spot only the relevant details: a crossing signal, a cyclist's sudden turn, that excited dog's path. By blending signals from cameras and sensors, the system makes quick sense of a scene that would over whelm slower minds. What stands out isn't just speed, but the network's knack for extracting just what's needed and dismissing the rest much like someone who knows what to focus on in a crowd.

Deep Learning:

Layers of Intelligence

Deep Learning, a branch of ML and a subset under AI, has powered many of AI's biggest breakthroughs. Built on multi-layered ANNs, it excels at uncovering patterns in vast amounts of raw, unstructured data, images, sounds, and human language. Each layer adds depth to its understanding, enabling it to tackle problems once thought beyond machines.

Today, Deep Learning reads medical scans, guides autonomous vehicles, powers translation tools and voice assistants, recommends what we watch or buy, and even predicts equipment failures before they happen. DL sifts through complex data to detect subtle patterns and hidden signals, turning raw information into actionable insight. From healthcare to entertainment to manufacturing, it has become a quiet force reshaping industries and everyday life.



Comparison of neural network architectures — showing a simple Artificial Neural Network (ANN) with a single hidden layer versus a Deep Neural Network (DNN) with multiple hidden layers.

Driving late, rain streaking across windows, you squint at the faded lane marks while headlights bounce oddly off wet roads. That's when the car's deep learning tech steps in. It burrows through layers of unclear information, searching for clues: a dim line, a flash of a taillight, movement where someone might be stepping out. It's like an artist searching for shape in a muddy canvas, teasing meaning out of ambiguity. This isn't superficial its training has c overed countless situations so it can cope when things get messy. You realise, watching it work, that the car doesn't "see" conventionally. Instead, it's learned to find order where humans struggle, which makes dark, wet drives surprisingly uneventful.

One of the most widely recognised tools in deep learning is the **Convolutional Neural Network (CNN)** a kind of specialist in the art of "seeing." Just as our eyes and brain work together to pick out a familiar face in a crowded place, CNNs are remarkably good at spotting patterns in images and videos. They power facial recognition systems, help radiologists detect tumors in medical scans, and enable self-driving cars to interpret their surroundings.

CNN

For creativity, there are Generative Adversarial Networks (GANs)

imagine two artists locked in a friendly competition, one producing new images and the other critiquing them, each pushing the other to improve. This dynamic makes GANs capable of creating astonishingly realistic artwork, designing clothing prototypes, generating lifelike game characters, and even restoring faded or damaged photographs.

RNNs

At the leading edge are **Transformers**, the master multitaskers of Al. They excel at understanding

complex relationships in data, allowing them to drive today's most advanced language models. With this ability, they can answer difficult questions, condense long reports into concise summaries, and even

Next are Recurrent Neural Networks

(RNNs) and their more advanced

(LSTM) networks. These can be

thought of as the "listeners" and

remember context. They listen to speech so a voice assistant can

understand sequences and

respond naturally, translate

trends.

relatives, Long Short-Term Memory

"storytellers" of the Al world built to

languages instantly in conversation,

and even anticipate stock market

movements by studying historical

write functional software code.

Transformers

GANs

What sets all of these deep learning models apart is their ability to learn directly from raw data, adapt as they encounter new information, and operate without being given a rulebook for every situation. This adaptability has sparked breakthroughs in countless fields from detecting fraud in banking and optimising global supply chains to personalising shopping experiences, predicting equipment failures in factories, and speeding up drug discovery for life-saving treatments.

Computer Vision —

Giving Machines the Power to See

If we think of AI as a living creature, then ANNs are like its individual nerve cells, Deep Learning is the brain, and Computer Vision is the eye. Just as our eyes naturally notice and focus on the most important details around us, Computer Vision allows machines to pick out useful information from an image zooming in on what matters and ignoring the rest.

This field combines digital images and videos with the power of machine learning and deep learning. The result is a system that can identify and classify objects, understand scenes, and even make decisions based on what it "sees." In the case of video, the process begins by breaking continuous footage into individual frames. Each frame is then examined, with objects detected, categorised, and interpreted. A practical example can be seen in Alpowered traffic lights. The live feed from a road camera is divided into frames, where the system identifies different types of vehicles, notes their colours, checks occupancy, and estimates traffic density. Based on this real-time understanding, the timing of the lights can be adjusted on the spot to keep vehicles moving smoothly and reduce congestion.

The same technology is essential in autonomous vehicles. Here, an onboard computer constantly receives input from multiple cameras and sensors. Computer Vision processes this visual data to guide safe driving spotting pedestrians, reading road signs, assessing road conditions, and calculating distances to nearby vehicles. In essence, it gives machines the "sight" they need to navigate the world with awareness and precision.



Al segmenting and identifying objects such as vehicles, people, and roads in a street scene to enable automated understanding and decision-making in real-world environments.

As the car navigates narrow streets cluttered with cyclists weaving through traffic, children playing close to the curb, and delivery bikes stopping at random, its computer vision hums in the background. It's sifting through thousands of images every second, picking out a stop sign, a hand wave, a moving shadow that could be a running dog. Its skill isn't just about identifying objects, but piecing together a fluid story of what's happening all around. It's almost like a very sharp-eyed person who never tires, always on the lookout for the next change. You get the sense there's a guardian force behind each decision one that's careful, observant, and always scanning for that one detail everyone else might miss.

Generative AI in Action:

Transforming Ideas into Reality

Generative artificial intelligence (GenAI) is a subset of AI that creates novel content, including text, images, audio, video, code, and more response to user provided prompts. Unlike traditional predictive models that classify or forecast based on known data, generative models learn the underlying patterns and structures within large data sets and then synthesise new, coherent outputs that resemble the training examples.



GenAl blending technology and art to explore the boundaries of machine-driven creativity.

Generative AI is a type of AI that can create new things like text, pictures, or even music by learning from lots of examples. GenAI works with the support of deep learning, particularly with models like Generative Adversarial Networks (GANs) and transformers. Advances in cloud computing have made GenAI commercially viable and available It works in two main steps:

Learning (Encoding) – The AI reads or looks at a huge amount of information like books, photos, or computer code and turns it into a kind of "map" where similar ideas or things are stored close together.

Creating (Decoding) – When you ask it for something (a prompt), the Al looks at that map and uses what it has learned to produce new content that fits with what it has seen before.

Think of it like this: first the AI "studies" millions of examples to understand patterns, then it uses that knowledge to "imagine" something new that still makes sense.

GenAl has opened the door to creating entirely new forms of content whether text, images, videos, or even audio while also enabling fresh ways of exploring and analysing complex data. Its applications have quickly moved from research labs into everyday life, transforming how businesses and individuals approach problem-solving. In real-world scenarios, we already see GenAl at work In multiple ways.

It powers customer support systems that can address grievances and provide complete resolutions often without direct human intervention. Businesses use it to evaluate and optimise scenarios for cost reduction in areas such as marketing, advertising, finance, and logistics. In product development, it can rapidly produce multiple prototypes based on given inputs and constraints and then refine or optimise existing designs using human feedback. In the world of software, GenAl can suggest code snippets or solutions to speed up application development. It can also support decision-makers by automatically generating reports, summaries, and forward-looking projections, helping leaders respond faster and more effectively.

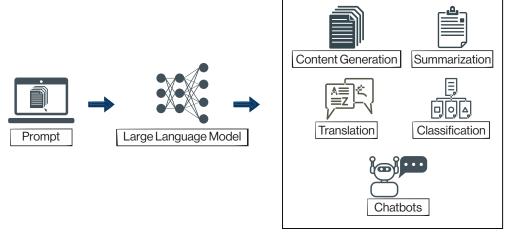
Across industries, generative AI models are becoming increasingly advanced. Current innovations focus on enabling these models to work securely with proprietary data while creating outputs text, images, video, and speech that are nearly indistinguishable from content made by humans. This progress brings not just efficiency, but also entirely new creative and analytical possibilities.

Before a single tire hits the road, the car spends hours in a world that's purely digital designed by generative AI. Imagine it navigating virtual rainstorms, dealing with crowds at simulated festivals, and dodging obstacles in test runs that mimic reality but stretch its imagination. These rehearsals aren't just theoretical; they shape the car's instincts, prepping it for rare, unexpected situations that real roads sometimes deliver. Onboard, a friendly voice assistant guides you through your trip, responding to simple requests and changing plans as needed. The blend of robust training and engaging interaction makes each journey feel more thoughtful and less mechanical. GenAI is the wizard behind the curtain, crafting both smarter vehicles and smoother experiences, unseen but essential.

Large Language Models:

The Powerhouses of Generative Al

A prominent example of GenAl is the Large Language Model (LLM). At its core, an LLM implements a simple yet powerful functionality: predicting the next word or "token" given a preceding context window. This is achieved by outputting a probability distribution over all possible words in its vocabulary, indicating the likelihood of each word following the given sequence. Examples of top generative Al tools include ChatGPT, Perplexity Al, Google's Gemini, Microsoft Bing Assistant, and DeepSeek.



Large Language Models (LLMs) transform a user's prompt into intelligent outputs—
enabling use cases such as content generation, text summarisation, translation, classification,
and chatbots across diverse industries and workflows.

A key advancement in LLMs, as seen with models like GPT-4, is in-context learning. This method allows pre-trained language models to be repurposed for specific downstream tasks by simply providing a few labeled examples within the prompt, rather than requiring computationally intensive adjustments to the model's parameters. This "conversation-like" interaction democratises the use of LLMs, making them highly attractive for end-users who do not need to customise the large model or run complex optimisation procedures. While fine-tuning may still offer performance gains, in-context learning significantly narrows this gap, making LLMs versatile for various applications.

Inside the car, conversations flow easily. Maybe you ask, "Let's take the scenic route and find a bakery along the way," or simply request, "Pause at the next gas station." The system doesn't just process the words it understands the intent, the mood, and makes practical decisions. If you're unsure, it'll ask for clarification, keeping things friendly and simple. Past requests and travel habits get tucked away for future rides, personalising the experience. It's less like talking to a robot and more like dealing with a polite, well-informed chauffeur who speaks your language and gets the nuances. This back-and-forth, full of small, human touches, turn severy ride into a partnership, not just a programmed journey.

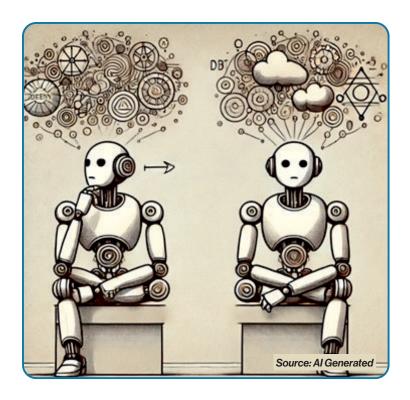
Large Concept Models:

Connecting the Dots at a Massive Scale

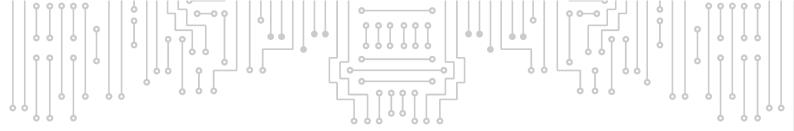
Imagine a person who has spent a lifetime learning someone who has read thousands of books, travelled across the world, studied science, art, and history, and can effortlessly connect a lesson from physics to a tale from ancient mythology. That's what a Large Concept Model (LCM) is like in the world of Al. It's not just about remembering facts it's about truly understanding ideas across different subjects and recognising how they link together.

Unlike AI models built for a single purpose LCMs are trained on vast and varied collections of information: text, images, diagrams, code, even sound. This diversity gives them the ability to think more broadly, make abstract connections, and apply knowledge from one field to solve problems in another.

If traditional AI is like a single-domain expert, an LCM is the "renaissance thinker" of the digital age a generalist that can move smoothly between topics, detect patterns most of us would overlook, and make connections that spark creativity. This makes it especially valuable for tackling problems where the solution lies in blending knowledge from multiple disciplines.



Approaching a tricky crossing, the car pays attention to the little things: a pedestrian who hesitates just outside the crosswalk, a cyclist who shifts suddenly, or even subtle changes in street noise. These bits of information aren't always clear or direct a person's glance, a hurried step, an oddly angled vehicle. But the system pieces them together, sensing which way things might go, and takes action well before a problem arises. Driving has always been about anticipating the unexpected, and here, the car feels almost prescient, adept at reading between the lines. This anticipation isn't flashy, but it's the quiet backbone of true road safety, letting technology blend seamlessly with instinct.

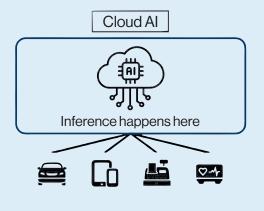


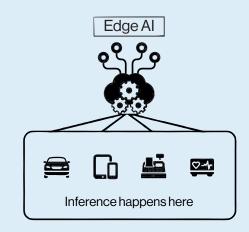
Edge Al:

Intelligence at the Source

Traditional cloud-based AI often works like a distant command centre data must travel back and forth before decisions are made. That delay can be costly when timing matters. Edge AI changes that. It's like giving your devices their own "mini-brains," right where the computation is happening. Smart watches doesn't need to send your heartbeat data to a server it can analyse it on the spot and warn you if something's off. Traffic cameras can spot congestion or accidents in real time and adjust signals immediately, without waiting for instructions from a central system. Even drones can navigate obstacles mid-flight without needing to "call home" for directions.

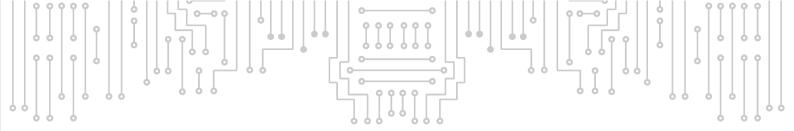
By processing data locally, Edge AI removes the lag, keeps sensitive data private, and makes AI feel more like a reflex than a delayed response.





Cloud AI performs inference on remote servers, sending results back to devices, while Edge AI executes inference locally on individual devices—enabling faster response and reduced reliance on constant internet connectivity.

Think of driving in a place with spotty cell service maybe along a rural stretch or through some mountains. Here, waiting for instructions from a far-off server could mean missing out on crucial decisions. Good thing the car's brain sits right inside, ready to analyse everything instantly. When a cyclist suddenly appears by the roadside or dense fog rolls in unexpectedly, there's no delay; the system on board responds at once, braking or steering as needed. What stands out is its independence it doesn't need to check in with a distant cloud to do the right thing. This self-sufficiency feels a lot more reassuring, especially in places where help isn't just a click away. Edge Al means peace of mind, knowing your technology won't let you down just because you've left the city behind.



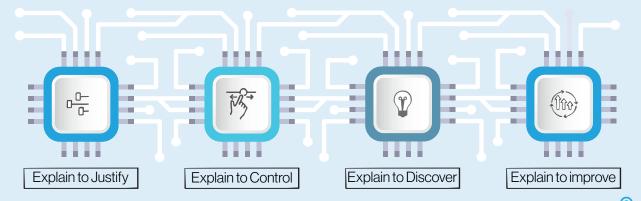
Explainable AI (XAI):

Opening the Black Box

Al systems can produce highly accurate results, but without insight into how those results were reached, it's difficult to know whether to trust them. Explainable Al (XAI) addresses this by showing the reasoning behind an outcome.

In healthcare, for example, an AI might flag early signs of pneumonia from a chest scan; and with XAI, it can also highlight the exact areas of the lung that led to that conclusion. In banking, rather than simply declining a loan, it can specify that factors like credit history and recent in come changes influenced the decision.

This kind of transparency isn't optional in many fields doctors need to justify treatments, financial institutions must follow strict regulations, and law enforcement agencies must ensure fairness. XAI transforms AI from a black box into a clear, auditable decision-making partner.



Explainable AI enables understanding and trust in AI systems through four key objectives: justifying automated decisions, controlling outcomes, discovering insights and knowledge, and improving models and processes for better performance.

.......

There's a certain relief in knowing why something happened, especially in unfamiliar situations. Picture riding in the car when, out of the blue, it brakes sharply at a quiet stretch. Instead of leaving you guessing or feeling nervous, a gentle alert pops up: "Stopped because a small animal was detected ahead." That transparency makes a world of difference. You aren't left wondering if a technical glitch occurred or if something's wrong. The car's reasoning is laid out plainly, helping you trust its choices. This pocket of honesty does more than soothe anxieties it provides engineers and developers with a window into how the system thinks. Safety checks become easier, and passengers can connect better with the technology knowing there's no hidden agenda, just thoughtful logic, always focused on real-world safety.

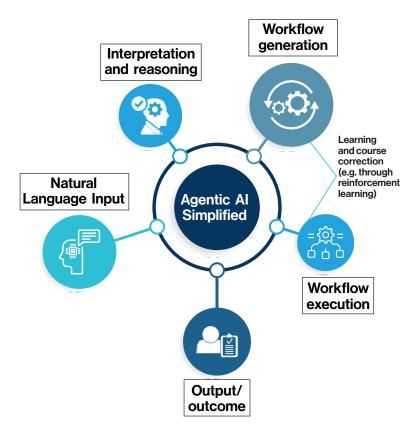
Agentic Al:

Al that takes the Initiative

Unlike traditional AI systems that wait for explicit instructions, Agentic AI can take initiative. You set the goal, and it determines the steps needed planning, acting, and adapting as conditions change.

For example, a fleet of warehouse robots might work together to sort packages. If one robot finds an aisle blocked, it can reroute itself and alert the others to avoid delays. A digital research assistant could take a request such as, "Summarise the latest studies on renewable energy storage," and instead of returning a list of links, it reads the papers, extracts the key points, compares methods, and identifies gaps in the research all without further prompting.

Agentic AI is particularly valuable in complex, unpredictable environments, from autonomous trading in finance to self-directing agricultural robots and virtual agents that manage customer support end-to-end. It's AI designed to move from passive responses to active problem-solving.



Agentic Al: Converts natural language input into outcomes by interpreting requests, executing tasks, delivering results, enabling human review, and continuously learning for improvement.

Sometimes roads get blocked maybe by a delivery van unloading, a maintenance crew, or just sheer traffic. The car, instead of sitting idle and waiting for explicit directions, analyses what's happening, weighs its options, and chooses its next move. It might nudge forward, wait for a gap, or signal its intent to merge. It's almost like having an assertive but considerate driver at the wheel, who reads the scene, anticipates what others might do, and acts accordingly. You don't have to micro manage or worry; the vehicle takes initiative, charting its own path through the everyday unpredictability of the road. This agency making choices independently feels like sharing the journey with a driver who isn't just skilled, but also wise in the ways of modern travel.

From Learning to Autonomy:

A Comparison of How AI Technologies Evolved and Connect

Al is a vast and diverse field, with each technology playing its own important role while working together to create smarter systems. At the base, we have ML, which is like a flexible learner that studies data and finds patterns to make predictions. It works well in many areas, such as spotting fraud or recommending products. But when the data gets complex—like images, sounds, or language DL steps in. DL uses layered neural networks to automatically understand subtle details without needing manual input. Neural Networks themselves are inspired by how our brains work, but they need large amounts of data and computing power to perform well.

Computer Vision builds on these to give machines the ability to "see" and interpret the world, helping in areas like medical scans, self-driving cars, and surveillance, where traditional methods fall short. Generative AI goes a step further by not just analysing but creating new content be it text, images, or code opening up exciting possibilities in creativity and innovation. LLM, a key example of generative AI, have become very powerful in understanding and generating human language, powering chatbots, translation tools, and even writing code.

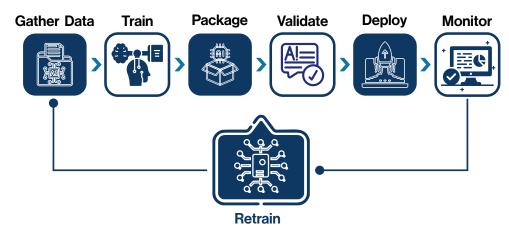
Meanwhile, Edge AI is about bringing intelligence closer to where data is generated on devices themselves. This helps with real-time decisions and keeps sensitive data private, without always depending on cloud servers. However, as AI models grow more complex, understanding their decision-making becomes difficult. Explainable AI tries to solve this by making AI's reasoning clearer and more transparent, which is crucial in sectors like healthcare and finance where trust matters a lot. Lastly, Agentic AI pushes AI towards independence it plans, acts, and adapts on its own, moving from being a passive tool to an active problem solver, useful in unpredictable and complex environments.

In short, ML offers adaptability, DL manages complex data, Computer Vision gives machines vision, Generative AI creates new content, Edge AI provides quick and private processing, Explainable AI builds trust, and Agentic AI brings autonomy. Together, they complement each other's strengths and overcome weaknesses, making AI systems more capable of understanding, creating, explaining, and acting in the real world.



Training and Testing Al Models: General Principles

The efficacy of AI models, irrespective of their type, hinges critically on the data they are trained on and the rigour of their testing.



Al Model Lifecycle



Training

At the core of every AI application lies the training phase, where the model is exposed to large volumes of data to learn underlying patterns, relationships, and features. For supervised learning, this typically means providing labeled datasets where inputs are paired with known outputs allowing the model to iteratively adjust its internal parameters to minimise errors between its predictions and actual results. This optimisation often involves sophisticated mathematical techniques such as gradient descent, stochastic gradient descent (SGD), and variants like Adam or RMSProp.

The importance of training on correct and relevant data cannot be overstated. Data must be carefully vetted and preprocessed to ensure accuracy, completeness, and representativeness. Errors or noise in training data can lead the model astray, resulting in poor or biased predictions. Because AI models often "learn" from patterns in their data, any flaws embedded in the data risk becoming amplified in their outputs.

Testing plays an equally critical role during development. After initial training, models are tested rigorously on unseen datasets to evaluate their generalisation capability. This phase helps identify weaknesses, such as overfitting to the training data or bias towards certain demographic groups. Based on test results, models often undergo retraining or fine-tuning where parameters are adjusted, additional data is incorporated, or algorithms are refined to improve accuracy and fairness. This iterative cycle of training, testing, and retraining is essential to producing robust Al systems that perform reliably in real-world conditions.

One common challenge is data scarcity when limited or incomplete data hampers effective learning. In such cases, techniques like data augmentation, transfer learning, or synthetic data generation may be employed to enrich the training set. Another major concern is bias amplification, where unrepresentative training data can cause models to perpetuate or worsen societal biases. For instance, a language model trained predominantly on Western text corpora may struggle to understand or appropriately respond to queries reflecting Indian demographics, geography, or cultural contexts. This gap can lead to irrelevant or even problematic outputs when deployed in diverse environments.



Testing serves as the checkpoint to verify that models not only perform well on familiar data but also generalise to new inputs. Performance metrics such as accuracy, precision, recall, F1 score, and fairness measures provide quantitative feedback on model quality. For LLMs, techniques like incontext learning allow testing by presenting new tasks as prompts and observing the model's ability to respond without further parameter updates.

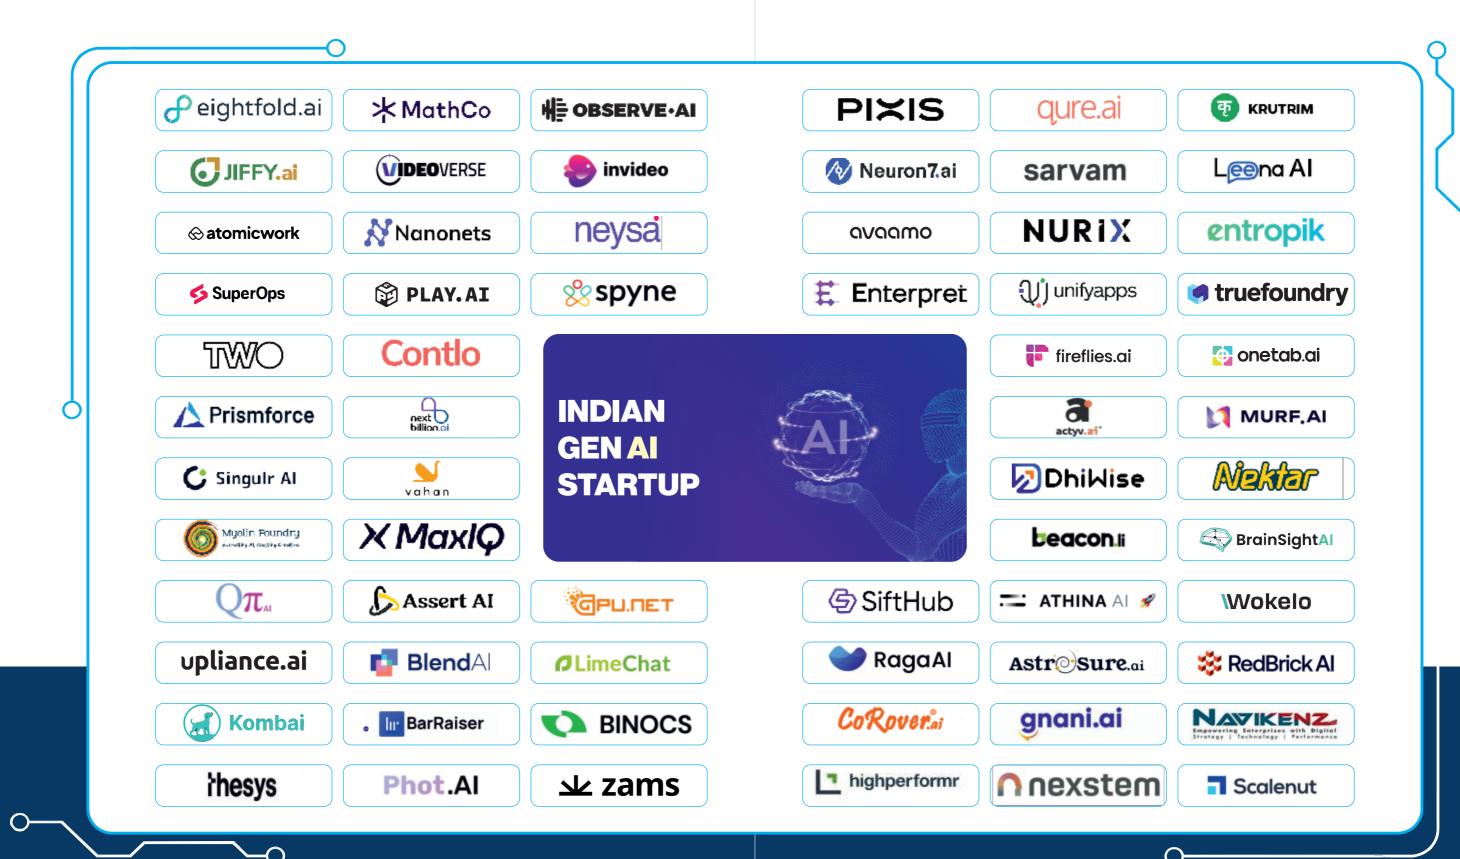
In essence, building effective AI systems is an ongoing process of training on high-quality, diverse data, testing for performance and fairness, and retraining to refine results. This cycle ensures models remain adaptable, trustworthy, and aligned with real-world needs.

From its theoretical beginnings over seventy years ago to its present role as a driver of global innovation, Artificial Intelligence has evolved from a bold idea into a practical, transformative tool. The journey has been marked by distinct phases from the pioneering work of visionaries like Alan Turing and John McCarthy, through decades of experimentation, to today's era of powerful, interconnected Al technologies that touch nearly every sector of society. Modern Al is no longer confined to research labs; it powers healthcare diagnostics, optimises transportation networks, supports climate modelling, strengthens governance, and fuels creativity in areas ranging from design to storytelling. The rich ecosystem of technologies Machine Learning, Deep Learning, Computer Vision, Generative Al, Large Language Models, Large Concept Models, Edge Al, Explainable Al, and Agentic Al each brings its own strengths, reinforcing the capabilities of the others. Together, they form a layered, adaptable system capable of observing, interpreting, creating, and acting in increasingly complex environments.

Yet, along with its promise, Al brings challenges that cannot be overlooked. Issues of bias, transparency, ethical deployment, and security highlight the importance of continuous oversight, responsible innovation, and inclusive access. As Al systems grow in capability, the emphasis must shift toward building trust ensuring these systems are fair, explainable, secure, and aligned with societal values.

Importantly, AI is now demonstrating exponential growth potential advancing not in slow, predictable steps, but with accelerating speed as computing power, data availability, and innovation compound on one another. For India's economy, this means more than technological progression it represents an opportunity for transformative growth. AI can boost productivity in manufacturing, agriculture, healthcare, finance, and governance; create new high-skill jobs in AI development, data science, and automation; enhance public service delivery; fuel a thriving tech-startup ecosystem; and position India as a global AI hub. If leveraged strategically, this exponential expansion can unlock significant GDP growth, foster innovation-led industries, and make India more competitive on the world stage. Looking ahead, the future of AI will be shaped not only by technological breakthroughs but by the way nations, organisations, and individuals choose to apply them. For India, initiatives like "AI for AII" reflect a vision of inclusive, sustainable progress using AI not merely for efficiency, but as an enabler of equitable growth and enhanced quality of life.

Al's story is still being written. The next chapters will depend on how we balance ambition with responsibility, capability with ethics, and automation with human insight. Guided wisely, the exponential growth of Al could become one of the defining forces in solving the most critical challenges of our time shaping a future that is intelligent, resilient, and inclusive.



ग्रेंड्रपुंबत वीhara | 23

Mapping the Al Ecosystem: Technologies, Trends, and Application **Early foundations include the Turing** Test (1950) and the coining of "Artificial Intelligence" at Dartmouth (1956) **Key Historical Milestones Progression through machine learning** and deep learning advances **Historical and Future Perspectives of Al** Emerging fields include explainable and agentic AI, neuromorphic computing, quantum AI, and edge intelligence Future trends point toward general Al **Emerging Areas and** with human-level cognition, **Future Trends** personalised medicine, integration with augmented reality and robotics, ethical frameworks, and self-improving

Healthcare uses Al for medical imaging, disease prediction, and virtual assistants. Automotive applications include **Healthcare and** autonomous driving, driver monitoring **Automotive** systems, and predictive vehicle maintenance. Finance benefits from fraud detection, **Real-World** Finance and **Artificial** credit scoring, and robo-advisers. **Applications** Retail/E-commerce Intelligence and Use Cases Retail uses Al for personalised recommendations (AI) and inventory forecasting **Energy sector applies AI to optimise supply** chains and smart grid management Security focuses on facial recognition cybersecurity and anomaly detection Energy, Security and General uses cases cover natural language processing (chatbots, sentiment analysis) **General Uses** predictive maintenance, customer Cases segmentation and medical diagostics. Perception focuses on vision and speech recognition for understanding input data. **Perception and** language Processing Language processing supports tasks like comprehension, translation and generation **Components of Encompasses general**

Problem Solving

and Optimisation

Techniques

AI Systems

Machine Learning includes supervised, unsupervised, and reinforcement learning for varied data-driven tasks. Machine Learning (ML) Deep Learning utilises neural networks like CNNs for images, RNN for sequences, autoencoders for feature learning, and GANs for generative tasks. Neural Networks vary from feedforward to advanced transformers, enabling complex pattern recognition. **Neural Networks** and Generative AI Generative Al focus on creating content such as text, images, and code, with applications in chatbots and media synthesis. Computer Vision handles object detection, facial recognition, **Al Technologies** autonomous navigation, and medical image analysis. and Sub-types Computer Vision, Large Language Models, and Large Language Models specialise in text generation, Edge Al translation, and summarisation. Edge Al enables on-device processing for fast, local decision-making in IoT environments. Explainable and This subfield emphasises Al transparency, trustworthiness, and Agentic Al autonomous agent functionality to improve user confidence and control Includes liner and logistic regression, support vector machines, decision trees, random forests, and ensemble methods like forests, and ensemble methods like XGBoost. **Machine Learning** Algorithms Clustering methods such as K-means support unsupervised pattern discovery. **Deep Learning Algorithms** and Optimisation **Core Al Algorithms** Covers CNN architectures (LeNet, AlexNet, ResNet), RNN variants (LSTM, GRU) transformers, and generative adversarial networks. Optimisation techniques include gradient descent methods (SGD, Adam) and Bayesian optimisation for hyperparameter tuning. -----Additional Anomaly detection algorithms identify unusual patterns. **Algorithms** Reinforcement learning algorithms like Q-learning enable -------Al to learn optimal actions through feedback

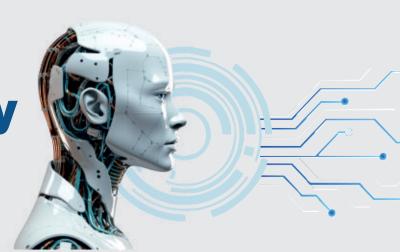
problem-solving capabilties

quantisation for efficiency.

Optimisation includes hyperparameter tuning, regularisation to prevent

overfitting, and model pruning or

Evolution of India's Al Policy and Strategies



ARTIFICIAL INTELLIGENCE TASKFORCE

Al Task force

2018



The Prime Minister's Science, Technology, and Innovation Advisory Council (PM-STIAC) Meetings on Al

The Ministry of Commerce and Industry set up a Task Force led by Dr. V. Kamakoti, consisting of experts from academia, research, and industry to explore the areas where Al could be leveraged for economic transformation, and to look at its various facets beyond the technology component.

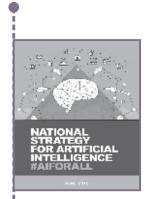
On 21 March, 2018, the task force released its first report acknowledging Al as a social and economic problem solver. It shed light on the Task Force's perception of Al, the sectors in which Al can be leveraged in India, the challenges endemic/native to India and certain ethical considerations.

The PM-STIAC serves as an overarching council to assess the status of various science and technology-related matters, comprehend challenges, formulate interventions, develop a futuristic roadmap, and advise the Prime Minister accordingly. The PM-STIAC in its initial meetings identified AI as one of its national missions and a key priority area. The council in its 21st Meeting held on 7 July 2022, recommended formation of the National Consortium and a dedicated Working Group to advance the deep-tech startup ecosystem, including AI.

India's Journey Towards an Al-Driven Future

As the world steps deeper into the age of Artificial Intelligence (AI), India has been steadily shaping its policies and strategies to build a resilient and forward-looking AI ecosystem. These efforts are not developed in isolation but are dynamically aligned with global trends and the pace of industrial adoption, ensuring that India remains both competitive and adaptive.

India's technological and research ecosystem for AI has been maturing for decades, with institutions like the Centre for Development of Advanced Computing (C-DAC), established in March 1988, laying early foundations in advanced computing and innovation. However, a structured approach to AI policy-making began in 2017, when the Government of India constituted a Task Force on AI India's Economic Transformation, led by Prof. V. Kamakoti. This initiative marked a turning point leading to India's current AI mission with the larger impact -bringing coherence, direction, and vision to India's AI journey by framing AI not just as a technological leap, but as a catalyst for economic and societal transformation.



NITI Aayog releases National Strategy for Artificial Intelligence (#AIForAll)



Suryanjay Singh Consultant, OPSA

2020

GPAL

India joins Global Partnership on Al (GPAI)

This strategy detailed how India might use AI for inclusive and social development, and work on maximising human capabilities and dealing with issues around accessibility, affordability, and skilled talent. In 2021, NITI Aayog released the "Principles for Responsible AI - Part 1" that handled issues of ethics and centering on implementation by policy interventions and capacity building.

On 15 June, 2020, India joined the GPAI as the founder. With aims to advance responsible development and use of AI with a focus on human rights, inclusion, diversity, creativity and economic prosperity; GPAI stands as a multinational and a multistakeholder project. After garnering over two-thirds of the first preference votes in 2022, India was chosen to take over as the Council Chair. India then served as the Incoming Chair in 2023, Lead Chair in 2024, and subsequently Outgoing Chair in 2025.

2023

Digital Personal Data Protection Act (DPDP), 2023

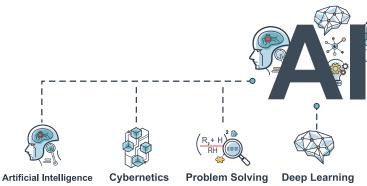


IndiaAl Mission

The Act aimed at regulating the "processing" of digital personal data. Since the DPDP applied to "processing" of personal data, any Al system that processes personal data, whether online or offline, is likely to fall under the purview of the DPDP Act. This is significant because Al systems now need to obtain explicit consent before using the personal data, unless there is an exemption. It also limited the extent of repurposing of data ie. data collected for one purpose cannot be reused for another purpose without further consent. This principle impacts the design and training of Al systems, ensuring careful consideration of data usage.

The IndiaAl Mission was approved by the Cabinet with an allocation of 10.371.92 crore for its seven core pillars, including IndiaAl Compute and IndiaAl Innovation Center. The mission envisions to build an ecosystem that helps Al innovation by democratising computing access, enhancing data quality, developing indigenous AI capabilities, attracting top AI talent, enabling industry collaboration, providing startup risk capital, ensuring socially impactful AI projects, and promoting ethical Al.

Set up in 2025, as part of the Safe and Trusted Pillar of the IndiaAl Mission, the IndiaAl Safety Institute focuses on advancing indigenous, science-based research on AI governance, pioneering innovative approaches to AI safety that are relevant to the developing world. It follows a hub and spoke model, where it will engage with multiple stakeholders including industry, academia, startups and Government Ministries/Departments to ensure safety, security and trust in Al.





Machine learning





Robotics

Neural networks



Al for India-Specific Regulatory Framework



India Co-chairs
Al Action Summit
(Paris)

In 2025, an Advisory group undertaking the development of "Al for India-Specific Regulatory Framework" was constituted. The advisory group chaired by the Principal Scientific Adviser (PSA) to the Government of India formed a subcommittee on 'Al Governance and Guidelines Development'.

The subcommittee's report highlighted the importance of coordinated, whole-of-government approach to enforce compliance and ensure effective governance with the evolving AI ecosystem in India.

The Al Action Summit is an international platform where world leaders, policymakers, technology experts, and industry representatives gather to discuss Al governance, ethics, and its role in society. The Summit in Paris was the 3rd in succession following the Bletchley Park Summit, United Kingdom, 2023 which talked about tackling the potential dangers of general-purpose AI and the need for international cooperation to ensure its safe development and deployment; and the Seoul Summit, South Korea, 2024 which largely discussed the responsible development and deployment of frontier AI, including risk management, accountability, and transparency. India co-chaired the Al Action Summit, Paris in February 2025, to advance inclusive and sustainable AI for people and the planet. India emphasised open, accessible AI in the service of national priorities. Drawing on its experience with responsible AI through endeavours such as IndiaAl Mission and GPAI, India advocated for bridging the digital divide, promoting ethical AI, and enhancing international cooperation for trustworthy AI ecosystems.

India to Chair the 2026 All Impact Summit in New Delhi

Building on the momentum of recent global AI dialogues, spanning the UK (2023), South Korea (2024), and Paris (2025), India is set to chair the upcoming AI Impact Summit in February 2026, New Delhi.

This high-level gathering signals a strategic evolution in the global conversation on Al. While earlier summits concentrated on frameworks for action, India's presidency will emphasise the transition from "action" to "impact." The focus will be on driving tangible outcomes, strengthening international collaboration, and ensuring the responsible deployment of Al at scale.

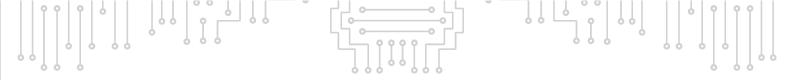
By hosting this summit, India positions itself not only as a beneficiary of Al innovation but also as a proactive contributor to shaping a globally inclusive, ethical, and impact-oriented Al ecosystem.

Understanding India's Al Ecosystem:Featured interview with CEO of

IndiaAl Mission



The detailed interview can be accessed on the given QR.



DECODING ARTIFICIAL INTELLIGENCE

Dr. Chagun Basha:

How do we decode AI as a disruptive technology, and what it means for a larger audience and society?

Abhishek Singh:

Al as a technology has existed for a very long time. The term was first coined in 1956, and since then, extensive work has been done in this field. Al evolved from basic data science to machine learning, then deep learning, and now, we are more familiar with its productised version known as Generative Al. Much of the recent interest was sparked by the ChatGPT movement in November 2022, when people began recognising that they were actively using Al. However, Al has long been part of many everyday experiences, whether in image recognition, shopping on Amazon, or content recommendations on YouTube and Netflix.

Al has been applied in various ways for many years, but the ease of use and the ability to create content such as videos or images have increased public interest. While the initial surge was driven by Generative Al, people are gradually realising its broader applications across sectors like healthcare, agriculture, manufacturing, and education. Al holds the potential to transform and disrupt nearly every sector of the economy. That's the power of Al, creating more opportunities and widespread interest in its possibilities.

In your opinion, which are the key sectors that would really benefit from large-scale adoption of AI, particularly in India?

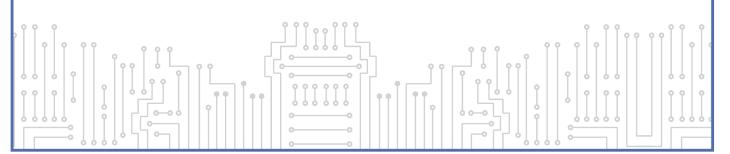
Abhishek Singh:

Natural language processing (NLP) is one critical sector where AI is already making a significant

impact, enabling Indians to access information and services in their native languages. In July 2022, we launched Bhashini, which supports speech-to-speech, speech-to-text, and text-to-text translations in all Indian languages. This platform allows users to access services in the language of their choice, supporting voice-enabled interactions.

In a country of nearly 1.4 billion people, around 900 million use online services, but approximately 500 million still do not. Bringing these citizens into the digital economy requires services in their native languages, particularly voice-based access. For example, a farmer in a rural area may seek information about crop prices, pesticides, or irrigation but might find it challenging to type URLs or navigate through apps. Voice-enabled services let users ask their query in their mother tongue and receive a spoken response in the same. Bhashini's APIs make this possible. Even without a smartphone, users can call a toll-free number, ask a question via speech, and get a real-time voice reply. This can empower, for instance, a woman in Jharkhand's tribal areas or a mother with a sick child in a remote location with limited internet or healthcare access, hence democratising access to vital information.

NLP is just one example of Al's wide-ranging scope, which extends across healthcare, agriculture, and more. Applications of Al include diagnosing diseases like tuberculosis or cancer through image analysis, overcoming radiologist shortages, predicting epidemics, facilitating drug discovery, advancing science, enhancing farming techniques, creating personalised learning plans, and supporting children with special needs. Al's interdisciplinary ability offers immense potential to transform services and improve lives.



There have been concerns about AI replacing the workforce, however these perceptions are now evolving. Could you talk about ways in which AI is empowering the workforce?

Abhishek Singh:

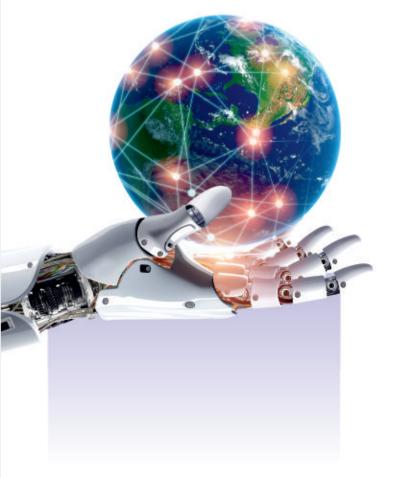
As with any technology, when a new one emerges, it replaces certain existing jobs or professions. Just as digital photography made traditional photo labs nearly obsolete, and automobiles replaced horse- and human-pulled rickshaws, Al's efficiency in performing tasks means some routine jobs may become redundant. This shift will require large-scale reskilling and upskilling to enable those affected by the technology. They will have to acquire new skills and participate in the growing data and Al economy, which is expected to generate many new jobs.

As AI advances from general AI to artificial general intelligence (AGI), Agentic AI, and physical AI, new roles will emerge in robotics,

data science, and data annotation. With proper training, we can harness these opportunities. Integrating AI education into school curriculums is also essential, so future generations are prepared for AI-related careers.

Al also enhances individual productivity. Tasks that once took six hours can now be completed in thirty minutes using tools like ChatGPT or Generative AI, boosting efficiency and income. For example, Al accelerates video editing, helping users earn more by saving time and reducing costs. For farmers, Al improves water and fertiliser management, which boosts crop yields and increases income. In MSMEs, Al enhances inventory tracking, price forecasting, and targeting the right markets efficiently. These gains in efficiency and income across sectors will cumulatively boost GDP and per capita income, helping India realise the vision of a Viksit Bharat and become a developed nation by 2047, as envisioned by the Hon'ble Prime Minister.

GLOBAL VALUE CHAIN OF AI



How is India positioned in the global AI value chain, especially as countries rush towards taking the lead role in areas like AI governance or AI adoption?

Abhishek Singh:

Al is a global technology, but currently, most Al development is concentrated within a few companies and countries. The largest advancements in model development, for instance, primarily come from the United States and China. India's key strength lies in its talent pool. According to the Stanford Al Index 2025, India ranks first globally, and on GitHub, India holds a 20% share of Al projects, ranking second worldwide. However, there is room for improvement in our research ecosystem and in the level of investment in Al. The IndiaAl Mission aims to address these gaps by providing affordable compute access to researchers and funding foundation models and applications, with the goal of being at par with global leaders.

We recognise that if AI models are trained on non-Indian datasets and deployed in India, they may adversely impact users in the country. Most global partners, including India, agree on the need for a global governance framework. This is important because models can be developed in one geography, applied in another, and used elsewhere. Such a framework needs to ensure that Al is safe, trustworthy, responsible, and ethical, supported by unbiased and representative datasets, rigorous testing, and minimal harmful effects. For example, Generative AI tools that depict women burning firewood when asked for a traditional "Indian kitchen" image reflect a bias. While these depictions may have been true in the past, they do not represent present realities. To address such biases, we need fairness assessment tools. We need mechanisms for detecting algorithmic biases, generating synthetic data, preserving privacy, and jointly combating harms caused by deepfakes and misleading Algenerated content. These harms pose risks including child sexual abuse, pornography, and financial fraud; hence a global governance framework is essential.

Given India's large user base and the high stakes involved, the country has taken a proactive role. As chair of the Global Partnership on AI (GPA) in 2023, India collaborated with the Organisation for Economic Co-operation and Development (OECD) and other partners to define a global Al governance framework. India also participated in the G7 Friends Group and contributed immensely to the Safety Summit Series, from meetings at Bletchley Park (November 2023) and Korea (May 2024) to the Al Action Summit in Paris (February 2025), co-chaired by Prime Minister Modi and President Macron. Key outcomes included establishing a forum to address current Al issues and launching sustainable Al initiatives. Now, India will chair the AI Impact Summit in February 2026, offering an opportunity to shape the global Al vision.

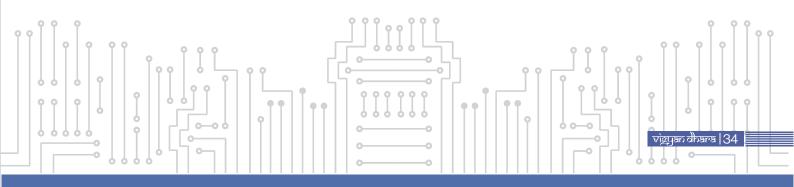
India has further contributed to UN deliberations on the Global Digital Compact, which aims to create a more equitable AI ecosystem by democratising AI resources, enabling access to computing power, datasets, algorithms, and models, and driving largescale capacity building. Many countries in the Global South look up to India for leadership and for joining the AI revolution. This momentum partly stems from our experience with large-scale digital public infrastructure, covering identity, payments, data empowerment, and service delivery in healthcare and agriculture, which has garnered global interest, especially after G20 (2023). Interest in India's approach to AI is growing worldwide, and the country is well-positioned to become a leading global provider of AI use cases and applications.

How do you assess the AI research and innovation ecosystem in India?

Abhishek Singh:

Many academic institutions in India have already undertaken AI development projects. The Department of Science and Technology has established technology innovation hubs, including the Technology Innovation Hub (TIH) at IIT Bombay, where Ganesh Ramakrishnan is leading the Bharat Gen project, and IISc Bangalore's ARTPARK initiative, which uses AI for robotics, among numerous others nationwide. The research community is eager to embrace AI initiatives, but limited funding has been a constraint, highlighting the need for increased investment.

With the establishment of the IndiaAl Mission, the Anusandhan National Research Foundation (ANRF), and the recently announced ₹1,00,000 crore research grant in the Union Budget (the RDI fund), funding for AI research is expected to improve significantly. Additionally, frameworks collaboration between academia, industry, startups, and venture capital must be developed, as academic research currently often occurs in isolation. Under the IndiaAl Mission, funding provisions for research and PhD scholars have been enhanced to align with research fellowships. Together, these measures aim to strengthen the research ecosystem and address existing gaps.



How do you see the India Inc. business enterprises and foreign tech corporations playing a role in India's AI ecosystem?

Abhishek Singh:

The global community looks at India as an AI hub, and as a key destination for sourcing most of their engineering talent. For example, Google DeepMind conducts a significant portion of its development in Bangalore under Manish Gupta's leadership. NVIDIA, which controls over 90% of the global chip market, has a development center in Pune with contributions from many Indian engineers. Indian engineers also make up nearly 20% of the global workforce in chip design. Major technology firms like Microsoft, Meta, and OpenAI have substantial Indian talent working on cutting-edge projects.

While global big tech companies will continue to show interest in India, it is crucial to encourage Indian firms such as TCS, Wipro, and Tech Mahindra, which have led global transformation initiatives in finance, transportation, and other sectors to also invest in and lead the Al space. Leveraging Al and taking leadership require sustained investment, including patient capital, because training Al models demands expensive GPUs. Indian startups like Sarvam and Socket are already investing heavily in Al, despite being relatively small. If these startups can take such risks, there is an opportunity and an imperative for India's larger companies to do the same. My message to the major Indian tech firms has always been to seize this opportunity; failure to do so will result in substantial long-term costs.





AI POLICY PRIORITIES

In terms of India's National Al Mission, what are the policy level priorities?

Abhishek Singh:

When we look at sustaining India's AI ecosystem, we look at the opportunity that lies with us. As I mentioned, our key strength is skills and talent. In 2023, we established seven multi-stakeholder working groups consisting of representatives from academia, the research community, startups, industry, and government. These groups conducted detailed analysis to identify gaps within the AI ecosystem. It became clear that building AI requires three essential components: skills and talent (which India possesses), compute resources, and data sets. However, significant gaps were identified in compute resources and data availability.

Two years ago, India had only 600 GPUs (A100 at CDAC Pune), compared to tens of thousands used by major companies like OpenAI, Meta, and Microsoft, which operate with thousands of GPUs (25,000 to 31,000) and plan to expand further. Addressing this, the AI Mission set a target of providing 10,000 affordable GPUs and, through public-private partnerships, achieved nearly 34,000 GPUs at Rs. 65 per GPU/hour, far below the global rate of \$2.5–3 per GPU/hour. This substantially resolved the compute challenge.

Regarding data sets, despite extensive digitisation, much of India's data was inaccessible or unusable. To overcome this, AI Kosh (aikosh.indiaai.gov.in) was launched. It hosts over 1,700+ data sets, 200+ models, and numerous tools and applications developed under the AI Mission, serving as a comprehensive repository of AI resources.

In addition to the two pillars addressing compute capacity and data sets, the third pillar of the mission focuses on building indigenous foundation models by supporting Indian startups to reduce reliance on foreign models such as Llama, Mistral, Claude, Gemini, and Grok. A call for proposals yielded 506 applications, and four were selected in the first phase to develop India-based large language models (LLMs).

The fourth pillar is the Application Development Initiative, which aims to build Al applications for farmers, women, healthcare seekers, students, and industry partners. An Innovation Challenge was launched around five themes: agriculture, healthcare, climate change, governance, and learning visibility. Thirty prototype applications were identified and will be scaled for national deployment as part of the Make in India initiative for Al. For example, an Al application designed to diagnose tuberculosis is currently in pilot phase but should be expanded to operate in real time across nearly all hospitals nationwide. Similarly, an Al tool for soil moisture assessment, linked to weather and crop stages, advises farmers on irrigation needs and is currently piloted but requires scaling.

The fifth pillar is skilling. While India excels in this area, ongoing support is crucial to maintain the advantage. Two main efforts are underway. First, fellowships are offered to undergraduate, postgraduate, and research scholars across disciplines including engineering, medicine, law, commerce, and humanities, reflecting wider Al adoption among diverse fields. An Innovation Challenge was also launched with Tata Memorial Hospital and the National Cancer Grid (CATCH grant) to fund startups and projects focused on Albased cancer care solutions, ten projects will each receive ₹50 lakh, with potential follow-up grants of ₹1 crore. Second, recognising that much of the Al

economy will involve data annotation and data science jobs, we are establishing over 500 data labs in Industrial Training Institutes (ITIs) and polytechnics across India, 25 of which are already operational. These institutions, which traditionally trained mechanics and printers, are now gearing up to meet demand in data science as traditional roles decline. A six-month, industry-partnered data science course with National Institute of Electronics & Information Technology (NIELIT) certification will prepare trainees for data annotation and analysis roles. The goal is to train 1.5 lakh students over three years. Hackathons and challenges are also being conducted, including collaborations with the Indian Cybercrime Coordination Centre (I4C) and the Geological Survey of India to develop Al-based mineral detection capabilities.

The sixth pillar is startup financing. Many AI startups require pre-seed and seed funding. To support this, a fund of funds is being established, and the IndiaAI Startup Global initiative has been launched. Through this program, ten startups were supported to attend an accelerator program at Station F in Paris, in partnership with HEC Paris, gaining exposure to European markets and building strategic partnerships.

The seventh and final pillar is safe and trusted AI, ensuring that AI solutions are responsible, ethical, trustworthy, and safe. To this end, an AI Safety Institute has been established, actively working on projects in machine learning, synthetic data generation, bias testing in data and algorithms, privacy-preserving technologies, deepfake detection and restriction, watermarking of AI-generated content, and development of ethical AI testing frameworks. These tools, developed in collaboration with academia and startups, will facilitate enforcement of regulations for unbiased, ethical, and responsible AI.

How do you see the large-scale adoption of Al and Al-based models in the government for e-governance?

Abhishek Singh:

Al is increasingly being used across multiple domains. A simple example is DigiYatra, which employs facial and image recognition to match a passenger's face with their Aadhaar details, and verifies the same against the boarding pass. This automates the CISF official's task and saves time.

Similarly, Jeevan Pramaan simplifies pensioner verification. Previously, pensioners had to visit government offices to submit a life certificate, which was especially difficult for the elderly. Now, using face matching and a liveliness test via mobile phone, pensioners can generate the certificate from home.

In law enforcement, Al aids crowd control by estimating crowd size at specific locations, and further helps optimise traffic management and prevent crime.

The startup Nayan Tech developed a city governance solution by engaging gig workers such as Zomato, Swiggy, Ola, and Uber drivers. Using GPS and phone cameras, they capture images of potholes, garbage dumps, missing manhole covers, or non-functional streetlights and traffic signals. These images, tagged with location data, are sent to control rooms for swift action. This method replaces traditional inspection teams with incentivised gig workers, improving efficiency. A similar system helps traffic police capture violations such as broken taillights, wrong-side driving, or red-light jumping, leading to automated challans and improved compliance.

At NIC, we offer Al-enabled services to institutions like the Supreme Court, where long-running petitions are common. Our tool summarises petitions, links them to related case laws, and assists judges. Al is also used to analyse documents of up to 500 pages. For example, Bhashini supports file notings and dictations in all Indian languages, enabling automated file preparation in native

languages.

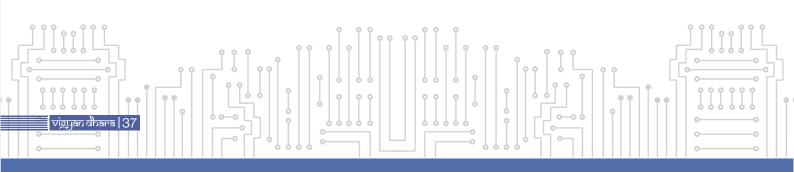
Police departments utilise AI for data analytics to enhance investigation reports. Defence forces employ AI in their operations, while AI-powered drones assist in surveillance. In agriculture, AI adoption is expanding across various applications. Overall, AI usage is growing rapidly across sectors. Going forward, AI will play an increasing role in addressing citizen grievances and improving governance and administration.

With respect to global standards and regulations, what are the standards that India is very keen on, and which standards are we pursuing for global interoperability?

Abhishek Singh:

Our principle from the outset has been to prescribe open standards and build interoperable solutions within the Digital Public Infrastructure (DPI) space. When considering AI, standards can be categorised into those for data sets, models, and other components. For instance, healthcare data sets adhere to Fast Healthcare Interoperability Resources (FHIR) standards, while agriculture follows global standards. India is represented in these standard-setting bodies through the Bureau of Indian Standards (BIS), ensuring that our specific requirements are integrated.

Regarding AI models, a significant debate exists between open source and proprietary approaches. Both coexist globally, but India prefers open source models while not prohibiting proprietary ones. Our regulatory focus is not on the technology itself but on its end use, particularly how the technology is applied. Our principle promotes pro-innovation and minimal restrictions: innovate with restraint; where restraint applies only to usage that could cause user harm. Otherwise, the default approach is to encourage innovation, as that is our strength. If we impose overly strict regulatory barriers, it will favor large players and limit the innovation that smaller companies and startups can achieve. This principle guides our approach to developing standards.



There was a Subcommittee report on Al Governance, under the guidance of an Advisory Group, which the PSA is chairing. What are the contours of that report and are we looking for that as a policy document for the country?

Abhishek Singh:

Yes, the committee thoroughly examined all aspects of AI development and regulation. It advocated a techno-legal and consent-based approach to data sharing and AI model development. The guiding principles were prevention of user harm and prioritising innovation without excessive restraint. The committee compiled its recommendations into a draft report that was opened for public consultation, which received a good response from stakeholders. The final report is expected to be completed within the next few weeks.

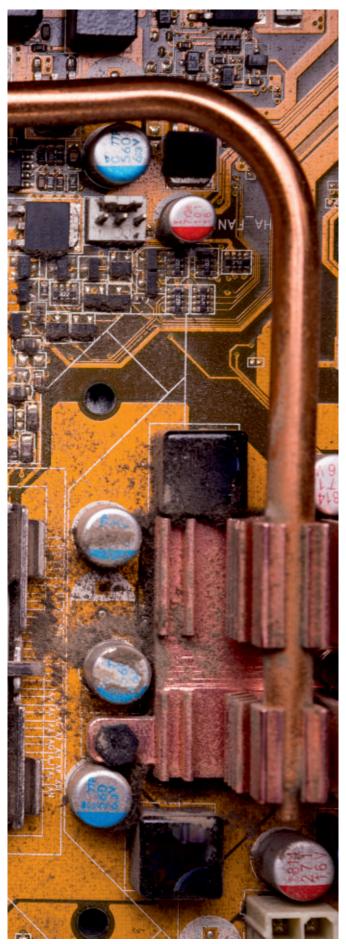
People talk about data sovereignty but now everybody is talking about Al sovereignty. How do we see Al sovereignty as a concept for India, given the national security implications that come with use of Al?

Abhishek Singh:

When discussing sovereignty in technology, it's important to consider sovereignty over the entire AI stack - from chips, data sets, and models to applications and use cases. India is strong in applications and is building indigenous models and hosting Indian data sets on its platforms. Indian companies manufacture servers, and we are progressing toward designing our own chips, although GPUs are still sourced from NVIDIA and others.

While relying on foreign GPUs for now, our servers are made in India, computation stays on Indian soil under Indian control, with opportunities for onpremise deployment. We emphasise hosting opensource foundational models within India to ensure no data leaves the country, while developing proprietary models for the future. Applications and use cases are built on Indian data, none of which is transmitted outside India.

This is the concept of sovereignty we are pursuing. We aspire for greater sovereignty, but until then, we focus on training and developing AI models and applications using Indian data, ensuring complete data residency. This forms the foundation of our approach today.



A Glimpse Inside a Graphics Processing Unit (GPU).

CONTOURS OF INDIA AI MISSION & AI IMPACT SUMMIT



How do we see the future of the IndiaAl mission? Do we have a phase two of the IndiaAl mission with some larger contours?

Abhishek Singh:

The IndiaAl Mission is currently one year old, with four more years ahead. We have made significant progress and expanded the program substantially. What began as a target of 10,000 GPUs has already grown to 34,000, and we expect to reach 50,000 by year-end, representing a fivefold increase within a year. In terms of applications, our initial targets have doubled, and for tools, we started with 20 but now anticipate needing over 100. We expect 5 to 7 indigenous Al models to emerge from India in the coming period.

Looking forward, we must address agentic AI and robotic AI. In Western countries, physical AI and robotics are prioritised due to limited workforce availability. In India, we focus robotics efforts only on areas where human labor is infeasible, such as hazardous jobs like explosives and landmine handling, or manufacturing sectors like electronics and automobiles where quality improvements are necessary. We avoid promoting robotics for tasks humans can perform, to prevent unnecessary unemployment. At the same time, we are aligning the engineering curriculum to train Indian engineers for global robotics job opportunities.

The IndiaAl Mission embraces a dynamic approach, recognising that technology evolves rapidly. When we began last year, H100 GPUs were state-of-theart; these swiftly became outdated and were replaced by H200, and now by B200. We continuously adapt while keeping our focus on delivering tangible benefits - improving education, healthcare, and agriculture. Our key questions are:

How is the common citizen benefiting? How is Al contributing to the Indian economy? Is it making the government more efficient? Is it reducing the drudgery of citizens and public servants? These are the questions we regularly ask ourselves as we assess our progress and plan future steps.

India will be hosting the upcoming AI Impact Summit. What are all the things planned for this major international event?

Abhishek Singh:

The AI Impact Summit is coming to a Global South country for the first time. After the UK, Korea, and France, India will host it. The main summit will be on 19–20 February in Delhi, attended by global heads of government and CEOs, with over 60 sessions across seven themes.

The core principle of the summit is impact, that is the central focus guiding our efforts. We frame this around three pillars, or sutras: People, Planet, and Progress. Within this framework, we have identified seven key themes, which we call the seven chakras. The first is Human Capital, which influences jobs, skills, and education. Inclusion is the second chakra, emphasising accessibility of Al including for people with disabilities and empowerment of women. The third is Safety and Trust, ensuring that AI remains ethical, responsible, and trustworthy. The fourth theme, Resilience, focuses on using AI to address climate change while minimising energy consumption to promote sustainable development. Science is the fifth chakra, acknowledging Al's role in advancing fields such as drug discovery and genomics. Democratising AI is the sixth theme, which aims to provide equitable access to computing resources and data for countries in the Global South. Finally, Al for Social Good applies Al to vital sectors including healthcare, education,

and agriculture. Together, these seven chakras span the three sutras and align under one guiding mantra: impact. This framework forms the comprehensive vision and plan for the summit.

In the run-up to the summit, nearly 200 side events are planned in partnership with academic institutions, government departments, industry, non-profits, and civil society. For instance, we will hold a research symposium featuring paper presentations from researchers worldwide. A global expo will showcase AI technologies for startups, including country pavilions and an exhibition at Bharat Mandapam. We are launching three hackathons: a global AI hackathon open to participants worldwide, a youth-focused hackathon to engage students, and a dedicated hackathon for women. Additionally, a pitch fest will be organised, inviting venture capitalists to facilitate fundraising. We are also inviting partners to propose their own initiatives for the event. The summit will feature keynotes and reports involving top CEOs and researchers, ensuring a rich and diverse program.

And while the main event is in Delhi, the side events will happen across the country. Globally as well, some partners will take it up. Everyone is welcome to check the impact.indiaai.gov.in website and contribute and join.

Do you have any final message that you would like to share with the ecosystem players?

Abhishek Singh:

Al is here, it is real, and is happening now. We are all using it and living in the age of Al. The cost of not learning Al will far exceed the cost of learning it. Therefore, investing in Al and investing in yourself is essential. Every student, individual, government department, business, startup, and academician must embrace this technology and benefit from it. Failure to do so will come with very high costs. This is the moment we are in, and I truly hope that we as a country will be able to take advantage of this and showcase its potential in the Al Impact Summit that we will be hosting in February.



PEOPLE • PLANET • PROGRESS

ABOUT THE SUMMIT

India will host the Al Impact Summit in 2026, building on the momentum of Bletchley Park, Seoul, and Paris to shape a future-oriented agenda for inclusive, responsible, and impactful Al cooperation.

Scheduled on the 19th and 20th of February 2026, the upcoming Summit will mark a strategic shift from "Action" to "Impact," focusing on tangible outcomes, global collaboration, and the responsible deployment of Al at scale. The Summit envisions a future where Al delivers real impact, advancing inclusive growth, social development, and a healthier planet.

Further details on the Al Impact Summit in India are included at page no. 113

OPSA's Role in Empowering India's Al Revolution

rtificial Intelligence (AI) is a branch of computer science that deals with the development of intelligent machines that can perform tasks that typically require human intelligence. Al involves the simulation of human intelligence processes by machines, especially computer systems. Specific applications of AI include expert systems, natural language processing (NLP), speech recognition, machine vision, and various applications augmenting life through human machine interactions.



The global view of AI is complex and multifaceted, with a variety of opinions and perspectives. AI evangelists are enthusiastic about the potential of AI and actively promote its development and use in various fields, while others express concerns and call for caution or regulation. The new applications of AI offer opportunities for increasing economic efficiency as well as quality of life and can be effectively used as a tool for social good.

The Government of India has recognised the potential of AI and has been actively working towards its development and application across various sectors. The country has a strong focus on research and development in AI, and the government has launched several initiatives to foster the development and adoption of Al. India, with its robust digital infrastructure, large talent pool of youth, and ambitious developmental goals, is seizing this moment to drive a national transformation, and OPSA is actively engaged in this transformational journey. The nation has been proactive in adopting this technology with a focus on fostering innovation while ensuring ethical standards. The Government of India initiated this path by releasing the National Strategy for Artificial Intelligence in 2018, titled "AlforALL," which emphasised inclusive and socially driven development across key sectors such as agriculture, healthcare, education, smart cities, and smart infrastructure. In the same year, Al was identified as one of the important national missions under the Prime Minister's Science, Technology, and Innovation Advisory Council (PM-STIAC), chaired by the PSA. Al remains an active topic of deliberation in PM-STIAC as well as the Technology Advisory Group (TAG) meetings under Empowered Technology Group (ETG).

India, by virtue of representing one-sixth of the global population, has a crucial role to play in shaping international discussions and policies on Al. As such, India has consistently engaged in the discourse pertaining to Al regulations. A wide range of developments have taken place in the Indian Al landscape in recent times. The National Strategy for Artificial Intelligence, focused on fostering Al development in specific sectors. This strategy laid the foundation for further advancements in utilising the applications of Al; however, it did not directly put forth Al regulations. In 2021, a draft for Responsible Al was introduced, featuring voluntary guidelines that prioritise ethical Al practices. These guidelines underscored the importance of transparency, accountability, and fairness. While enforceable regulations were absent, the emphasis on ethical considerations was a significant step forward. In March 2021, the Indian government adopted a "light touch" approach to Al regulation, demonstrating a thoughtful and measured approach to ensure responsible Al usage.

Overall, India has, thus far, taken a pro-innovation approach towards Al. The government, through the seven pillars of IndiaAl mission, intends to create an enabling Al ecosystem in India.

Operationalising AI Governance:

A Hybrid Techno-Legal Approach

Recognising the need for an India-specific approach to Al governance, the Ministry of Electronics and Information Technology (MeitY) established an Advisory Group chaired by the PSA to guide the development of a comprehensive regulatory framework for Al in India in November 2023. The Advisory Group recommends the government on Al governance and provides insights into the India-specific regulations required to ensure the ethical and responsible advancement of Al technologies. A subcommittee under the advisory Group examined key challenges, conducted a gap analysis of existing frameworks, and proposed recommendations for a robust, trustworthy, and accountable Al ecosystem.

The gap analysis by the advisory committee underscored the opportunity to strengthen existing legal and regulatory frameworks related to deepfakes, malicious content, cybersecurity threats, and IPR violations enabled by rapidly evolving AI technologies like generative AI. Additionally, concerns around AI-driven bias, discrimination, and a lack of transparency and accountability across the AI ecosystem were also highlighted.

In this context, the sub-committee outlined eight foundational principles to guide AI governance in India. These include Transparency, Accountability, Safety, Reliability & Robustness, Privacy & Security, Fairness & Non-Discrimination, Human-Centred Values with a focus on the 'Do No Harm' principle, Inclusive & Sustainable Innovation, and Digital-by-Design Governance. Aligned with the initiatives of the Organisation for Economic Co-operation and Development (OECD), NASSCOM, and NITI Aayog, these principles are intended to promote the development of AI systems that are secure, ethical, fair, and serve the broader public interest.

To effectively operationalise these principles, the subcommittee report emphasised the need for adopting a hybrid, techno-legal, lifecycle-based, whole-of-government approach towards AI regulation. Creating such a regulatory framework necessitates involvement of all AI ecosystem actors from data principals, developers, deployers to end users. This would include voluntary commitments on transparency by industry, incident reporting mechanisms, clearly defined regulatory responsibilities, and technology-driven compliance overseen by an inter-ministerial committee. These efforts support the Sovereign AI vision, ensuring intellectual property and sensitive data remain secured domestically while developing AI that's tailored for India's unique needs.

IndiaAl Mission

Building on the discussions held initially in PM-STIAC, the IndiaAl Mission, a national initiative to develop a robust and inclusive Al ecosystem, was introduced. The Cabinet approved the IndiaAl Mission on 7th March 2024. The mission aims to harness Al's transformative potential across sectors and democratise access to computational resources, improve data quality and access, foster indigenous Al capabilities, attract top talent, support startups through risk capital, encourage industry collaboration, promote socially impactful Al projects, and ensure ethical Al development and use.

Additionally, the IndiaAl Mission also includes initiatives like the Al Competency Framework for public officials and a Global Acceleration Program for startups, aiming to expand Al education, talent, and mentorship. Recently, IndiaAl mission also launched a call for proposals for the Safety Institute, which will work with all relevant stakeholders, including academia, startups, industry and government ministries/departments, towards ensuring safety, security and trust in Al. The Institute aims to advance indigenous research and development, based on Indian datasets and contextualised to India's social, economic, cultural, and linguistic diversity. These Institutes will be incubated by IndiaAl Mission and set up in a hub-and-spoke model with various research and academic institutions and private sector partners.

Collaborations under the IndiaAl Mission include the establishment of the Center for Generative Al, Srijan, at IIT Jodhpur in partnership with Meta, and the launch of the "YuvAi Initiative for Skilling and Capacity Building" with AICTE, empowering future Al researchers through open-source development and skill enhancement.

In May 2025, MeitY expanded national AI compute capacity to 34,333 GPUs through public-private partnerships, increasing access to high-performance computing, under IndiaAI mission. IndiaAI mission has supported 4 proposals to develop indigenous Large Language Models (LLM) and calls for proposals for supporting small Language Models (SLMs) are underway. National Centres of Excellence for Skilling will be created with global partners to equip youth with AI and skills. Additionally, MeitY selected eight responsible AI projects in October 2024, with half slated for deployment on AIKosh by December 2025, emphasising AI safety tools and benchmarks. OPSA actively participates in Program Monitoring and Evaluation Committees and screening of foundational and sector-specific AI models.

Further, the Department of Higher Education, Ministry of Education, has set up three Al Centres of Excellence (CoE) in Healthcare (IIT-Delhi and AIIMS), Agriculture (IIT Ropar), and Sustainable Cities (IIT Kanpur). Additionally, a fourth CoE for Al in Education with a budget of ₹500 crore was announced in Budget 2025.

Towards a responsible and collaborative AI ecosystem

To strengthen the responsible development of AI, OPSA is also playing a key role in international AI ethics discussions. In partnership with the Indian Institute of Science (IISc) Bengaluru and MeitY, OPSA organised two high-level AI policy roundtables with French counterparts that brought together senior policymakers, researchers, and industry leaders. The roundtables were organised during the AI Action Summit in Paris, France, held on 10–11 February 2025 co-chaired by Hon'ble Prime Minister Narendra Modi.

India's AI policy is also informed by extensive global engagement, ensuring alignment with international standards while pursuing national priorities. Active global collaborations spanning AI Policy Roundtables, the US-India TRUST, India-EU TTC, QUAD initiatives, and the AI Impact Summit - India, a sequel to AI Summit in Paris underscore India's dual ambition of being a knowledge hub and a responsible global AI leader.

With the focus on collaboration through the Triple Helix Model, the Government of India is also pushing research, development and innovation in multiple emerging technologies, including Al, ML and NLP. OPSA's Manthan Platform has continuously been at the forefront to facilitate such collaborations. To date, there have been more than 36 collaborative projects executed via Manthan Platform.

The World Economic Forum (WEF) under the guidance of OPSA and MeitY is developing a framework aimed at advancing responsible and impactful AI adoption. The AI Sandboxes framework emphasises the need for establishing safe and structured environments where startups and innovators can test, validate, and refine AI solutions using real-world data, compute infrastructure, and regulatory guidance. WEF, under discussion with OPSA, has prepared AI Playbooks that serve as a comprehensive guide for sector-specific AI adoption covering areas such as agriculture, and micro, small, and medium enterprises (MSMEs). The Playbooks also include practical use cases, readiness assessments, and targeted policy recommendations to support effective and context-aware implementation.



Hon'ble PM Narendra Modi co-chaired the Al Action Summit in Paris 10–11 Feb, 2025 driving global collaboration on artificial intelligence.









Al Policy Priorities for India

I's potential to address complex challenges like building climate resilience to providing healthcare access is widely acknowledged. But, at the same time, it raises critical questions around safety and equity. Al is not just another emerging technology, it's a "general-purpose capability" that has an impact across sectors, institutions, societies and borders.

Even as use of AI continues to grow, defining "Artificial Intelligence" remains an ongoing conversation. AI encompasses multiple tools, techniques and methods - from machine learning (ML) and natural language processing (NLP) to computer vision and robotics. The core characteristics within these systems may share things like problem-solving or learning, but their applications and implications vary widely. This diversity in its use cases makes AI difficult to govern from a single perspective. Hence it is better to understand AI as a field of ongoing investigation and research that is shaped by new capabilities, shifting contexts, and evolving social expectations, rather than a fixed technology.



Definitions and Dilemmas: The Challenge of Categorising Al

In the global policy space, definitions of AI are often drafted to be forward-looking or "future-proof." However, the process of defining AI can also bring ambiguity. If definitions are too broad, they may inadvertently capture traditional software or rule-based systems. If too narrow, they risk missing emerging applications that don't fit neatly into predefined categories.

India's regulatory institutions have typically taken a technology-agnostic approach, focusing less on defining AI per se, and more on managing its outcomes. The recent sub-committee report on AI governance guidelines, which is chaired by the Principal Scientific Advisor (PSA) to the Government of India, reflects this perspective. While acknowledging the value of definitions for specific regulatory or risk contexts, the report also notes that rigid classifications may not keep pace with the technology's rapid evolution.

Instead of offering a singular answer, these deliberations raise broader questions:

When and where are definitions of AI useful?

What trade-offs do they introduce?

And how can they remain adaptable in a constantly changing technological landscape?



These questions remain open and worth deeper engagement.

The Broader Canvas of Al Governance

Before diving into India's specific policy focus areas, it's important to understand how governance in the AI space is being conceptualised. Governance encompasses a wide range of interconnected elements, enabling policies that shape the environment for innovation; standards that guide technical and operational practices; and legal frameworks that ensure accountability and trust. These elements don't operate in silos; they overlap, complement, and sometimes complicate each other. Striking the right balance requires continuous dialogue among stakeholders.

India's development-focused approach to AI governance is shaped by this complexity. It emphasises practical applications of AI in areas where public value is most visible such as healthcare delivery, financial inclusion, agriculture, and language technologies. Several regulators, including the Reserve Bank of India (RBI), Securities and Exchange Board of India (SEBI), and the Indian Council of Medical Research (ICMR), have issued sector-specific guidance to support responsible experimentation. At the same time, national investments in capacity building through the IndiaAI Mission, Digital India Initiative, and AI Centers of Excellence signal a long-term commitment to strengthening the ecosystem. These initiatives reflect a belief that AI must be built on a strong foundation of infrastructure, skills, and inclusive design.



Emerging Priority Areas for Al Policy in India

As part of a wider effort to engage with global Al governance discourse, the Office of the Principal Scientific Adviser (OPSA) to the Government of India partnered with the Indian Institute of Science (IISc), IndiaAl Mission, and Sciences Po University, Paris, to organise the 2nd India-France Al Policy Roundtable in February 2025. Held in Paris as a side event to the Al Action Summit 2025, the roundtable followed the inaugural edition in Bengaluru in January during Technology Dialogue 2025. These dialogues brought together experts from government, academia, industry, and civil society to reflect on shared challenges and opportunities in Al governance. Informed by consultations with a wide network of Indian stakeholders, several emerging policy priority areas were discussed - not as fixed agendas, but as starting points for deeper engagement. Some of these have been listed below:



Advancing Indigenous Foundation Models

Foundation models are capable of performing multiple tasks and adapting across domains, becoming central to Al innovation. Indigenous foundation models offer a strategic frontier for India to localise, secure, and align Al with national contexts especially in terms of language, culture, and data sovereignty. This leads to questions of how can India build and sustain its ecosystem of general-purpose Al systems? And what institutional and infrastructural pathways could support India's vision of Sovereign Al?



Strengthening Al Governance through a Techno-Legal Framework

Integrating technical safeguards directly into legal and regulatory structures is gaining traction. Techno-legal frameworks aim to embed safeguards and accountability into AI systems from the outset. The current thinking explores whole-of-government approaches, inter operability between legal and technical systems, and adaptive models that can respond to sector-specific risks. It becomes critical to deliberate on how these approaches can foster innovation while managing systemic risk?



Democratising Access to AI Infrastructure

Extending India's DPI (Digital Public Infrastructure) model into AI raises critical opportunities for equitable access through open-source tools, shared compute, and inclusive data ecosystems. This further leads to the exploration of how might a DPI-for-AI approach unlock regional AI development, both within India and globally?



Powering India's Al Growth through Sustainable Al Ecosystem

Sustainability in Al invites reflection not just on efficiency, but on equitable access and long-term ecological impacts. So, can India craft an Al development strategy that aligns with its renewable energy goals and positions it as a leader in green computing?



Building a Future-Ready AI Talent Ecosystem

With its strong base in STEM education and IT services, India is well-positioned to lead global AI skilling efforts. Expanding training through public-private partnerships and integrating AI curricula across education levels are ongoing conversations. But how can capacity-building be made inclusive, responsive, and future-ready especially for emerging and under served segments?



Balancing Data Sovereignty with Strategic Cross-Border Data Flows

Balancing openness with control remains a key issue in data governance. India's leadership in secure, interoperable digital systems positions it to contribute to global frameworks. As digital trade and AI development rely on data flows, how can data policies enable innovation and trade, while safeguarding national and individual interests and without creating fragmentation?



Enabling Widespread Al Adoption through Systemic Preparedness

Adopting AI across government services and enterprises requires robust data systems, skilled personnel, and clear ethical guidelines. As public and private institutions scale up adoption, what governance structures are needed to manage risk, ensure transparency, and support feedback loops?



Building Inclusive AI for Equitable Digital Transformation

India's social and linguistic diversity makes inclusivity both a necessity and an opportunity. Strategies like multilingual AI models, open-source development, and diverse datasets are being explored. What frameworks can help embed inclusivity-by-design across AI lifecycles and ensure AI reflects and serves all communities?



Shaping Equitable and Representative Global Al Governance

India's role in shaping global governance mechanisms for AI is still taking form. Questions around interoperability, ethics, access, and voice for developing countries are central to this conversation. What global norms are needed to ensure AI serves shared human goals rather than reinforcing divides?



Al Across Sectors: Mapping the Impact of Artificial Intelligence in India

hether it is waking up to the adaptive cooling of air conditioners, or a smart-watch tracking our steps and health, real-time directions on our commute, summaries of complex papers, or recommendations for our next watch or purchase, checking in with DigiYatra, automated customer support, smart home assistants or farmers diagnosing crop, AI has become a part of our routines. It is transforming every sector, raising the bar of efficiency and convenience, while driving innovation.



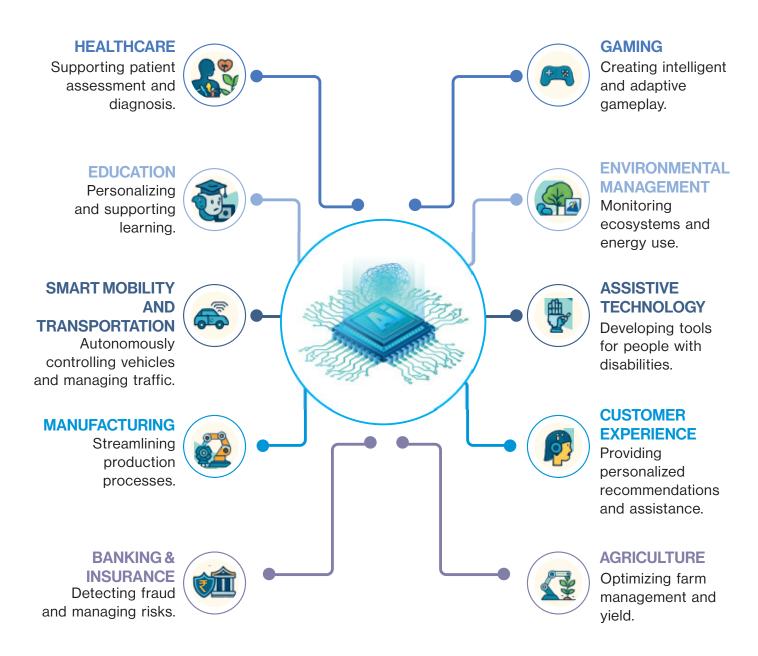
Talking about the influence of AI across sectors, Prof. B. Ravindran, Head, Department of Data Science and Artificial Intelligence (DSAI), IIT Madras highlights, "A new form of AI, i.e., generative AI (GenAI), has taken the world by storm. For the first time, we do not need to learn a new computing language to talk to the AI. Instead the language is human. This general-purpose technology has the potential to be an active, transformative force reshaping every sector of our economy and society."

As the world witnesses transitions and new models of problem-solving with AI, India's ability to integrate AI into its core sectors is also enhancing.

"In India, we are witnessing the deep impact of AI in language translation, where AI models trained on Indian languages are bridging the digital divide and enabling citizens to access information in their native tongues. Platforms like Bhashini and the research at AI4Bharat Centre at IIT Madras enable such innovations," Prof. Ravindran added. "Using these sovereign AI models, the government can enhance public service delivery at scale in areas such as grievance redressal, welfare delivery, and financial inclusion. For instance, UIDAI has launched AI-powered voice services in 10 Indian languages, making Aadhaar related help more accessible and adding real-time fraud detection for better security during authentication and updates. As we embed AI deeper into critical and social systems, it is essential to ensure these technologies are transparent, aligned with human values, and minimise harm. Responsible AI is not optional; it is foundational. India has the unique opportunity to shape AI not just for itself, but for the world."

Al is rapidly reshaping industry, research, and public service, with adoption expanding across a broad range of sectors. From advancing smart farming to driving breakthroughs in medtech, drug discovery, life sciences and more, Al has become a key driver of scientific progress. Its influence spans strategic areas like aerospace and defence, and digital domains like financial services, as well as cross-domain fields such as quantum computing, cybersecurity, and hardware development. As the country builds resilient disaster management infrastructure and fosters inclusive, ethical education and e-governance, Al is central to enabling data-driven insights, automation, and real-time decision-making for widespread impact.

Al is driving transformation across a wide array of sectors, from healthcare and agriculture to education, manufacturing, and banking. Its adoption is improving crop management, supporting smarter healthcare systems, enhancing personalised learning, and streamlining production processes. Al is also reshaping smart mobility, customer experience, assistive technology for people with disabilities, environmental management, and gaming. Al's integration into various domains is causing disruptive shifts in how sectors operate, innovate, and deliver value across traditional as well as emerging fields.



This article presents some of the advancements across agriculture, drug discovery, disease modelling, gene mapping, space, finance, quantum, cybersecurity, hardware product development, resilient infrastructure, education and e-governance. Through insights shared from experts at the forefront of AI integration, this curation highlights some key sectors where AI is making a significant difference. While these examples showcase AI's expanding potential, they offer only a glimpse of a much-larger, evolving AI landscape in India.

AI IN AGRICULTURE:

Transforming India's Farming Future

Agriculture has always been central to India's growth story and survival. India is home to vast agro-ecological diversity and is one of the world's leading producers of food. As of 2023-2024, the sector contributes nearly 18 per cent to India's GDP, employs about 40-45 per cent of its workforce, and feeds a population of over 1.4 billion, according to the Ministry of Statistics and Programme Implementation, Government of India. With 160 million hectares of arable land, India is the second-largest agricultural producer globally, leading in milk, pulses, and spices, while ranking second in rice, wheat, fruits, and vegetables. According to the Agricultural and Processed Food Products Export Development Authority (APEDA), agricultural exports crossed \$50 billion in 2022-23, underscoring its role in both food security and the economy.

India's agriculture sector demonstrates strong potential, supported by a large population of small and marginal farmers. Efforts to optimise fertiliser and pesticide use, promote sustainable irrigation, and diversify cropping practices are progressively improving soil health, groundwater conservation, biodiversity, and food safety. Expanding irrigation coverage is enhancing



Pushpendra P. Singh Project Director, ANNAM.AI

resilience to climate variability. Strengthening the link between agricultural productivity and nutrition offers promising opportunities to reduce food insecurity. As the population grows and arable land becomes a precious resource, innovative resource management and technology adoption will help India sustainably feed its people and ensure food security for the future.

This is where AI presents transformative opportunities. AI can enhance productivity, increase efficiency, and support sustainability in agriculture. Applications for these include AI-powered crop monitoring using drones and satellites, early detection of pests and diseases, precision farming to optimise water, fertiliser, and pesticide use, and predictive models for crop yield and weather forecasts. AI-driven automation can support irrigation, supply chain management, and sugar mill operations, reducing costs for farmers and industries.



Tools like livestock monitoring devices (e.g., Buff Sense and MoohSense) track animal health and reproductive cycles, while Bee Sense technology monitors beehives for early interventions to support pollination and biodiversity. Additionally, Al-driven plant health analysis, advanced weather stations, and agricultural data platforms provide farmers with real-time, data-based tools to make informed decisions. Together, these technologies can optimise yield, improve resilience to climate change, reduce environmental impact, and raise rural incomes.

Centres of Excellence in Al

Recognising the transformative potential of AI, the Government of India launched an initiative in 2023 under the Ministry of Education to set up Centres of Excellence (CoEs) in three key areas - healthcare, agriculture, and sustainable cities, with an outlay of 990 crores. After a competitive selection process, three consortia were chosen - IIT Delhi and AIIMS Delhi for healthcare, IIT Kanpur for sustainable cities, and IIT Ropar for agriculture. These CoEs are designed to serve as hubs for interdisciplinary AI research, combining academic institutions, industry partners, startups, and government agencies to deliver scalable, Indiaspecific solutions.

Centre of Excellence in Agriculture at IIT Ropar

The AI Centre of Excellence for Agriculture at IIT Ropar, ANNAM.AI Foundation, supported with ₹ 300 crores, is building solutions for India's farming sector. Guided by the vision to strengthen food security and improve farmer incomes, the CoE is focusing on projects across livestock management, plant health diagnostics, agricultural data resources development, crop yield forecasting, and crop classification through AI.

Our mission is to enhance food security and raise farmers' incomes using technology, thus positioning India as a leader in agriculture and food security," said Professor Rajeev Ahuja, Director, IIT Ropar.

Some of the key projects include:

Buff Sense: Uses tail sensors to monitor buffalo health and reproductive cycles, aiding livestock management, animal health, dairy productivity, and sustainable agriculture.

Bee Sense: Monitors beehives for early detection of issues, ensuring healthy colonies for pollination, biodiversity, and agricultural production.

AI-Powered Plant Health Analysis: Provides farmers with tools to monitor plant health, anticipate conditions, and reduce losses. AI-based chlorophyll sensors enable targeted insecticide use, reducing environmental degradation.

Advanced Weather Station: Captures real-time weather data to support precise farming decisions and resource optimisation.

Agriculture Data Resource Platform: Delivers timely Al-derived updates on soil, crops, and best practices, with a chat interface for easy farmer access.

Al-Driven Crop Yield Prediction: Uses satellite, soil, and meteorology data to forecast yields at the district level, enabling better harvest planning and resource management.

Crop Classification Technology: Classifies crops for precise resource allocation, tailored interventions, maximised output, and minimal misuse

This initiative aims to benefit 500,000 farmers and expand impact across crops, demonstrating Al's scaling potential in agriculture. By building partnerships with startups, industry, and global institutions, the CoE also supports agritech entrepreneurship and workforce skilling. By combining cutting-edge technology with farmer-centric solutions, India can strengthen food security, empower smallholders, and build a resilient, future-ready agricultural system.

AI IN MEDTECH & DRUG DISCOVERY:

Engineering the Future of Global Health

Health security today rests on two equally critical pillars: medicines and vaccines on one hand, and medical technologies on the other. While drugs treat diseases therapeutically, medical technologies (such as diagnostics, imaging systems, ventilators, oxygen devices, surgical robots, and wearable monitors) facilitate early detection, patient management, and lifesaving interventions. Together, these parallel streams form the foundation of modern healthcare.

What has changed in the last decade is the role of Al. Where traditional device development and drug discovery relied on long cycles of prototyping, testing, and approval; Al now enables researchers and engineers to compress these cycles dramatically, moving from years to weeks in certain cases. For a country like India, with its imperative of delivering affordable, scalable, and equitable health solutions, Aldriven MedTech innovation is not just an advantage but a necessity.

Across the MedTech landscape, Al is catalysing transformation. In diagnostics, machine learning



Dr. Jitendra Sharma MD and Founder CEO of AMTZ

models now optimise assays, refine RT-PCR primer design, and simulate mutations, ensuring testing technologies remain relevant as diseases evolve. In imaging, Al-enabled Radiology systems interpret scans faster and more accurately, reducing diagnostic errors. Device engineering has also been revolutionised, Alassisted 3D printing platforms rapidly prototype components for implants, prosthetics, and biosensors, dramatically cutting development timelines. In hospitals, Al-driven physiologic monitors predict patient deterioration, while wearable sensors provide clinicians with realtime data. Even in drug discovery, Al and MedTech work together; diagnostic tools guide clinical trials, and computational models help find new uses for existing molecules to treat emerging health conditions.



Recent developments redefining MedTech

India's leadership in this transformation is exemplified by the Andhra Pradesh MedTech Zone (AMTZ). Envisioned as a one-stop ecosystem for medical technology, AMTZ brings together research, manufacturing, validation, sterilisation, logistics, and customer support within a single campus, something unprecedented globally. Supported by 17+ Ministries of the Government of India, AMTZ has become a living laboratory for Alenabled MedTech resilience and innovation.

The specialised centers of excellence within AMTZ demonstrate how integration accelerates innovation.

- ADDIT, the Centre for 3D Printing, uses Alguided additive manufacturing to produce everything from high-precision implants to mass-manufactured device components.
- BIOME, the Biomaterials Testing Centre, employs Al-driven analytics to evaluate implants, PPE, and consumables, ensuring compliance with global safety standards.
- The ICMR Validation Lab applies data processing to validate diagnostic kits, adhering to international benchmarks of accuracy.

This consolidated approach redefines how medical technology innovation is delivered. Instead of scattered facilities working in silos, AMTZ demonstrates how scientific infrastructure, Al-enabled systems, and cross-sectoral collaboration can coexist within one ecosystem, improving efficiency, resilience, and scalability.

Globally, similar trends are reshaping healthcare. All breakthroughs such as AlphaFold2 in protein structure prediction, Al-guided biosensor design, and predictive maintenance for MRI and surgical robots are accelerating both pharmaceutical and device pipelines. In India, AMTZ and allied institutions like the BioValley Incubation Council, MediValley Incubation Council and World Health

Innovation Fellowship are nurturing startups that deploy AI in imaging, point-of-care diagnostics, prosthetics, and wearable health devices, areas that will define the future of personalised and preventive medicine.

Yet, as AI permeates deeper into MedTech, governance becomes paramount. Devices powered by AI must undergo rigorous validation to safeguard patient safety, particularly in life-critical environments like ICUs and operating theatres. Equity must remain central, AI-driven innovations should reach district hospitals and rural health systems, not only metropolitan centers. Transparency and explainable AI frameworks are essential to building trust, particularly when algorithms directly inform patient care decisions.

Towards the Future

The next frontier lies in anticipatory health systems. Imagine Al-powered portable labs that can detect pathogens in rural areas before outbreaks spread; smart ventilators that auto-calibrate to individual patient physiology; or wearable biosensors that pre-emptively flag risks of chronic disease, preventing hospitalisation altogether. These are not distant possibilities but near-term realities, and they align closely with India's vision of Viksit Bharat @2047.

Al in MedTech and drug discovery is no longer aspirational, it is operational, transformative, and accelerating. From Al-driven diagnostic validation and 3D-printed device prototyping to predictive monitoring and logistics, India has shown how innovation, policy, and collaboration can converge at scale.

The way forward is clear - expand indigenous AI-MedTech capabilities, embed ethics into innovation, and position India as the MedTech capital of the world.

By doing so, India can shape a future where AI not only accelerates discovery but also safeguards lives, for India and for the world.





AI IN DISEASE MODELLING

In the times where climate change and emerging pathogens could pose unparalleled challenges to public health, AI and related technologies, which are already supporting human decision-making across sectors, could emerge as a transformative advancement in our approach to safeguarding population health.

Yet, this pursuit of anticipating and controlling disease spread is not new. Understanding how diseases spread has long been a human pursuit – from Bernoulli's 18th-century calculations on smallpox to the compartmental models of the 20th century that formed the basis of modern epidemiology to that of sophisticated computation approaches that help anticipate, understand and control disease outbreaks.

However, although foundational in understanding disease dynamics, they are constrained by their ability to capture the complex dynamics of real-world disease transmission, homogeneity, fixed parameters, computational capacity, and the challenges of integrating multimodal, real-time data sources, evolving disease landscapes and policy priorities.

Potential Opportunities & Early Insights

It is in this context, AI in disease modelling could offer potential opportunities. At AI and Robotics Technology Park (ARTPARK), our pilot project on dengue outbreak risk forecasting demonstrated the use of data science and machine learning. We integrated epidemiological line list data with meteorological variables to enable early detection of dengue outbreaks and





Harish Nalawade Program Manager

Health and Climate Initiatives, Al & Robotics Technology Park, I-Hub @ IISc

identify temporal and spatial clusters. This information equips city administrators to strategically coordinate interventions such as fever surveillance, sanitation, vector control, and public awareness campaigns.

While this offered promising early insights, it also revealed key data gaps, particularly the absence of entomological inputs (vector surveillance) crucial for capturing vector densities and dynamics. Discoveries like this underscore the importance of approaches that can incorporate diverse and dynamic data sources.

The ability to ingest and iterate on vast amounts of heterogeneous data in real time, leverage advanced computational power, and build on flexible model architectures makes it possible to continuously refine predictions. This also opens up opportunities to incorporate non-traditional surveillance inputs such as social media signals and satellite imagery. These diverse inputs allow models accommodate real-world complexities, potentially extending the capabilities of conventional mechanistic and statistical approaches. It would facilitate translation into actionable intelligence, and strengthen the operational utility of data as well as model driven approaches in public health.

Further, while traditionally applied to infectious disease contexts, this Al-enabled modelling paradigm could now be extended to emerging climate-sensitive health risks. In particular, we are applying these approaches to extreme heat, where health outcomes are driven not by pathogens but environmental exposures temperatures, humidity, heatwave duration, etc.,), health vulnerabilities, and system-level capacities. Our current effort is to develop a deep-learningbased hyperlocal heat-health risk model that integrates health data with weather predictions to provide sub-district to village level forecasts up to 7 to 10 days in advance. The aim is to go beyond temperature thresholds by incorporating multiple meteorological parameters along with health and social vulnerability factors. This approach will enable effective adaptation measures and targeted interventions for the most at-risk populations.

Moving from Promise to Impact

For AI in disease modelling to move from promise to impact, it must be grounded in the realities of how data is generated, interpreted, and used. Currently, data is captured based on frontline reporting practices that reflect the realities of the health system. However, the frontline workers

collecting this data are often not the direct beneficiaries and may not understand how their work is being used. At the same time, modellers and decision-makers work with different assumptions, timelines, and data needs. For Al in disease modelling to be effective, these distinct actors must be seen as participants in a shared process.

The challenges of using AI in disease modelling reflect the broader challenges in AI itself. These go beyond data quality or algorithm performance, and include issues of representation, bias, responsibility, and explainability, as well as how AI models influence decisions and outcomes. In public health, these challenges become even more important.

It is crucial to envision AI not as a stand-alone layer, but as part of a system-integrated solution. AI models should sit within well-designed products that translate insights into actionable tools. These products must, in turn, be embedded within broader solutions – integrating both technological and non-technological components such as protocols, workflows, and human capacity. These solutions then operate within the broader systems and programmes, such as institutional design, public health policies, and governance structures.



AI IN GENE MAPPING

Over the last couple of decades, genetics has firmly cemented itself as a central part of biomedicine and human health. And now, the field is poised to see massive changes with the Al wave.

Human genomics specifically is a field ripe for disruptions. Genomics, the study of the entire genome as opposed to isolated genes, has many features that make it particularly suitable for many of the methods and tools of Al and machine learning. First, it has always been a field of 'big data'. Sequencing technologies continue to become cheaper and more efficient; we now have pocket-sized DNA sequencers that we can carry into jungles or into space. This has resulted in a goldmine of publicly-available data on genetic variation and genome function, which makes for a substantial amount of training data available. Second, the underlying language of DNA is an 'alphabet', which makes this a natural field to test the recently-successful large language models that work well on large bodies of text.

Understanding the Human Genome and the Challenges

The human genome is 3 billion letters long. Of these 3 billion positions, which ones are important for function? If we have a mutation at any of these, will it cause a disease? These questions are at the heart of human genetics. When a patient goes into the clinic with a rare disease, clinicians often turn to genome sequencing to understand the nature of the disease and its underlying cause. The patient's DNA is sequenced, and geneticists look for mutations found in the patient's DNA that aren't typically seen in healthy people. While this may sound simple, the reality is more complicated.



Dr. Shweta Ramdas Assistant ProfessorCentre for Brain research,
IISC Campus, Bengaluru

Every patient carries a handful of variants-whose effects on biological function remain unclear. We call these 'variants of uncertain significance'. How do we then know which ones are important and causal? Testing these out individually in the lab can be a long and expensive process.

The first few models aimed at these questions are already being tested. Alpha-missense, a descendant of the Nobel-prize winning AlphaFold, aims to predict the effect of mutations on protein structure and thus predict if an observed genetic variant is disease-causing. A more recent development is AlphaGenome, which tries to predict the function of every position in the genome, and not just the proteincoding region. It doesn't just take protein structure into account, but other information like how the genome folds, and how different parts of the genome are turned active in different tissues of the body. The hope is that models like this will be beneficial in rare disease diagnosis, where the current diagnosis rate from genetic data is about 40%. Having a genetic diagnosis can help patients find personalised treatments and intervene early, and even develop drug targets.

Generative AI can not only be used to discover patterns in the human genome, but use identified patterns to generate novel gene and protein sequences that are optimised for drug efficacy.

Another area with the potential for improvements is the use of non-genetic data to predict genetic information. Using gene expression data in cancer to find out which genes are differently active in different parts of a cancerous tissue, is used extensively to stratify or distinguish cancer patients, and personalise therapies. However, these gene expression datasets are expensive to generate. New research is showing that image data from patient tissue samples can be used to 'predict' their gene expression data to a reasonable degree, specifically for the genes that are involved in the underlying cancers. These approaches for cross-data predictions can eventually lead to rapid first-pass patient stratification that is affordable and more accurate.

Towards an Al-improved future

The path to this Al-improved future is still being paved. The improvements to the questions posed above so far have been incremental. With more datasets and more fine-tuning, these models are expected to keep improving. In the Indian context, we are likely to have larger datasets sampling a larger fraction of Indian human biodiversity over the next decade. This increase in data (which is a bottleneck currently) is likely to lead to better performance of foundation models applied to Indian datasets. Such efforts are also likely to require new models of doing science like teaming science with multiple complementary partners, often bridging academia, clinicians, and industry.







AI IN SPACE

India's space journey has always been marked by innovation, resilience, and impact. As the world enters a new era of space exploration, defined by small satellites, reusable launchers, and a growing role of private players, AI is also emerging as a powerful enabler. AI is now central to how we observe the Earth, explore distant planets, and manage our expanding presence in orbit. Its integration is helping transform space missions into faster, smarter, and more efficient endeavours, while also making them more relevant to daily life back on Earth.

At the heart of this change is data. Space is a complex, data-heavy environment. Earth observation satellites generate terabytes of imagery and sensor feeds every day, capturing everything from crop health and water body levels to forest cover, glacial retreat, and urban expansion. Without intelligent systems to interpret this information, much of its potential cannot be captured. All through machine learning and deep learning, is allowing us to process this data faster, draw sharper insights, and act in real-time.

Al's impact is already visible across key pillars of India's space activities.

- In Earth observation, AI-powered image recognition is being used to detect crop patterns, predict yield, identify mining activity, and flag deforestation. These applications are particularly useful for building climate resilience, enabling precision agriculture, and supporting long-term resource planning.
- In satellite operations, Al is improving how we manage large constellations of small satellites. It supports real-time health monitoring, bandwidth optimisation, and autonomous collision avoidance, helping keep space safe and sustainable. This is becoming increasingly important as India sees growing private participation in LEO (Low Earth Orbit) services.

Lt. Gen. Anil Kumar Bhatt (Retd.)
Director General

Indian Space Association

- By combining data from satellites and ground sensors, AI helps predict floods, forest fires, cyclones, and earthquakes with improved accuracy. These insights support emergency services and safeguard lives and livelihoods across rural and urban India.
- In planetary exploration, AI plays a critical role in enabling autonomous navigation. With long communication lags between Earth and distant celestial bodies, spacecraft need to make decisions independently. During the Indian Space Research Organisation (ISRO's) Chandrayaan-3 mission, AI-supported onboard systems helped the Vikram lander analyse terrain and execute a precise soft landing on the Moon's south pole. Future missions to Mars and beyond will increasingly rely on such intelligent autonomy.
- Alis also becoming essential in communication and navigation systems. It is being used to enhance spectrum efficiency, manage signal quality, and ensure reliability of Satcom, and Position, Navigation, and Timing (PNT) services which are key to sectors like mobility, defence, and infrastructure.

Al enabling Missions in Space

The National Aeronautical & Space Administration (NASA) uses neural networks to identify exoplanets, while the European Space Agency's (ESA) Gaia mission is cataloguing over a billion stars using Al-powered photometric analysis. In this evolving landscape, India has a unique opportunity to lead through cost-effective, scalable innovation. India's space economy is projected to grow to \$44 billion by 2033, according to the decadal vision released by IN-SPace (Indian National Space Promotion and Authorisation Centre).

Al will be central to realising this vision, not only by enhancing mission efficiency but also by creating downstream value in sectors like agriculture, climate science, defence, telecom, and urban development. Equally important is India's potential to serve other developing nations in the Global South, by offering Al-enabled space solutions that are both affordable and contextually relevant.

India's Al-in-space ecosystem is growing not only through government programmes but also through a strong wave of private innovation. Startups and enterprises are using AI to build solutions that are both technically advanced and commercially viable. Pixxel is building a constellation of hyperspectral imaging satellites, using Al analytics to generate detailed insights for environmental and agricultural monitoring. SatSure is integrating geospatial AI with financial analytics to provide decision support for agritech, banking, and insurance, helping assess crop loss, deliver targeted subsidies, and optimise rural credit risk. Skyroot Aerospace and Agnikul Cosmos are applying AI in launch vehicle design and propulsion system optimisation. Al is helping these startups improve performance predictions, simulate flight paths, and accelerate the development of reliable, small satellite launch platforms.

ISRO is also leveraging AI in launch planning, Earth data analysis, and robotics. Its Bhuvan platform

integrates AI for urban mapping and infrastructure monitoring. India's collaboration with NASA on the NISAR (NASA-ISRO Synthetic Aperture Radar) satellite will generate one of the largest AI-ready datasets on Earth's land and ice surfaces, vital for studying climate change and tectonic dynamics.

With regulatory support from IN-SPACe and industry coordination through the Indian Space Association (ISpA), India's private players are now able to access satellite data, ground infrastructure, and research facilities, further fuelling innovation at the intersection of Al and space.

Convergence of Al and space in future

As we embrace AI, we must also recognise the ethical and governance challenges that come with its autonomy and scale. Training datasets, particularly those used in surveillance and resource monitoring, must be representative and free of bias. Systems that make independent decisions in orbit must be accountable and transparent. India can take a leadership role in setting frameworks for responsible AI in space including clear guidelines on data privacy, system auditability, cybersecurity, and ethical design.

For India, the convergence of AI and space holds the key to unlocking new levels of capability, autonomy, and impact. With the right investments, policies, and partnerships, India can lead the world in building a responsible, scalable, and inclusive AI-powered space future.



The NASA-ISRO Synthetic Aperture Radar (NISAR).

AI IN FINANCIAL SERVICES

In the financial services sector, Al is evolving from a niche experimental tool to the very core of strategic growth and innovation. Globally, Aldriven platforms are revolutionising financial operations. JPMorgan's LOXM autonomously executes trades, while BlackRock's Aladdin provides real-time risk modeling across assets worth USD 10 trillion. In India, legacy institutions such as HDFC and ICICI have integrated Alpowered analytics into credit and investment products. Meanwhile, fintech innovators like BharatPe, Lendingkart, Fibe, Kissht, and Revfin leverage Al-based alternative data underwriting to extend credit access to underserved customers in Tier-2 and Tier-3 cities, democratising financial inclusion.

Al is shifting finance from reactive transaction handling to proactive life-stage orchestration, transforming client journeys into dynamic, anticipatory experiences. Financial institutions worldwide are leveraging Al insights to revolutionise credit evaluation, risk management, and investment decisions, delivering hyper-personalised products and proactive engagement. Indian digital lenders' use of alternative data exemplifies efforts to broaden financial inclusion, while global asset managers harness AI predictive models to optimise portfolio performance in volatile markets.

The emergence of agentic AI, which are autonomous systems capable of advising, executing, and continuously learning, signals the arrival of autonomous banking and deeply embedded finance across ecosystems. This shift necessitates comprehensive rearchitecting of operating models, investments in workforce reskilling, and embedding AI ethics and regulatory compliance as core strategic priorities rather than peripheral concerns.

Institutions embracing AI as a strategic co-pilot



Deepak Sharma
Board Member
Digital Leader,
Founder & CEO-Venture Studios

across innovation, governance, and customer trust are likely to dominate the next decade. While the adoption might take some time, the future of finance appears to be real-time, personalised, and increasingly Al-native.

Banking, Insurance, Payments, and Wealth Management

Banking & Lending: Fintech firms like Fibe (formerly EarlySalary) and Kissht provide instant loans via Al-powered digital underwriting with minimal documentation, collectively serving over 25 million users across 350 cities. Revfin's psychometric underwriting model, developed with IIT Kharagpur, is expanding underserved markets by financing electric vehicle purchases.

Insurance: Insurtech companies such as Ditto Insurance employ advisory-first AI to automate plan matching and claims guidance, enhancing customer experience through data-driven recommendations.

Payments & Infrastructure: Razorpay and Zeta utilise AI-based routing and reconciliation systems. Zeta's AI-enhanced cloud platform processes over 208 million credit accounts in 40 minutes, exemplifying the scale and speed AI brings to financial infrastructure.

Wealth & Capital Markets: Platforms like Groww democratise investing through robo-advisors; exchanges including NSE deploy Al-driven surveillance tools to maintain market integrity by detecting fraudulent activities in real-time.

AI, Productivity, and Reimagined Roles

Contrary to fears of job displacement, Al is expected to redefine roles within financial services, generating an estimated USD 1 trillion in annual value worldwide, according to a McKinsey report. A 2024 World Economic Forum report states that while automating approximately 40% of routine tasks, AI creates new career pathways including bias auditors, explainability specialists, financial coaches aided by algorithmic insights, and hybrid advisors blending human empathy with Al analytics. Bankers in the future will have to reflect the fusion of quantitative skills and human understanding in an Al-augmented financial landscape. ICICI's AI Centre of Excellence, for instance, is training initiatives fostering digital literacy and collaboration between humans and Al systems.

Tech-First Innovators

Technology-native firms are uniquely positioned to redefine finance. Agentic Al systems autonomously manage portfolios, negotiate insurance, and monitor transactions, evolving into true financial assistants acting on users' behalf. Completely autonomous banking workflows, covering everything from lending and KYC to fraud resolution are emerging in neobanks abroad and hold significant promise for India's smartphone users.

Al-powered vernacular interfaces and voice-chat enabled assistants like Niki.ai are bridging usability gaps in Bharat, enabling financial inclusion through intuitive interaction in native languages. Startups such as Neurofin, with backing from UNLEASH (USD 1.6M seed funding), develop GenAl-first platforms to automate compliance and operations within BFSI ecosystems. Generative Al chatbots reduce support costs dramatically; for instance,

brokers like Dhan automate KYC queries, while Zeta employs large language models for ultrascale account reconciliation. Further, Indian fintech unicorns PhonePe, Razorpay, Slice, Groww, PolicyBazaar, and BharatPe (which turned profitable in 2025), continue to push innovation frontiers in payments and lending.

Navigating Ethical Design, Risk, and Societal Impact

The adoption of AI in finance introduces profound ethical challenges demanding rigorous strategy. Algorithmic bias threatens equitable credit access, while data security risks mandate robust defenses against adversarial manipulation. Explainability and transparent governance have become regulatory imperatives worldwide, exemplified by guidelines from bodies like the RBI and the EU's AI Act. Emerging tools such as federated learning and continuous auditing facilitate responsible Al deployment. India's Digital Personal Data Protection Act (2023) and RBI regulations for digital lending establish a foundation for ethical accountability; ongoing oversight and clear Al governance remain essential to maintaining public trust.

The global AI market in BFSI (Banking, Financial Services and Insurance) is projected to reach USD 97 billion by 2030, with India's share expected at USD 4.2 billion (NASSCOM). Meanwhile, India's fintech market is set to hit USD 1.3 trillion by 2025 (BCG). Together, these figures highlight the rapid AI-driven transformation reshaping financial services in India and across the world. AI stands poised to move financial services beyond automation, reinvent customer experiences, streamline operations, and create new growth frontiers. India's financial ecosystem anchored by legacy banks, innovative fintechs, and a vibrant developer community is strategically positioned to lead this transformation.



AI IN QUANTUM

Al and Quantum Computing are two of the most transformative technologies of our time, and their intersection promises exponential breakthroughs. As Al applications become increasingly powerful and widespread, their demands on compute resources, energy, and efficiency are escalating. Simultaneously, quantum computing is maturing from theoretical curiosity to practical scientific tool. The convergence of these two domains is not just an academic possibility, it is a rapidly evolving area with significant implications for industry, science, and national capability.

India, with its deep talent pool and emerging quantum ecosystem, is particularly poised to benefit from integrating AI with quantum technologies. AI is already disrupting how quantum algorithms are written, how quantum systems are managed, and how scientific discoveries are accelerated. This dual impact of AI for quantum and quantum for AI, is opening entirely new avenues for national compeitiveness.

Integration of Quantum & Al

There are two sides of the equation when it comes to the integration of quantum and AI.

 One side is exploring how the potentially transformative power of quantum computing could accelerate the field of AI, such as new machine learning algorithms, re-imagined neural networks, and how we can interact with large

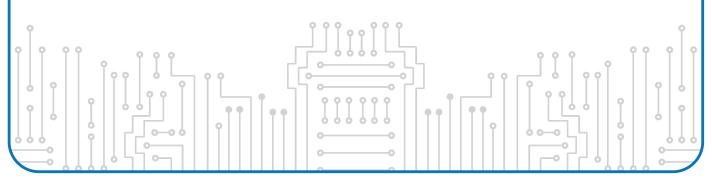




Kishore Balaji Executive Director Government & Regulatory Affairs, IBM India/South Asia

language models (LLMs). Today, current Al systems are still explored at a rapid pace, and their progress can be limited by factors such as energy consumption, long processing times, and high compute power demands. As quantum computing advances in scale and speed to open up new and previously inaccessible computational spaces, it could hold the potential to help Al process data at more practical and efficient speeds.

 The other important side of the equation, how Al can accelerate quantum, is one that can be advanced now. We can apply advanced Al models, such as generative Al, to accelerate the development of quantum code. In turn, this will catalyse the creation of quantum algorithms that industries are pushing forward now to discover quantum advantage - the point at which a quantum computer can solve a problem better than any classical method can.



Utilising the Quantum Advantage

Quantum systems are now at an important inflection point, which we refer to as utility-scale. This is the point at which they can serve as scientific tools to explore new classes of problems that are difficult or impossible for classical systems to access. To move forward into quantum advantage, our users across industries need to use generative Al tools, such as those within Qiskit Code Assistant, for easier creation and discovery of the algorithms that will push forward the complexity and relevance of quantum applications.

Quantum computing is a completely new branch of computing that is opening up new computational territory for industries. IBM has a very clear path to scaling quantum systems in a way which delivers real value to users. As they scale, they will have increasing capabilities to process incredibly complex datasets. This gives them natural potential to accelerate AI applications that require generating complex correlations in data, such as uncovering patterns that could reduce the training time of LLMs.

Pushing quantum computers to this point - where they can truly benefit AI acceleration - requires us to move forward the quantum circuits being run on systems today. This is why we are bringing generative AI tools to the more than 250 members of our IBM Quantum Network to do this as the experts in their fields. This network spans governments, industries, universities, and research institutions across the world-all of whom can use generative AI tools to improve and research new quantum algorithms that can bring us closer to quantum advantage.

As both AI and quantum advance in tandem, another aspect of their integration is bringing the strengths of both paradigms together for enterprises in a new architecture of high-performance computing. At IBM, this is what we describe as bits, neurons, and qubits, all working together as a tightly integrated system. Our strengths are in building this architecture for enterprises that will allow companies to move data between GPUs, CPUs, and QPUs seamlessly, which will in turn enable them to explore how they can apply different parts of a problem that are best suited for each paradigm.

The ongoing integration of AI and quantum computing is reshaping the technological landscape. Successfully & responsibly addressing these will enable AI and quantum technologies to drive sustainable national advancement and maintain global technological leadership.



IBM Quantum Network.

AI IN CYBER SECURITY

Cybersecurity is the core of digital resilience and is essential across traditional IT environments and critical sectors like finance, healthcare, infrastructure, public governance, and e-commerce. As threats become more automated, globally coordinated, and AI-powered, legacy defence models no longer suffice. As critical national systems become increasingly interconnected and attacks grow smarter, AI has emerged as a key defender of vital infrastructure. AI is disrupting cybersecurity through capabilities such as predictive threat detection, automated response, behavioral profiling, and identity protection.



- Behavioral & Anomaly Analytics: Al models analyse network communication, user patterns, endpoint actions, and system logs to establish baselines, instantly flagging deviations such as credential misuse or unusual lateral movement, reducing false positives and alert fatigue.
- Predictive Threat Intelligence & Attack Surface
 Management: All platforms continuously scan for
 vulnerabilities, dark-web changes, and
 configuration gaps, providing predictive risk
 assessments before attacks occur.
- Agentic AI & Automated Incident Response:
 Autonomously triages alerts, isolates
 compromised endpoints, and terminates malicious
 processes under pre-approved policies.
- Deepfake & Content-Genuine Verification:
 Forensic AI tools detect manipulated audio, video, or images in real time using metadata, heatmaps, and confidence scoring, supporting law enforcement, media integrity, and enterprise security.
- Cloud-Native & Container Security: ML systems analyse container behaviour and traffic to detect threats, isolate compromised microservices, and secure cloud deployments, even in production.



T.R. Reshmi SETS Chennai



M. Swathi Mithran SETS Chennai



A. Suganya SETS Chennai



Al Shields for Securing Critical Infrastructure in the Era of Advanced Threats

Power & Energy Grids: Al systems in leading utilities analyse real-time grid telemetry and network logs to establish behavioural baselines, detecting anomalies like unusual voltage spikes or control logic deviations before escalation. Predictive maintenance identifies equipment issues such as transformer degradation in advance, preventing downtime and reducing cyber-physical risk.

Industrial Manufacturing: Al-based anomaly engines detect subtle deviations in sensor data or control loops using explainable ML models, alerting operators with high confidence and minimal false positives.

Water Treatment & Waste Management: Al monitors chemical dosing, pump operations, and tank levels, with centralised ML platforms comparing real-time behaviour against expected patterns to detect threats from sensor spoofing to actuator tampering.

Healthcare & Health-IoT Networks: Al models at the edge monitor device and user behavior, detecting anomalies even in encrypted data streams to protect patient safety and confidentiality. Federated learning enables hospitals, including smaller facilities, to jointly develop models without sharing patient data, improving detection in compliance with privacy-preserving policies.

Physical & Drone Security: Al video analytics, thermal sensors, motion detection, and facial recognition integrate cyber, operational technology, and physical sensor data into unified defence layers to detect intrusions at scale.

India & Global Advancements in Al assisted Cybersecurity

In the U.S., the Department of Homeland Security issued AI safety guidelines for energy, water, and air travel infrastructure recommending human-centric value alignment and privacy protections. Globally, investments and innovation continue at pace like the recent reports estimate the global AI cyber market will grow from \$15 billion in 2021 to \$135 billion by 2030. Industry leaders emphasise open-source AI

security models and agentic Al systems for autonomous cyber defence. Google's Big Sleep and Sec-Gemini illustrate how AI is now detecting unknown vulnerabilities and accelerating global deployment of secure-forensics tools. Globally, financial firms like JPMorgan use AI analytics to prevent fraud. Vendors of platforms like CloudSEK and FireCompass have solutions that transform vulnerability management into a proactive process such as scanning continuously, surfacing emergent risks, and enabling organisations to address their weakest points before they are exploited. Similarly, security vendors like Microsoft, CrowdStrike are deploying agentic Al systems for automated incident response. Industries have solutions like Vastav AI for the content-Genuine Verification. Examples include Dragos, Darktrace for OT, and specialised tools like Jizo Al and ThreatDefence, which integrate with industrial control protocols to proactively hunt threats.

In India, Al adoption is rising across both public and private sectors. Behavioural analytics is getting adopted in many services which helps in identifying insider threats and anomalous access in government and enterprise systems. Also, Aadhaar and biometric authentication now integrate AI for identity verification. Similarly, UP's 'Al Pragya' initiative has trained over 1 million people including officials and educators in Al, ML, and cybersecurity skills by mid-2025. CERT-In has scaled Al-driven threat monitoring for critical infrastructure, while the Cyber Surakshit Bharat initiative boosted skills in Al-based cyber defence. State Bank of India also leverages AI to cut fraud by around 30% in real time. Similarly, Indrajaal Infra uses Al-driven drone detection and neutralisation systems deployed around ports, refineries, and nuclear plants. The system covers up to 4,000 km² using a combination of sensors, jammers, and intelligence tools.

Start-Ups and R&D initiatives

Several Indian startups are delivering cutting-edge Al-driven cybersecurity solutions. Lucideus (Safe Security) offers SAFE, an Al-powered risk quantification platform that translates cyber risk into financial terms for prioritised mitigation. Seconize provides a continuous Al-based risk and compliance engine automating discovery, predictive intelligence, and remediation aligned with ISO-27001, NIST, and RBI/SEBI standards.

CloudSEK's XVIgil delivers real-time threat intelligence and predictive analytics through surface, deep, and dark web monitoring, contextual alerts, brand abuse detection, and takedown services. FireCompass offers an Al-powered Continuous Automated Red Teaming (CART) platform that maps the full digital attack surface, including Shadow IT, simulates multi-stage attacks, and ranks risks for remediation. Zero Defend Security's Vastav Al, India's first cloud-based deepfake detection system, achieves up to 99% accuracy and is deployed by government and media agencies. SecureThings.ai's xSecureSquadron ensures vehicle cybersecurity with ML-driven threat intelligence, vulnerability scanning, in-vehicle protection, and secure OTA updates. Innefu Labs provides AuthShield, an Al-powered multi-factor authentication tool, and Prophecy, a predictive intelligence engine to mitigate risk. Sequretek offers an Al-enabled Endpoint Detection and Response (EDR) and threat monitoring platform.

Government & Skill Initiatives

India launched the IndiaAl Mission in March 2024 with a five-year plan to build a full-stack AI ecosystem. Skill-building programmes such as IndiaAl FutureSkills and iGOT-Al provide training in Al, ML, cybersecurity, and data analytics. The government has also funded AI Centres of Excellence in healthcare, agriculture, and sustainable cities. National missions like National Mission on Interdisciplinary Cyber Physical System (NM-ICPS) and National Supercomputing mission (NSM) foster AI-powered cyber defence through startup accelerators and sector-wide cybersecurity projects.

Harnessing Artificial Intelligence for Cybersecurity: Projects at SETS

As part of this strategic push under NSM, Society for Electronic Transactions and Security (SETS), in collaboration with five partner institutions including IIT Madras, IIT Jodhpur, IIT Delhi, and CDAC Bengaluru is involved in developing cutting-edge Al-driven cybersecurity technologies. The SETS initiative focuses on the following key areas:

- Multimodal Misinformation Detection and Network Analysis: Focuses on detecting fake news across modalities (text, memes, videos) by building curated datasets from factchecking and news sites. Executed by Scientist M. Swathi Mithran (SETS) with IIT Jodhpur, and to be hosted as a service under NSM.
- Framework for Al Applications (DARFA): A framework developed by CDAC Bengaluru with SETS contributions for robustness validation of Al models. Dr. Reshmi T.R. and M. Swathi Mithran (SETS) are applying it to Albased ransomware detection using defend-attack simulations to improve resilience.
- Secure and Private Machine Learning for Ransomware Forensics: Led by Dr. Reshmi T.R. (SETS), this project enables secure uploading and analysis of infected data by victims. Al models are developed to classify ransomware variants, support prevention/recovery, and integrate advanced cryptographic methods to preserve sensitive forensic data privacy.
- Physical Attacks on Al Systems & Al-enabled Defence Mechanisms: Investigates side-channel attacks on Al models implemented on FPGA/Microcontrollers. The study covers timing analysis, power consumption profiling, and EM emissions to demonstrate how leakages can reveal architecture details, parameters, or sensitive information. The project is led by Senior Scientist A. Suganya and Scientist Eswaridevi (SETS), in collaboration with IIT Madras.
- Malware Threat Detection and Defense in OT Environments: Funded by PGCoE-IISc Bengaluru, this project led by Dr. Reshmi T.R. and M. Swathi Mithran (SETS) leverages Alassisted behavioral threat intelligence for smart grid infrastructures, enhancing cybersecurity of critical OT systems.

Al in cybersecurity is emerging as a vital shield for safeguarding critical infrastructures. The adoption of Al in cybersecurity is accelerating across public and private sectors, enabling proactive defence against increasingly sophisticated threats. With ongoing investments and innovation, Al is set to become the foundation for securing digital ecosystems in the face of advanced cyber challenges.



HARNESSING AI IN HARDWARE PRODUCT DEVELOPMENT

Reshaping tasks across fields, AI has rapidly emerged as a prominent topic, prompting industries to rethink their processes and integrate AI into their workflows. However, hardware product development amidst this ongoing transformation presents challenges that require companies as well as AI developers to go beyond the current state of the art.

Developing a new product involves a series of iterative 'design-build-test' loops. The process typically involves loops for exploration & ideation, detailed design, manufacturing readiness and finally production. The role, effectiveness, and impact of AI are determined by the specific stage of development under consideration.

Al's use case in Hardware Product Development

Today, AI has shown clear value in the early stages of hardware product development, particularly during ideation and rapid prototyping. At this stage, the risks are relatively low, failures are quick, inexpensive, and anticipated, and requirements are less nuanced than in later phases. This "blank canvas" environment fosters rapid ideation and quick iterations.

efficiently sifting through scientific papers to pinpoint relevant findings and examine patents to identify design boundaries. Al-assisted deep search tools help identify suitable parts for given specifications and potential suppliers, aiding in initial cost estimations and vendor identification. Image generation tools contribute by transforming preliminary paper sketches into polished concept designs and schematics. Al can also optimise geometries for 3D printing, accelerating the prototyping of early concepts.

As ideas solidify, AI continues to add value in textbased processes, generating documents based on initial requirements.







These outputs provide helpful baselines but must still be checked and iterated upon by experienced engineers. Nevertheless, Al greatly accelerates the often expensive and time-consuming process of initial concept generation, prototyping and testing. Al is also starting to show promise in the later stages of product development, specifically during testing, validation, and optimisation.

Where AI Hits the Wall today

As hardware development progresses into detailed design and production phases, the limitations of AI become more apparent. Current Al tools cannot fully account for the practical challenges that engineers encounter, such as trade-offs between performance, reliability, manufacturability, and cost. Physical restrictions, supply chain quirks, and tacit knowledge accumulate in ways that machine learning models struggle to grasp. Any design or software deployed in the field requires meticulous validation to precede it. Unvetted Algenerated code or designs pose unacceptable risks, and companies are yet to organise their knowledge in ways that Al can use effectively to bridge the gap between concept and production readiness.

The limits of what is possible with AI is however being stretched every day. There are many efforts to use Generative AI to automate the core parts of product development such as system design, PCB design, CAD modeling, creation of assembly instructions, design for manufacturing and algorithm development. A future where AI develops a fully defined product with schematics and final blueprints based on text inputs is not outside the realm of the possible anymore.

Alat Ather Energy

At Ather Energy, we design, develop and manufacture smart, connected electric two-wheelers (E2W), E2W charging infrastructure, associated software, and accessories. We take a vertically integrated approach to product development, with in-house capabilities across battery packs, vehicle management systems, software integration, and power electronics. This combination of deep hardware engineering and continuous data feedback from connected vehicles positions us uniquely to leverage Al in hardware development; whether in optimising battery life through predictive models, enhancing vehicle performance via simulation-driven design, or accelerating iterative testing.

We build products with a focus on quality and user experience. This is achieved through rigorous product testing and validation, a process today enhanced by the use of Al-based tools. Al tools help define comprehensive test cases, covering a wide range of operating conditions, including edge cases. After testing, Al helps in analysing vast amounts of measurement data from numerous channels, identifying anomalies and errors that could otherwise lead to wasted effort, retests, or undetected faults. In addition, Al plays a role in optimising and designing the user interface for our applications and dashboards, a key component of the overall user experience. Al is increasingly playing a role in enhancing the vehicle service experience, another crucial aspect of user experience. It facilitates rapid identification of the root causes of customer issues and accelerates deployment of solutions via OTA updates.

These AI applications highlight how Indian hardware companies can contribute to the nation's 2047 vision by building indigenous, datadriven platforms where AI strengthens both innovation speed and reliability at scale.

The Future and opportunities for India

Al is integrating into the hardware product development process, particularly in its initial creative stages by accelerating research, ideation, and documentation. However, in the later, higher-maturity phases, the existing limitations of Al, specifically its current inability to handle complexity, nuance, and the need for hands-on experience and physical validation, pose significant hurdles. For Al to be more widely adopted, companies need to re-organise their data and build methods to validate it thereby building engineers' confidence in its outputs.

However, these very challenges provide a whitespace for Indian companies to utilise. Al presents an opportunity to innovate and enhance the software tools used in hardware product development, moving beyond the established practices of existing companies. Al has the potential to exponentially improve the pace of product development and reduce time to market. The rules of hardware product development are on the cusp of being re-written in the Al age.



Ather energy assembly station

HARNESSING AI TO BUILD A RESILIENT FUTURE FOR ALL

When a disaster occurs, resilient infrastructure ensures an effective response and quick recovery, thereby preserving the continuity of vital services. Imagine communities where Alpowered early warnings enable timely evacuation, and resilient bridges withstand flash floods. Such outcomes aren't futuristic; they are achievable with appropriate investment, inclusive technology, and global collaboration.

Across continents, there is a growing resolve to fortify public infrastructure against climate volatility and geophysical shocks. Resilient infrastructure is not merely a safety net; it's a foundation for sustainable development. That's also the core vision of Coalition for Disaster Resilient Infrastructure (CDRI), a global coalition launched by India in 2019 at the UN Climate Action Summit. Today, CDRI brings together over 50 countries and a dynamic network of experts, UN agencies, development banks, academics, and private players to strengthen infrastructure resilience against climate and disaster risks.

Addressing the Infrastructure Gap with Digital Technology

When integrated ethically and inclusively, AI is revolutionising how nations assess and respond to risks, turning reactive crisis management into proactive resilience.

AI and machine learning (ML) are already reshaping infrastructure planning in vulnerable regions. In South Asia and the Pacific, projects are yielding significant impacts. In India, ML models are assessing the seismic vulnerability of heritage structures. In Bangladesh, image classification algorithms support cyclone risk modellingand analysis of housing patterns. Nepal's mountainous terrain is being mapped with ML and satellite imagery to assess landslide vulnerability. The Maldives is using AI models to monitor sea level rise, while in Fiji, drone footage



Director General
Coalition for Disaster
Resilient Infrastructure (CDRI)

is swiftly processed to direct recovery efforts post-disaster.

These field-tested, scalable tools are driving real change by using AI/ML to reshape how we design and manage infrastructure under stress.

Yet, the promise of AI remains unevenly distributed. In Least Developed Countries (LDCs) and Small Island Developing States (SIDS), lack of digital infrastructure, data fragmentation, and limited institutional capacity often hinder the ability to fully harness AI.

To address these challenges, CDRI's 'Data and Technology Programme' is aimed at helping governments strengthen resilience by embedding AI & technology across infrastructure lifecycles, from design and regulation to operation and financing.

Some communities begin with basic dashboards mapping flood zones; others leap ahead with Alenhanced hazard simulations. But regardless of the technology tier, the objective remains constant - empower institutions to understand and act on risk.

Resilience isn't solely about technology; it's rooted in sound governance. CDRI aids member nations in building infrastructure asset registries, training officials in data analytics, and using platforms like the Global Infrastructure Risk Model and Resilience Index (GIRI). Increasingly, these efforts incorporate AI to enhance predictive modelling and early warning capabilities.

Towards Collaborative, Ethical, Inclusive Al Innovation

CDRI harnesses academic research, private sector innovation, and global experience to ensure member nations access the latest in geospatial intelligence, IoT, remote sensing, and immersive technologies like AR/VR. More importantly, these partnerships promote equity. Technology isn't delivered top-down; it's co-created. This ensures solutions are culturally informed, locally relevant, and tailored to diverse risk profiles.

For example, a CDRI Fellowship grant helped develop an AI simulator for the Dominican Republic to model climate threats, analyse impacts, and plan responses, while assisting communities and leaders in enhancing readiness, resilience, and sustainability.

Al's potential to revolutionise infrastructure must be matched by ethical responsibility. Al must be inclusive, transparent, and designed with community realities in mind because its purpose isn't just efficiency, it's equity. India, with its prowess in technology and diplomacy, is uniquely positioned to lead this shift by investing wisely, partnering inclusively, and governing responsibly.

A Resilient Tomorrow

Infrastructure should never be a barrier; it should be a bridge. Al, when harnessed thoughtfully, transforms risk into resilience. But technology alone won't solve our challenges. It requires collective will, visionary governance, and a shared commitment to safety and prosperity.

As the climate crisis intensifies, the resilience of infrastructure will no longer be judged by how well we build, but by how wisely we prepare for disaster management and risk reduction. With AI as our ally, we can shape systems that protect, empower, and adapt, not just for some, but for all, transforming today's vulnerabilities into tomorrow's strengths.



HARNESSING AI FOR INCLUSIVE AND ETHICAL EDUCATION IN INDIA:

UNESCO's Vision and Engagement

As India accelerates its digital transformation, AI stands at the heart of a profound societal shift, reshaping education, governance, healthcare, and even the ways in which citizens engage with institutions. Few technologies in recent history have carried such transformative potential, but it also comes with its fair share of risks.

At the United Nations Educational, Scientific and Cultural Organization (UNESCO), our mandate for over eight decades has been to build peace through education, science, culture, and communication. In this new era of technological disruption, that responsibility requires ensuring that innovation advances our shared values of inclusion, justice, and sustainability rather than undermining them.

Al's promise and the paradox

Al's potential to accelerate sustainable development is undeniable. From improving early disease detection in remote communities to supporting climate-resilient agriculture, from bridging learning gaps in classrooms to making public services more efficient, Al can serve as an equaliser. Yet, this very progress carries the risk of deepening divides. Unequal access to Al technologies, algorithmic bias, lack of transparency, and weak institutional safeguards could reinforce existing inequities rather than overcome them. The challenge is therefore a global one: how do we prepare AI for humanity, and societies for AI?

UNESCO's global ethical framework on Al

Recognising the urgency of this moment, UNESCO's 193 Member States unanimously adopted the Recommendation on the Ethics of Artificial Intelligence in 2021, the world's first global normative standard in this field. This historic agreement is grounded in human rights and sets out principles of fairness, accountability, inclusivity, and environmental sustainability in the design and deployment of Al.





UNESCO Regional Office for South Asia

But adoption on paper is only the first step. Putting values into practice requires action. To that end, UNESCO has developed the Readiness Assessment Methodology, now implemented in more than 70 countries, which helps governments identify gaps in policy, regulation, and institutions. In India, this is being carried out in partnership with the Ministry of Electronics and Information Technology (MeitY), ensuring that the country's AI ecosystem is not only innovative but also ethically robust. In June 2025, the National RAM Stakeholder Consultation in New Delhi, brought together over 200 actors from government, academia, civil society, and industry, underlining the shared urgency of integrating ethical guardrails into the nation's Al journey.

Education at the centre

Education remains UNESCO's founding mission, and its relevance is sharper than ever in the Al era. Recognising the transformative potential of Al in classrooms, UNESCO launched its Al Competency Frameworks for Teachers and Students in 2024. These frameworks provide a roadmap for countries to develop ethically oriented Al curricula.

For teachers, the framework sets out 15 competencies spanning ethics, pedagogy, and professional learning, with progressive levels of mastery. For students, it emphasises critical thinking, data ethics, and environmental sustainability, ensuring that the next generation are not passive consumers of technology but active, responsible shapers of it.

India, with its vast demographic dividend and digital public infrastructure, is uniquely placed to harness these tools. Al can support personalised learning, help bridge urban-rural

gaps, and enhance administrative efficiency in schools. But it can also reproduce social inequities if not approached with care. UNESCO's engagement seeks to ensure that India's education system treats AI not just as a tool, but as a subject of ethical reflection, fostering informed and empowered citizens.

Beyond classrooms: justice, climate, and inclusion

UNESCO's ethical vision for AI extends far beyond education. In the judicial system, for instance, the proliferation of AI tools raises pressing questions of fairness, transparency, and due process. Our Global Judges Initiative and Toolkit on AI and the Rule of Law equip legal professionals with the knowledge to interrogate algorithms that affect human rights. In South Asia, UNESCO co-hosted with National Law University Delhi a landmark international conference in 2024, bringing together judges and policymakers from South Asia Region, to examine AI's role in upholding the rule of law.

To build global momentum, UNESCO, in partnership with LG AI Research, is developing a Massive Open Online Course (MOOC) on the ethics of AI under its Empowering Minds: Transforming Lives initiative. The course equips technologists, students, and policymakers with practical tools to embed ethics-by-design in AI systems, fostering a worldwide community of responsible AI champions.

Similarly, UNESCO is advancing climate literacy and environmental education recognising that technological progress cannot be isolated from ecological responsibility. Initiatives such as the Greening Education Partnership, launched in 2022, seek to integrate climate consciousness into school systems, ensuring that the same young people learning to code AI systems are also learning to steward our planet responsibly.

A critical dimension of fairness lies in gender equity. Women remain significantly underrepresented in Al design and leadership, with just 30% of global Al professionals being female. UNESCO's Women for Ethical Al initiative underscores the risks this entails, gender-blind or gender-biased Al systems, and provides policy recommendations to promote diversity in STEM education, research, and governance. By empowering women and girls to

participate equally, India and the world can shape technologies that reflect the full spectrum of human experience.

India's opportunity and responsibility

These initiatives are not abstract ideals but urgent imperatives. UNESCO's 2022 State of the Education Report for India on Al in education laid out ten concrete recommendations, from making Al ethics a national priority to expanding Al literacy across all social groups. These remain as pressing today as they were three years ago.

India, already an acknowledged leader in digital public goods and committed to becoming a hub for Al innovation, has an opportunity to model a path where technological ambition is matched by ethical foresight. The lessons from the rollout of tools such as Aadhaar, UPI, and Digital Health IDs show that inclusion must be designed from the start, not retrofitted later. The country's ambitions in Al must likewise be grounded not only in technical capacity but in shared values.

Towards an inclusive Al future

The story of Al is yet to be written. It could be one of empowerment, democratised knowledge, and expanded opportunity, or one of exclusion, opacity, and widening divides. The path we take depends on whether human values remain at the heart of technological progress.

At UNESCO, we believe that peace begins in the minds of men and women, and that includes the way we design and deploy technologies. We stand ready to work with India, through policy dialogue, technical assistance, education frameworks, and global networks, to ensure that Al serves as a force for good.

By placing dignity, diversity, and sustainability at the core of AI development, India can not only shape its own future but also show the world how innovation can be inclusive, equitable, and serve the common good.



https://www.unesco.org/en/fieldoffice/newdelhi

AI IN E-GOVERNANCE: FORGING A VIKSIT BHARAT

Technology has long been a powerful foundation for a bridge connecting citizens and the government, enabling a more efficient, transparent, and citizen-centric governance framework. At the heart of this digital transformation is Artificial Intelligence (AI), a technology now augmenting this new era of egovernance. India's unique demographic and socio-economic landscape presents both challenges and unparalleled opportunities for AI, making its adoption in public services a priority for a Viksit Bharat.

The AI Resolution in Governance

The AI disruption has already begun, with the government's focus on leveraging AI to enhance the reach and effectiveness of its Digital Public Infrastructure (DPI) platforms like Aadhaar, UPI, and DigiLocker. Historically, governance has been a complex, resource-intensive process. Manual data collection and a one-size-fits-all approach often created barriers for citizens, and AI is now providing solutions to these deeprooted challenges.



Kavita Bhatia COO IndiaAl

For example, in India's judicial system, thousands of courts and multiple languages pose significant operational and administrative challenges. A case filed in a regional court would often be in a local language, requiring time-consuming and expensive manual translation to be understood at a higher judicial level, like the Supreme Court. This creates delays and a knowledge transfer gap, hindering the justice delivery system.

To address this issue, the Supreme Court of India recently announced the launch of SUVAS (Supreme Court Vidhik Anuvaad Software). It facilitates the instant translation of judgments, court orders, and legal documents from English to nine Indian languages and vice versa. This not only speeds up judicial processes but also empowers citizens by making legal proceedings accessible in their native tongues, ensuring a more equitable and efficient justice system.



Beyond the courtroom, the Digital India Bhashini Initiative builds on this principle to break language barriers more broadly. Bhashini aims to foster multilingual digital inclusion by developing Alpowered language models that allow citizens to access government services, information, and schemes in their native languages using voice commands and receiving voice-based responses. This is significant for a country like India with over 22 official languages and thousands of dialects spread across various geographies, ensuring that no citizen is left behind due to a language barrier. The seamless integration of Bhashini into various e-governance platforms demonstrates how Al is making governance more inclusive and accessible.

Al is also strengthening governance and public service delivery by improving efficiency, personalisation, and accessibility, transitioning from a reactionary to a proactive citizen-centric approach. It can analyse vast amounts of data to identify patterns and predict outcomes, enabling data-driven policymaking. This transition would make policies more effective and ensure they are tailored to the specific needs of different regions and demographics.

IndiaAl Mission: Road to Viksit Bharat

The IndiaAl Mission is also aligned with the vision of facilitating Al innovation to improve governance and public service delivery. In particular, the IndiaAl Application Development Initiative is a core pillar that focuses on developing and scaling Al solutions to address problem statements from various government ministries and departments. This initiative uses Al to solve India-specific challenges including critical sectors.

The IndiaAI Innovation Challenge, launched under this initiative, serves as a testament to this mission. It invited innovators, researchers, and startups to build impactful AI solutions. In its first round, governance was identified as a key thematic area for the challenge, alongside other areas like healthcare, agriculture, learning disabilities and climate change. The challenge specifically sought AI-driven solutions that could enhance public service accessibility and grievance redressal. One of the shortlisted solutions is an AI-driven transcription service designed to address India's overwhelming judicial delays. This legal tech solution automates court processes, including the collation of testimonies, orders, and judgments in multiple languages, and integrates with existing court systems to enhance efficiency.

As AI advances, the ethical aspects of its use in governance become paramount. The IndiaAI Mission, through its pillar on Safe & Trusted AI, is also committed to ensuring that AI is used responsibly. This involves developing ethical frameworks, promoting data privacy and security, and building explainable AI models to ensure transparency and accountability. We must guard against biases in AI systems and ensure they are fair and equitable, especially in decision-making processes that affect citizens' lives. The establishment of the IndiaAI Safety Institute underscores India's unwavering focus on ethical and safe AI development.

The impact of AI in e-governance is multi-faceted. It streamlines administrative processes and empowers citizens by making services accessible to all, regardless of language or location. Globally, India's experience can serve as a model for other nations, particularly those in the Global South, on how to leverage technology for inclusive and responsible governance. As AI advances, the ethical aspects of its use in governance become paramount. The IndiaAI mission, through its pillar on Safe & Trusted AI, is also committed to ensuring that AI is used responsibly. We must guard against biases in AI systems and ensure they are fair and equitable, especially in decision-making process that affects citizens lives.

Al as a Growth Engine

Al is no longer a futuristic enabler; it is the backbone of sectoral disruption, resilience, and inclusive growth in India and worldwide. The synergistic integration of Al across sectors such as agriculture, disaster management, life sciences, defence, and the digital economy is driving transformative change, enhancing farmer incomes, enabling smarter healthcare, and building resilient cities. Strategic investments in Al-driven sectors are expected to impact GDP, enhance global competitiveness, and redefine employment opportunities. Hence, national programs, government-industry-academia collaborations, and regulatory clarity become essential for Al scale-up. India is among the leading nations in developing frugal, scalable Al solutions and is increasingly shaping global Al policy and partnerships, positioning itself to harness the full potential of Artificial Intelligence





his article has been curated based on the conversation with Dr. Shivkumar Kalyanaraman, the Chief Executive Officer (CEO) of the Anusandhan National Research Foundation (ANRF) on Al's influence in the transitioning architecture of scientific inquiry, engineering practice, and education.



Early AI, programes relied on expert systems that held knowledge bases responding to questions with pre-codified answers. The advent of more sophisticated algorithms led to landmark moments such as IBM's 'Deep Blue' defeating chess grandmasters, based on the ability to understand and process natural language queries. The fundamental leap, though, came around 2010 with the rise of deep learning, a set of techniques involving neural networks capable of modeling complex, layered relationships in data. Unlike earlier linear models that essentially drew simple lines through data points, deep learning networks learn intricate, multi-layered "patterns" akin to human cognition. These networks can adaptively learn parameters at various levels of abstraction, which enhance Al's ability to process images, text, audio, and video.

A critical advancement occurred around 2017, with the introduction of transformer architectures, pioneered by Google. These models established powerful ways of pre-training on enormous datasets and then fine-tuning for specific tasks. This approach, called transfer learning, allows an AI trained to understand language broadly (for example on Wikipedia) to adapt to more specialised and diverse applications efficiently. Techniques like self-supervised learning, where AI predicts masked parts of data instead of relying solely on labeled examples, accelerated this growth.

This brought about tools like ChatGPT, AI systems capable of textual generation with increasingly sophisticated reasoning, education assistance, and directly modeling scientific concepts. "What is happening now is that you can perform more reasoning-like tasks; ultimately science itself can be more directly modeled. That is the threshold where we are today," explains Dr. Shivkumar.



Scan QR code to tune into the entire conversation



Democratising & Demystifying Scientific Research with Al

Al's transformative potential for research and education is also immense. Traditionally, acquiring knowledge meant reading books and internalising complex concepts. Today, Al can process that same knowledge, demystifying dense scientific papers, simplifying complex content, and delivering it in regional languages or personalised contexts. Projects like "Saral" by ANRF, exemplifies this mission of turning technical research into accessible material that engages a broader spectrum of learners and researchers.

Dr. Shivkumar states, "Democratising research and accelerating knowledge diffusion gives new researchers faster paths to discovery, confidence, and insight." With AI, educational and research materials, ranging from technical courses to extensive YouTube tutorials, are now at students' fingertips across the country. However, AI's true value lies in making learning more contextual, personalised, and actionable.

But access alone is not enough. "The real goal is not only to unlock potential, but also to help individuals discover it themselves," and to help them take the leap from passive learning to active contribution. The challenge, then, becomes to make learning relevant as well as applicable. "A problem is nothing but a knowledge statement and AI can actually expand those knowledge statements and add more context." By using structured "problem banks" sourced from live industry challenges and NGO experiences, India is fostering a pipeline where theoretical understanding aligns with real-world impact. Knowledge repositories, when enhanced by Al's contextual reasoning, deliver personalised pathways for learners' unique needs. Platforms like OPSA's Manthan and the forthcoming Anusandhan Community Platform play pivotal roles by fostering mentorship, talent collaboration, research internships, fellowships, and grant opportunities.

Facilitating Academic Research with Al Tools

Many commercial AI tools available today, including free versions, offer a "research mode" that is particularly relevant for the academic community. In AI, there are two main stages: training (where the model learns from large amounts of data) and inference (where the model answers questions or solves problems).

A recent development is allowing more compute time during inference. Instead of providing instant answers, AI systems in research mode can take longer to return more reasoned and comprehensive responses. This capability is very useful for tasks like conducting literature surveys or related work reviews for research papers. Users can interact, ask follow-up questions, and guide the AI to reach a meaningful conclusion. Human guidance is essential throughout the process, ensuring correct conclusions through judgment that cannot be substituted by any AI system, underlined Dr. Shivkumar. He further stated that these advancements help reduce mundane tasks for researchers, especially when access to specific resources may be limited and platforms like One Nation One Subscription (ONOS) can complement this by making publications more accessible.

Once the right papers are suggested, researchers can quickly find and read them. This integration of AI with discovery platforms streamlines the process of understanding, and makes research, iteration, and new findings faster and more efficient. Ultimately, AI accelerates any knowledge-intensive process, and there is a whole new set of emerging tools dedicated to this purpose. "The fundamental point is that the process of discovery itself is a knowledge-intensive process. Any knowledge-intensive process can be accelerated through AI-that is the basic idea," he added.

Models and Large Scale Simulations for Scientific Advancements

A crucial area for AI in science and engineering is the direct modeling of scientific equations and engineering systems. Instead of traditional approaches, where textual descriptions or equations like Newton's F=ma or the Navier-Stokes equations are interpreted and solved numerically, AI can now use neural networks to approximate not just functions but also operators. This breakthrough enables the creation of "neural operators". With neural operators, partial differential equations can be solved as a forward pass in a neural network, vastly speeding up complex simulations. For example, car body or wind turbine blade designs that once required extensive numerical modeling and high-performance computing can now utilise AI simulations that are ten thousand times faster. While there may be some approximation error, the gains in speed mean that many more iterations are possible, and some models can even be used in real-time.

This opens the path to embedding digital twins directly into products. A digital twin in a car, for instance, could enable real-time optimisation for performance or efficiency. Similar approaches can apply to wind turbines, engines, compressors, and other engineering systems. The vision is to build an open "India Stack" of science and engineering models, allowing plug-and-play access for innovators, startups, and researchers across domains. This creates a science and engineering stack, enabling direct modeling and rapid innovation.

Measuring scientific progress in the context of Al also demands a careful rethinking of key performance indicators (KPIs). At its heart, the goal is simple- to accelerate discovery, expand exploration, and achieve outcomes that are faster, better, and more affordable. Yet, as Dr. Shivkumar points out, "KPIs should capture both positive advances, like acceleration and impact, and potential negatives, such as bias or unintended consequences."

While Al's efficiency in generating reference reviews and preevaluating proposals is promising, there are clear cautions against removing human oversight. "Balancing automation with human insight is vital," he adds.

India is uniquely positioned to leapfrog in science and research through the adoption and accessibility of these Al tools. Just as the country moved rapidly from limited landline availability to widespread mobile phone usage, Al tools can allow Indian researchers to bypass traditional limitations and advance quickly. The opportunity exists to revolutionise how science is done, and ANRF aims to support India to take that leap and become a leader in this new paradigm. "Science is going to get revolutionised; how we do science is going to get revolutionised. The opportunity for India is to leapfrog. And the best way to leapfrog is to doit."





Al for Science and Engineering

ANRF is on the threshold of launching another mission, 'AI for Science and Engineering'. This initiative is designed as a national open-source program, the first of its kind in the world. "The plan is to encourage the development of new AI models and make datasets widely available," with collaboration from partners such as the IndiaAI Mission and key government departments including Department of Science & Technology (DST), Ministry of Earth Sciences (MoES), Department of Biotechnology (DBT), Department of Health research (DHR), Defence Research and Development Organisation (DRDO), and others. More partnerships with external agencies are also expected.

The mission's first goal is to build a stack of open-source AI models. It aims to extend into advanced domains, accelerating research in areas as diverse as astrophysics, biology, enzyme design, and cell and gene therapy. The next stage intends to deliver value to verticals by reimagining how processes like high-performance computing, CAD, and experimental design are conducted. This shift could enable startups, support online experimentation, and transform the entire loop of science. To achieve this, a variety of projects will be supported to push boundaries, positioning India to lead in science-driven innovation.

The mission also seeks to accelerate the journey for deep tech startups, which typically face long timelines, high capital needs, and extensive experimentation. By converting more workflows into computation, especially through AI, the cycle from deep science to deep tech can be significantly shortened.

From a deep tech perspective, the landscape offers multiple verticals for investment and growth, with AI as a major focus. ANRF aims to gradually liberalise and expand its support across these areas. "ANRF is liberalising progressively so that we can have academia, national research labs, startups, industry, and MSMEs come together."

India boasts one of the largest developer ecosystems, ranking among the top users of AI platforms, globally. Young users and builders from universities drive this momentum. In the national open-source program, model-building will be democratised through contests such as Kaggle-style challenges. Participants will be tasked with building models using specified data sets and ty pes, with prizes awarded to the best solutions. GPU resources will be provided, and collaborations will be facilitated with MIT and the IndiaAI program to ensure models are robust.

The resulting models will become public goods, freely available worldwide under broad open-source licensing terms. India aims to contribute substantially to the global pool of open AI resources, reciprocating the benefits received from international open source projects. If successful, these efforts will empower ordinary students to become future champions and redefine the application of science, both in India and around the world.

Academic Ethics in the Age of Al

As Al becomes increasingly integrated with research and digital publishing, questions of academic ethics grow more complex. Ethical guidelines and undertakings are required in the granting process itself, emphasising a shift towards quality over quantity. Reforms are underway in rankings and evaluations to promote meaningful outcomes and minimise negative practices.

Al has inherent limitations and biases, just as any tool does. It is crucial to highlight these concerns, opening new areas of study within academic ethics. To address this, ANRF recently announced a program for centers that combine social sciences and humanities with science and technology, aiming to support research into these emerging ethical challenges.

Dr. Shivkumar observes, "It is early days. Fundamentally, we aim to accelerate scientific discovery and open new directions by making research faster, better, and cheaper." Appropriate KPIs must reflect whether the process is truly being accelerated, and also monitor possible negative outcomes, such as bias or complacency. Unexpected effects often accompany new technology, making it hard to predict all measures upfront. "Let us continue to progress while keeping an eye on this as we go forward," he suggests.

Governance in mission mode programs must prioritise outcomes, but also balance quality, impact, and tangible results. Programs cannot use a "fund and forget" or abrupt withdrawal approach. Instead, a mentoring-oriented governance model is needed, focused on learning and adaptation. "What matters is not just failing fast or succeeding fast, but how many cycles you go through, how much learning occurs, and the adaptation rate," he explains. This is guiding ongoing reforms to make governance more responsive and supportive of meaningful progress.

Al is set to transform the foundations of scientific research and engineering. With advanced modeling, Al can accelerate the discovery process, enable machines and humans to reach new findings more rapidly, and uncover hidden insights from massive datasets in fields like astrophysics. Science is at a threshold where both the pace and scale of discoveries can be greatly enhanced.

Once new scientific knowledge is created, AI can help rapidly model and scale innovations, readying them for practical applications such as digital twins and real-world products. AI enables innovators and researchers to scan patent databases, identify opportunities, and evaluate product-market fit far more efficiently. Though certain elements still need development, the future possibilities are vast.

While there is ongoing dialogue about AI replacing jobs, Dr. Shivkumar emphasises to see AI not as a threat, but as a catalyst to empower scientists, students, and the broader population to advance more quickly and achieve what was previously out of reach. "How can we get empowered with AI to do new things, more things, and different things?"

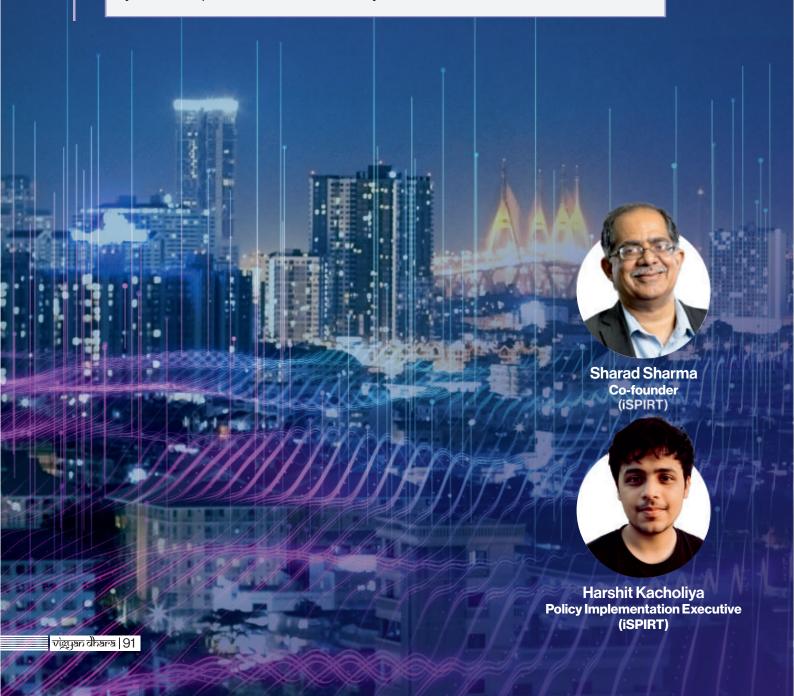
The transformative story of AI in research is one of empowerment, where science, technology, and innovation can reach new heights, helping a billion people and accelerating India's journey toward leadership in the global research ecosystem.



07

Digital Public Infrastructure (DPI) for AI

he era of building AI models on the vast, open frontier of the public internet is coming to an end. The foundation models that captivated the world, trained on the seemingly infinite expanse of scraped web data, represent the first and, ultimately, last generation of their kind. Leading research projects indicate that the available stock of high-quality public text will be exhausted by 2032. Compounding this scarcity is the degenerative risk of "model collapse," where AI trained on its own synthetic outputs loses touch with reality.



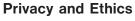
We are witnessing an evolution through three distinct generations of Al data. The first was the public internet, whose potential dangers have been outlined. The second, an intermediate step, is the era of aggregated, non-personal data portals. These are large-scale repositories, which are coming up across the world, including India - India's Al-Kosh and GDI (geo-spatial data-sharing interfaces). While they offer much better, more structured insights than the raw internet, their use is still limited as they do not provide insights that can be derived from private training datasets. But the true revolution, the third generation, will be built on a foundation that seems paradoxical: a world where the most valuable personal and proprietary data is aggregated and used for training, but in a privacy-preserving manner and without ever becoming public. India, by design and foresight, has built the infrastructure to leapfrog directly to this third generation and lead the world.

In a lot of ways, this transition is already underway. In finance, models like BloombergGPT, trained on private archives, are outperforming their public counterparts. In healthcare, institutions are fine-tuning models on de-identified patient records to create highly accurate diagnostic tools. In manufacturing, companies are leveraging proprietary sensor data to achieve remarkable efficiency gains. These examples reveal a clear pattern that when models are trained on rich training data, especially the one with personally identifiable information, it leads to superior, more reliable, and more valuable results.

The Challenge

While an AI trained on one hospital's records is good, a model trained on the combined records of one hundred hospitals, multiple pharmaceutical companies, and national health registries would be revolutionary. It could spot disease outbreaks in real-time, identify novel drug interactions, and create personalised treatments at a scale previously unimaginable.

This is the potential of the third generation of AI, but significant challenges pose constraints. Such rich, cross-sectoral data, often containing personal and proprietary details is unlikely to be made publicly available, due to important considerations. To avoid such critical risks, it is safeguarded by a host of concerns:



Patient records, financial transactions, and personal communications are protected by law and a fundamental right to privacy.





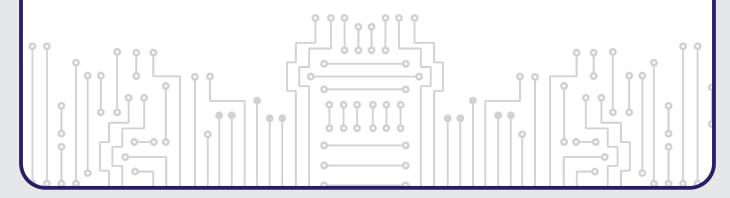
Business Secrecy and IP

Proprietary data on manufacturing processes, market strategies, and customer behavior is the lifeblood of competitive advantage.

National Sovereignty

The data on citizens, critical infrastructure, and economy are confidential national assets for the Governments.

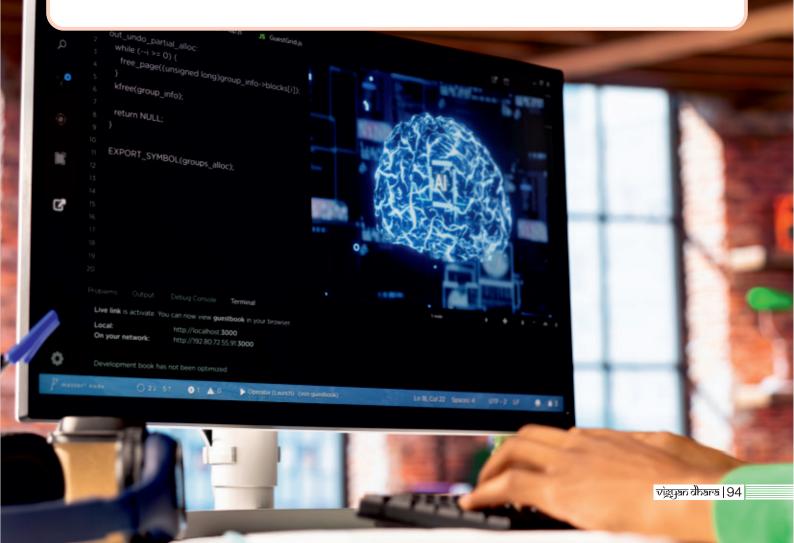






Initiatives like India's AI-Kosh or the India Urban Data Exchange (IUDX) provide significant value by making aggregated non-personal datasets accessible. The Geospatial Data-Sharing Interface (GDI), established under the National Geospatial Policy 2022 by the Department of Science and Technology (DST), brings together datasets from public and private partners across agriculture, livelihoods, transportation, and logistics. This sector-specific platform exemplifies how focused data-sharing models can help unlock government data effectively. These platforms mark an important advancement in data sharing. At the same time, some of the deepest insights in AI require sensitive, private data. By design, such data will remain confidential and unavailable for public model training, presenting an opportunity to explore innovative and complementary solutions.

The third generation requires a new plumbing, a DPI that provides a way to train models on aggregated personal and proprietary data from multiple, competing sources without any of those sources having to expose their data to each other or to a central repository, thereby addressing concerns related to privacy preservation, business secrecy, intellectual property, and national sovereignty. It must enforce access controls, ensuring that sensitive information flows only to authorised recipients under specified conditions. And it must provide audit trails, creating immutable records of how data is used without compromising the privacy of the underlying information.



DEPA for Training:

India's Leapfrog to the Future

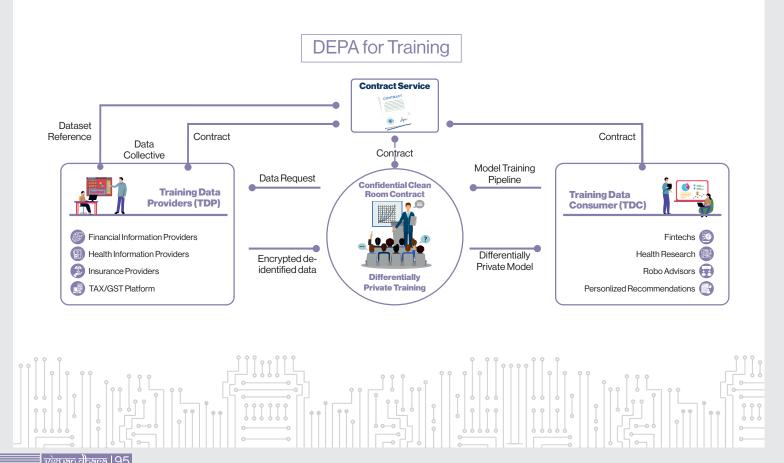
India's Data Empowerment and Protection Architecture (DEPA) addresses these challenges through an integrated approach that treats data sharing as a technical and legal problem simultaneously. Rather than building another platform or protocol, DEPA creates a comprehensive framework for privacy-preserving collaboration that works across sectors and borders.

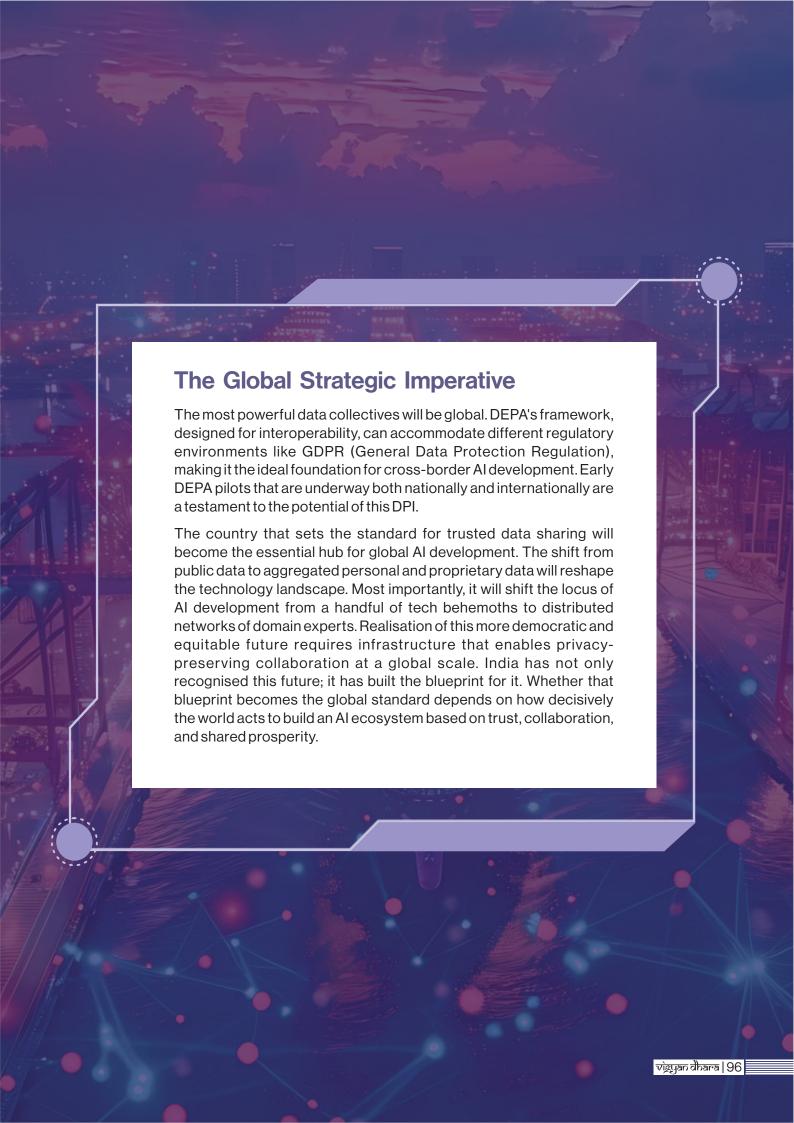
The system starts with data labels, rich descriptions that turn every dataset into a self-describing asset. These labels specify not just what the data contains, but how it was collected, what controls including privacy and monetisation constraints apply, and what uses are permitted. This separates annotations of data from the enforcement of these controls by unbundling them, DEPA also creates flexibility for innovation. Privacy-preserving techniques then enable the actual collaboration. Differential privacy adds mathematical noise that obscures individual records while preserving statistical patterns, and secure multi-party computation allows joint analysis without revealing private inputs.

All of this sits on top of techno-legal electronic contracts. When organisations want to collaborate, they negotiate terms that automatically generate secure computing environments where the actual data processing occurs. These environments enforce the agreed constraints technologically, making compliance automatic rather than optional.

Finally, Al-chains provide end-to-end audit trails that track every step from raw data to the final model. These immutable logs create accountability without compromising privacy, enabling organisations to prove compliance while maintaining competitive advantages.

The result is infrastructure that makes data collectives practical at scale. Organisations can contribute private information to shared training processes, receive insights derived from the combined dataset, and maintain complete control over their individual contributions throughout the process.





Al for Arts and Heritage

ach passing day brings with it more news about the rapid strides being made in Artificial Intelligence (AI) and its spread into every conceivable domain with deliberations raging about its effects on people, societies, countries and cultures.

The increasing penetration of AI into our society, and especially in the realms of art and heritage, signals a transformative era for global culture. Of course, there are many nuances and regulatory guidelines that need to be discussed, but AI technologies hold the potential not only to preserve, protect, and restore artistic and historical artefacts, but also to reimagine how culture is experienced, studied, shared and monetised. From digitising endangered manuscripts and reconstructing lost architectural marvels, to enabling immersive virtual exhibitions and unlocking new forms of creative expression, AI could act as both a custodian of the past and a catalyst for future innovation.





Sanjay Anandaram Co-founder Network of Indian Cultural Enterprises (NICEorg.in)

So what can Al do for culture, manifested through art and heritage, for the world?







Indus Script Source

Language is intrinsic to large language models (LLMs) used in Al. But language emerges from the substrate of the mind, which is a result of culture, values, experiences and embodies unique socio-cultural context, meaning and codes. A general-purpose LLM would, therefore, not be suitable for delivering meaningful outputs in a country like India, where over 30 languages are spoken by at least one million people each, and where unique Indic Knowledge Systems also need consideration. India offers a solution via its Bhashini initiative that aims to transcend the barriers of language. It leverages natural language processing (NLP) to build language datasets and multimodal Al models for 14 Indian languages. Translations of Indian cultural artefacts e.g. epigraphs, books, manuscripts, into English and other Indian languages, while preserving the socio-cultural context of the original, will enable the preservation of diversity. A recent example being the efforts of translating the Modi script into Devanagari using an Al framework for Visual Language modelling.



IIT Roorkee Modi script



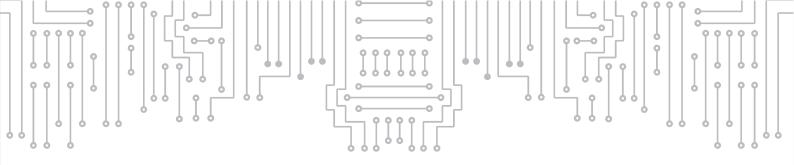
Indus Script Source

The creation of dictionaries. grammar resources and archives, will enable the creation of content. knowledge corpuses and dissemination of cultural insights across languages. NLP and RLHF (Reinforced Learning with Human Feedback) based Al offer great possibilities for decoding ancient scripts like the Harappan/Indus one, which has baffled scholars for decades. Deep neural network learning with custom databases and coupled with computer vision, image recognition and pattern matching are being proposed to determine the meanings encoded in the Indus script.



Rembrandt Nightwatch restoration

Heritage monuments, affected over time by factors such as environmental conditions, neglect, vandalism, or even deliberate damage, require thoughtful preservation and restoration to safeguard our shared cultural identity. Al and other digital tools play a significant role in this regard with 3D digital modelling, diffusion algorithms and accurate image reconstructions. Al can also be used for early detection of damage to art works, provide helpful guidance on the colours and patterns used in the original alongside reverse engineering painting techniques such as brush strokes of a particular painter for more realistic restoration works.





Rijksmuseum, Netherlands

Al, high resolution image and machine learning technology was used by the Rijkmuseum in the Netherlands, for example, for the restoration of Night Watch, one of the best known works of the Dutch painter Rembrandt, 370 years after it was cut down in size. In the case of damaged art work, collecting digital datasets, curating and labelling them with the help of domain experts and then using Al technologies is critical to ensuring fidelity to the original.



Louvre Museum, Paris

The Louvre Museum in Paris used AI to analyse, identify, and assist in the restoration of damaged sections of Leonardo da Vinci's The Virgin and Child with Saint Anne, a masterpiece over 500 years old. Though concerns were raised around the extent of cleaning, the use of AI marked a significant step in blending technology with art conservation.

Al played an important role in the restoration of the The Notre Dame cathedral in Paris that was destroyed in a 2019 fire thanks to an accurate digital depiction of the building.



Notre Dame Cathedral, Paris

In addition, AI is beginning to also play a role in addressing heritage theft by identifying stolen antiquities. For example, the Khmer Statuary Project uses AI trained on images of Cambodian art from multiple sources to detect bogus artifacts. Indian Institute of Technology -Madras is partnering with the Idol Wing of the CID, TN Police to use AI software to uniquely identify idols belonging to the Sangam, Chola and Pandya periods.

Usage of AI along with immutable ledgers and IoT is also building trust and credibility in the authenticity of traditional arts and crafts across the supply chain. Companies like Kosha.ai are leading this effort. By integrating generative AI with traditional crafts such as block printing, which are witnessing a decline, it is possible to introduce innovative designs and color palettes while preserving their essence.

Art and heritage are audio-visual immersive experiences. Al can provide such rich 3D experiences through mash-ups with Augmented and Virtual Reality, enabled by wearable devices. Museums around the world are at the forefront of such experiences with personalisation, interactivity and rich experiences especially when coupled with multimodal and local languages. Museums like the British Museum, Louvre and the Smithsonian are using Al driven apps for enhancing user experiences. The residence of the President of India, Rashtrapati Bhavan, has the Navachar Al gallery while the National Museum's miniature paintings section provides enhanced visitor experiences using machine learning based AR and robotics in its technology aided gallery. The private Museum of Art and Photography in Bengaluru uses Al and VR/AR to create interactive digital personas of celebrated painters using speech synthesis and deep learning. This opens up an entire new set of experiences that appeals to an entire digital-native generation of museum goers.



Digitally experiencing cultural heritage across the country and world, for example, via "tours" provide virtual education, entertainment, emotional participation that, in turn, drives physical experiences through a virtuous cycle. There are multiple emerging platforms that enable such experiences and adopting and enhancing these with citizen participation will make for richer and personalised experiences. Developing Al-enabled, district-wise digital maps of the country that highlight heritage points of interest such as arts and crafts, cuisine, monuments, and other cultural landmarks enriched with crowd-sourced content, interactive features, and 3D experiences, holds transformative potential for both the tourism sector and local communities.

All of these require not just the unlocking and availability of data across different sources but also context, source and meanings that only metadata can provide. All enables intelligent cataloguing and pattern matching to identify trends and relationships across different datasets. This enables faster and more accurate discoverability and searchability including semantic search, especially across different text, image, speech and video data.

But there are Challenges

The availability of accurate, comprehensive datasets and commonly accepted frameworks such as those for metadata and storage, across different data custodians and geographies complicates matters. Without such standards, establishing trusted data spaces to use, share, and collaboratively develop cultural heritage assets becomes difficult, particularly in light of privacy, security, and regulatory concerns. Ensuring data interoperability, supported by an identification chain, is therefore essential. In this context, Aldriven data mining can play a vital role in analysing visitor traffic, pricing trends, outreach channels, cultural consumption patterns, and exhibit preferences to better understand and engage diverse audience segments.



The EU's Europeana programme attempts to create a common data space for EU countries. The need for enhanced digital infrastructure, training, and funding to achieve the goals laid out in its strategy cannot be overstated. Creating awareness about the uses, challenges and governance of cultural artefacts - libraries, museums, galleries, private collections, public domain and so on - becomes a challenge.

Deep fakes are also a serious problem too. How does one determine the authenticity of a work of art or craft as a work of an actual artist or crafts person and not mass produced by an AI workshop somewhere? Should works of art be watermarked or should they be tokenised and put on an immutable ledger? How does one deal with the challenges of copyright? Identifying copyrighted material or providing for data management of such material or managing rights becomes another important aspect to look into. Who is liable if an AI model hallucinates and this is discovered after a major artefact is repaired or restored?

These are crucial questions that need deliberation and a mindful implementation, keeping all sensitivities in mind as AI evolves further. The possibilities of AI in art and heritage are immense and can revolutionise the cultural sector in India. Imagining this new world requires new institutional architectures with the public and private sectors, technologists and domain experts from different cultural fields, coming together to unlock data, funding, talent and awareness to make the new world of AI mesh with the traditional.



Building Al for a Sustainable Tomorrow The revolutionary impact of Artificial Intelligence (AI) is no longer a distant prospect. Al-driven innovations have redefined traditional boundaries and are accelerating societal progress in unprecedented ways. However, this potential is deeply intertwined with the need for advanced compute

infrastructure, vast datasets, and high-speed network elements that are increasingly energy-intensive and environmentally consequential. As AI systems become more powerful, their energy demands and material dependencies intensify, raising urgent

questions about sustainability and ecological resilience.

Dr. Divya Vashisht, OSD, Office of PSA र्णेड्रपुंबत वीतवन्व | 105

Al's Ecological Paradox:

Al holds transformative potential in addressing climate challenges. Its advanced modeling techniques enable high-resolution forecasting of glacial melt, iceberg trajectories, drought progression, carbon fluxes and facilitating more timely and targeted interventions. For instance, Al-driven systems analysing polar ice dynamics have demonstrated the ability to compute meltwater discharge rates up to 10,000 times faster than conventional manual approaches, substantially improving the precision and speed of climate impact assessments. However, this promise is accompanied by a profound paradox. At the heart of Al's environmental footprint lies its digital infrastructure. Data centers, which underpin the training, deployment, and operation of AI models, are among the most resource-intensive facilities globally. In 2022, data centers consumed approximately 460 terawatt-hours of electricity. [11] Training frontier AI models such as GPT-4 entails exceptionally high computational and energy demands. The model was developed using approximately 25,000 NVIDIA A100 GPUs, each consuming around 6.5 kilowatts of power. [iii] Over a training period spanning several weeks, this infrastructure resulted in an estimated energy consumption of 50 gigawatt-hours (GWh). The associated carbon emissions from this process are projected to range between 12,456 and 14,994 metric tons of CO₂ comparable to the annual emissions of over 1,500 average U.S. households. McKinsey estimates that global data center capacity could nearly triple by 2030, reaching up to 219 gigawatts (GW) with approximately 70% of this growth driven by Al workloads. This surge is attributed to the rapid expansion of generative AI and advanced machine learning applications, which require significantly higher computational density and energy throughput than traditional IT systems. [iv] These figures underscore the critical need for carbon-aware scheduling, energy-efficient model architectures, and sustainable infrastructure in the development of largescale Al systems.

Beyond electricity consumption, data centers produce substantial thermal loads, necessitating sophisticated cooling systems to maintain operational integrity particularly for high-density AI workloads. Water-based cooling has emerged as a dominant solution, yet it introduces significant environmental trade-offs. In 2023, U.S. data centers consumed approximately 17 billion gallons of water, with forecasts suggesting this figure could quadruple by 2028 due to escalating AI-driven demand. As the compute intensive infrastructure continues to scale-up, it will drive a proportional increase in water demand for cooling power systems. These trends raise acute sustainability concerns, especially in water-stressed regions, where competition between agricultural, industrial, and digital sectors may intensify. AI technologies could contribute up to 3.4% of global carbon emissions, a figure that has surged more than elevenfold over the past decade. [vi]

Al's Ecological Footprint Spans both Software and Hardware Domains.

On the software side, energy consumption spans the entire AI lifecycle from data acquisition and preprocessing to model training, inference, and deployment with training historically being the most energy-intensive phase, especially for large-scale architectures like GPT-4 and BLOOM that demand tens of gigawatt-hours of electricity during development.

Hardware-related implications are equally significant in shaping the ecological footprint of AI systems. These technologies rely on high-performance components such as GPUs and specialised accelerators, which are manufactured using critical minerals including rare-earth elements. To give a reference, constructing a single 2-kg computer can require up to 800 kg of raw materials, underscoring the resource intensity of AI infrastructure. Moreover, the accelerated obsolescence cycle often every two to three years driven by rapid advancements in chip design and performance benchmarks, contributes to mounting electronic waste. Recent projections estimate that AI-related e-waste could reach between 1.2 and 5 million metric tons by 2030, depending on hardware demand trajectories and upgrade cycles. Notably, generative AI alone is expected to account for 2.6 million metric tons annually by 2030, highlighting the urgent need for sustainable hardware lifecycles and circular design strategies. [ix]

This infrastructure race is not merely technological—it is geopolitical. Nations that fail to scale Al infrastructure risk erosion of digital sovereignty, diminished global competitiveness, and reduced strategic influence.



India's Dual Mandate for AI and Sustainability

India stands at a pivotal juncture. As it advances digital sovereignty through IndiaAI mission, it must simultaneously uphold its sustainability commitments, including net-zero carbon emissions by 2070. Data center energy demand in India is projected to increase tenfold, reaching 9.2 GW by 2030. While this growth is essential for AI competitiveness, it presents complex challenges in balancing energy security, environmental stewardship, and strategic autonomy.

India's AI strategy, anchored in inclusive growth (#AlforAII) has the responsibility of balancing the environmental costs with scaling infrastructure particularly in water-stressed regions and coal-dependent grids. While the country's e-waste recycling rate has improved from 22% in 2019–20 to 43% in 2023–24, yet a significant 57% remains unprocessed^[x], highlighting the need for robust circular economy mechanisms. The Union Budget 2025 further reflects this dual commitment to technological advancement and ecological responsibility, with increased allocations for Geographic Information Systems, smart urban planning, and climate-resilient infrastructure. India's strategic posture thus positions AI not only as a driver of innovation and sovereignty, but as a lever for sustainable development. [xi]

The Twin Pillars of Sustainable Al

The imperative for environmentally responsible AI has catalysed two complementary frameworks: Green AI and Sustainable AI.

Green AI prioritises computational efficiency by employing techniques such as model pruning, quantisation, knowledge distillation, and low-rank factorisation. These approaches streamline neural networks by reducing the number of parameters and operations, thereby lowering memory requirements and power consumption. Empirical studies show that such optimisations can cut energy usage by up to 50%, while maintaining over 95% of the original model's accuracy especially in high-demand applications like natural language processing and computer vision. [XIII] This balance of performance and efficiency makes Green AI a critical strategy for sustainable deployment at scale.

Sustainable AI takes a comprehensive approach by embedding environmental, ethical, and social principles across every stage of the AI lifecycle. It emphasises responsible sourcing of materials for hardware, the use of energy-efficient architectures, equitable access to AI technologies, and robust ethical governance. This framework ensures that AI systems are not only technically advanced but also aligned with long-term societal well-being and planetary sustainability.

Strategic interventions can contribute to reshaping the sustainability landscape of AI:

- Building carbon neutral infrastructure: Transitioning to solar, wind, and hydroelectric power can make
 Al development more sustainable. Additionally, Small Modular Reactors (SMRs) offer a reliable, lowcarbon alternative capable of supporting the continuous energy demands of Al infrastructure.
 Integrating SMRs alongside renewables not only enables scalable Al deployment but also reinforces
 national objectives around environmental leadership and technology sovereignty.
- Strategic placement of data centers: Building data centers and high-performance computing hubs in proximity to renewable energy sources or naturally cooler coastal regions can substantially lower cooling requirements and associated carbon emissions. This strategic siting enhances operational efficiency and has been widely adopted by leading technology companies such as Google and Amazon as part of their long-term sustainability commitments.
- Edge Al processing: By enabling data processing directly on devices such as sensors, smartphones, and embedded systems, Edge Al reduces dependence on centralised data centers for small-scale tasks and significantly reduces energy consumption for routine inference operations.

- Circular approach to equipment uses: Strategic reuse of AI hardware helps balance cost, performance, and energy demands, extending the lifecycle of components while reducing environmental impact. Manufacturing a single high-performance GPU can emit up to 600 kg of CO₂ before it is even deployed, making recycling a practical solution to curb hardware waste. [XIIII] Although older GPUs may not match the efficiency of newer models, they remain effective for lightweight training and data preprocessing tasks. Manufacturers and operators can support a circular AI economy by disclosing supply chain emissions, extending equipment warranties, and designing systems for modular upgrades. These measures are essential to reduce end-of-life waste and promote responsible resource stewardship.
- Resource sharing: Shared data center infrastructure enables economies of scale, optimised hardware utilisation, and streamlined cooling and power systems—resulting in lower Power Usage Effectiveness (PUE) and a reduced environmental footprint. Recognising this imperative, the IndiaAI Mission has adopted a public-private partnership model to establish a national AI compute facility comprising over 18,000 GPUs, housed within shared data centers. [xiv] These facilities are designed to meet stringent energy efficiency benchmarks, including PUE targets below 1.35, aligning with global best practices. [xv] This approach not only democratises access to high-performance computing for startups, academia, and public institutions but also supports India's broader climate commitments, including its net-zero target by 2070.



Regulatory Framework for Sustainable Al Deployment

The United Nations Environment Programme (UNEP) has taken a pioneering stance by advocating for end-to-end lifecycle assessments (LCAs) of Al systems. In its 2024 report, UNEP emphasises that Al's environmental impact must be evaluated for raw material extraction, hardware manufacturing, model training, deployment, and eventual decommissioning. [xvi] The report calls for:

- Science-based methodologies to quantify Al's carbon, water, and material footprints.
- Standardised metrics for energy consumption and emissions across AI workloads.
- Integration of LCAs into national digital strategies to align innovation with ecological resilience.

India has made significant strides in embedding sustainability into corporate governance through the Business Responsibility and Sustainability Reporting (BRSR) framework, mandated by the Securities and Exchange Board of India (SEBI)^[xvii]. As of FY 2023–24, the top 1,000 listed companies are required to disclose:

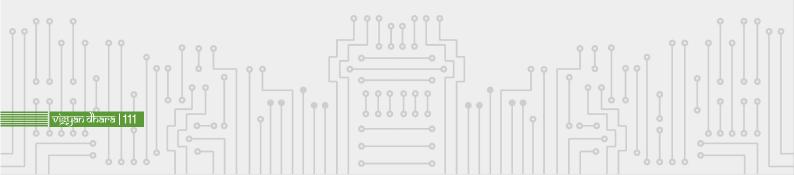
- **Scope 1 emissions:** Direct GHG outputs from owned or controlled sources (e.g., diesel generators, HVAC systems in data centers).
- Scope 2 emissions: Indirect emissions from purchased electricity, heat, or steam particularly relevant for AI workloads hosted on energy-intensive cloud or on-premise infrastructure. [xviii]
- Water and energy consumption: Including metrics on cooling systems, which are critical for high-density Al compute environments.
- Waste management and circularity: Addressing e-waste from GPUs, servers, and other Al hardware.

The BRSR Core format enhances transparency by requiring intensity-based disclosures and alignment with global standards such as GRI (Global Reporting Initiative).

Global Governance for Sustainable Al

- The European Union's AI Act introduces substantive measures to advance the sustainability and environmental accountability of AI systems. Among its key provisions is the mandatory reporting of energy consumption thereby embedding ecological transparency into the regulatory framework. Furthermore, the associated standardisation request mandates the development of deliverables focused on documentation and reporting protocols aimed at enhancing resource efficiency. This includes lifecycle-based assessments to reduce energy and material consumption in high-risk systems, as well as promoting energy-efficient design principles in the development of general-purpose models. [xix]
- The AI Action Summit 2025 culminated in a landmark Joint Declaration on Inclusive and Sustainable AI for People and the Planet, endorsed by 64 nations and organisations. The declaration prioritised environmental sustainability, acknowledging the rising energy demands of large-scale AI models and data centers. Discussion on global knowledge exchange to drive investment in energy-efficient AI systems spanning hardware, infrastructure, and model while advancing international dialogue on AI's environmental impact. Key outcomes included the launch of an AI energy observatory in partnership with the International Energy Agency and the showcasing of energy-friendly AI innovations. [xx]
- The OECD's (Organisation for Economic Co-operation and Development) Recommendation on AI, adopted in 2019 and updated in 2024, serves as the leading intergovernmental framework for trustworthy and sustainable AI. Though non-binding, it guides policy development and regulatory alignment across 47 countries, fostering global interoperability. Its principles centered on human rights, transparency, and sustainability are widely adopted by the EU, UN, Council of Europe, and the United States, including its standardised definition of AI systems and lifecycle. The Recommendation promotes AI design aligned with the UN Sustainable Development Goals and encourages international cooperation to share best practices and regulatory tools. [xxii] [xxiii]
- Established in 2020, the Global Partnership on Artificial Intelligence (GPAI) is a multi-stakeholder initiative of 44 member countries advancing human-centric AI through collaboration among governments, academia, industry, and civil society. In July 2024, GPAI integrated its governance with the OECD AI Policy Observatory, aligning with the OECD Recommendation on AI to promote trustworthy and sustainable AI. [XXXIII]
- India, Brazil, and South Africa as part of the G20 Troika issued a Joint Declaration on Digital Public Infrastructure, AI, and Data for Governance, endorsed by several G20 and guest countries. This declaration highlights how well-designed DPI augmented by AI can accelerate progress toward the SDGs, especially in the Global South. It emphasises inclusive digital transformation, citizen-centric systems, and equitable technological deployment. [xxiv]

With the vision for responsible innovation, India's AI trajectory must be both visionary and responsible. As the nation scales its digital capabilities, it must lead in integrating sustainability into the core of its AI ecosystem. The future of AI must be not only intelligent and inclusive but also resilient, equitable, and ecologically sound.



References:

- [I] https://www.weforum.org/stories/2024/02/ai-combat-climate-change/
- [ii] https://www.datacenterdynamics.com/en/news/global-data-center-electricity-use-to-double-by-2026-report/
- [iii] https://www.genspark.ai/spark/how-many-gpus-are-needed-to-train-gpt4-o/56c3b449-f206-495f-bc9b-d06812bf4a33
- [iv] https://www.mckinsey.com/industries/technology-media-and-telecommunications/our-insights/ai-power-expanding-data-centercapacity-to meet-growing-demand
- [v] https://theconversation.com/data-centers-consume-massive-amounts-of-water-companies-rarely-tell-the-public-exactly-how-much-262901
- [vi] https://www.accenture.com/content/dam/accenture/final/corporate/corporate-initiatives/sustainability/document/Powering-Sustainable-Al-CanvasV10.pdf
- [vii] https://unric.org/en/artificial-intelligence-how-much-energy-does-ai-use/
- [viii] https://www.nature.com/articles/s43588-024-00712-6?utm_campaign=The%20Batch&utm_medium=email&_hsenc=p2ANqtz _PBjy3BA0I JFWivM-bVCodsk10mWG1qbFV6AC_JkHAIVRNjB7V-aQN7TGWbQ_5AP06SWsaeU7I7N98UMOvM2APT6n2hB7pUZ_ z1xb34wPE_ B2V5B0&_ hsmi=2&utm_content=2&utm_source=hs_email
- [ix] https://spectrum.ieee.org/e-waste
- [x] https://www.downtoearth.org.in/waste/indias-e-waste-surges-by-73-in-5-years
- [xi] https://government.economictimes.indiatimes.com/news/economy/union-budget-2025-industry-pins-high-hopes-on-ai-tech-sustainability-to propel-indias-digital-future/117391589
- [xii] https://link.springer.com/article/10.1007/s43621-024-00641-4
- [xiii] https://www.analyticsinsight.net/tech-news/extending-ais-power-without-the-waste-rethinking-gpu-use-for-a-greener-future
- [xiv] https://indiaai.gov.in/article/union-minister-of-electronics-it-railways-and-i-b-announces-the-availability-of-18-000-affordable-ai-compute-units
- [xv] https://energy.economictimes.indiatimes.com/news/renewable/indias-sustainability-push-for-data-centers-a-progressive-move-towards-decarbonisation/114307870
- [xvi] https://wedocs.unep.org/handle/20.500.11822/46288;jsessionid=99C93C9FC24F1FBD5594EC602AE5571F
- [xvii] https://www.arbor.eco/blog/business-responsibility-and-sustainability-reporting-brsr
- [xviii] https://www.sebi.gov.in/sebi_data/commondocs/may-2021/ Business%2 Oresponsibility%2 Oand%20sustainability %20reporting%20by%2 Olisted%20entitiesAnnexure2_p.PDF
- [xix] https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng
- [xx] https://www.elysee.fr/admin/upload/default/0001/17/0efc53b2025d5bc1f0fc68556c41370cee166e7f.pdf
- [xxi] https://www.ailawandpolicy.com/2024/05/oecd-updates-guidance-on-responsible-ai/
- [xxii] https://oecd.ai/en/ai-principles
- [xxiii] https://oecd.ai/en/
- [xxiv] https://www.mea.gov.in/bilateral-documents.htm?dtl/38551/Declaration+on+Digital+ Public+Infrastructure+Al+and+Data+for +
 Governance++Joint+Communiqu233+by+the+G20+Troika+India+Brazil+and+South+Africa+endorsed+by+several+G20+countries+guest
 +countries+and+international+organisations





PEOPLE · PLANET · PROGRESS

Date: 19-20 February, 2026

Bharat Mandapam

Pragati Maidan, New Delhi – 110001, India.



Scan for Details on the Al Impact Summit 2026



Amit A. Shukla Joint Secretary, Cyber Diplomacy Division, Ministry of External Affairs, India that the transformative power of artificial intelligence is not only safe, ethical, and accessible but is also attuned to the diverse realities of our world. By prioritising developmental perspectives and the unique aspirations of the Global South, the Impact Summit aims to shape the global AI discourse into one that is inclusive and human-centric, ensuring that voices and priorities from all regions, especially the Global South, are meaningfully represented. As a responsible stakeholder in the international community India seeks to amplify the voices and needs of developing nations, foster international collaboration, and work towards a future where AI promotes equitable growth and collective progress for all.

As the AI technological revolution grows, it offers an opportunity to leapfrog traditional developmental pathways, especially for developing countries. Al is seen as a strategic tool to enable inclusive growth and expand access to opportunities. The past four years have witnessed a remarkable surge in multilateral initiatives dedicated to harnessing AI's potential responsibly and equitably. From the G20 AI Principles to the UN and Global Partnership on AI (GPAI) Resolutions on AI, the African Declaration on AI, and most recently, the Hamburg Declaration on Responsible AI, these efforts reflect international recognition regarding AI's transformative impact that transcends borders and requires coordinated global action. Despite these high-level commitments, a fundamental gap persists between global aspirations and the reality on the ground; the Global AI Divide continues to widen. Simultaneously, AI's rapid proliferation across society is creating new challenges that demand urgent attention, underscoring the need to move beyond aspirational frameworks toward concrete, measurable impact that addresses both AI's promise and its perils.

THE IMPACT OF AI

The world stands at a pivotal moment. Al is reshaping societies, economies, and ecosystems at an unprecedented scale, offering transformative potential for inclusive growth and sustainable development. The choice before us is clear - we can allow Al to widen existing divides, or we can work together to ensure it becomes a force for equity and shared prosperity. Achieving this vision requires global consensus on how Al is developed and applied. This includes rethinking how data is managed, how Al models are designed and deployed, and how these technologies are integrated into our social and economic fabric.

THE MANTRA

A Time for Impact

In 2026, leaders of the world will gather in India for the first Global AI Summit of this series to take place in the Global South. The India - AI Impact Summit represents a critical inflexion point, building upon existing multilateral initiatives while advancing new priorities, deliverables, and cooperative frameworks that shift focus from high-level political statements to demonstrable impact and tangible progress in global AI cooperation.

The Three Sutras

The India - AI Impact Summit will be anchored on three foundational principles, known as 'Sutras' — a Sanskrit term meaning guiding principles or essential threads that weave together wisdom and action. These three Sutras define how AI can be harnessed through multilateral cooperation for collective benefit:



PEOPLE

This sutra is grounded in the principle that AI must serve diverse human needs while upholding the highest standards of safety, trust, and inclusive participation.



PLANET

This Sutra embodies the principle that technological advancement must align with planetary stewardship, ensuring that AI contributes to rather than compromises global sustainability goals.



PROGRESS

This sutra also emphasises the application of AI to accelerate socioeconomic development across key sectors, including healthcare, education, governance, and agriculture.

These three Sutras: People, Planet, and Progress, will permeate the deliberations and activities in the run-up to the Summit through seven focused themes, termed 'Chakras'—centers of focused energy and action in the Sanskrit tradition. Each Chakra represents a specific domain of multilateral cooperation that channels collective efforts toward tangible outcomes and deliverables, ensuring that the global Alagendamoves decisively from aspiration to impact.

The Seven Chakras

Building upon the three foundational Sutras, the deliberations at the Allmpact Summit will be organised around seven interconnected Chakras, areas of multilateral cooperation designed to channel collective energy toward holistic societal transformation. These Seven Chakras represent focused domains of international collaboration that translate the broader principles of People, Planet, and Progress into concrete areas of action:



Human Capital

Impact of AI on Employment, Skilling & Workforce transformation

The Human Capital Chakra envisages addressing these challenges by developing global frameworks/mechanisms/platforms for equitable AI literacy and skills development, enabling workforce transition, and establishing mechanisms that ensure the benefits of AI-driven productivity gains are shared broadly across society. This chakra is anchored in the principle that technological advancement must enhance rather than undermine human potential, ensuring that AI becomes a tool for inclusive economic growth and social mobility rather than a source of division and displacement.



Inclusion

Inclusive, Culturally Grounded, and Accessible Al

The Inclusion Chakra addresses these challenges by fostering the development and deployment of AI systems that are inclusive by design, locally relevant, culturally respectful, and reflective of diverse knowledge systems. This chakra focuses on the design, development, and deployment of AI systems that actively account for and accommodate the diverse needs, identities, and experiences of all users, ensuring accessibility for individuals with disabilities, cultural and linguistic representation, and meaningful participation in AI development and usage. This chakra aims to ensure the meaningful participation of all individuals and communities, particularly those traditionally left behind, enabling them to actively shape and benefit from AI's transformative opportunities while preserving and celebrating the rich diversity of human experience.



Safe and Trusted AI

Technology-Led Governance for All

The Safe and Trusted AI Chakra addresses these challenges by democratising access to technology- enabled governance tools and frameworks. Rather than relying solely on regulatory instruments that may stifle innovation, this chakra focuses on empowering all nations and stakeholders with the technical capabilities needed to govern AI effectively. This chakra ensures that robust AI governance becomes accessible globally, enabling the Global South to participate meaningfully in AI oversight while fostering continued innovation and technological advancement.



Resilience

Sustainable Al for Resource-Efficient Innovation

The Resilience Chakra seeks to promote innovations that enhance the energy efficiency and sustainability of the entire AI lifecycle, from model development to deployment. This is essential to ensure that all countries, including those with limited resources, can equitably access and benefit from AI-enabled solutions.



Science

Al to Advance Global Science

The Science Chakra focuses on expanding avenues for the responsible use of AI to accelerate scientific innovation and discovery. It promotes inclusive collaboration and cooperation within the global scientific community, particularly by strengthening research ecosystems and partnerships in the Global South, enabling AI to be more effectively leveraged to address shared and region-specific challenges. This chakra also emphasises the importance of open, interdisciplinary research, ensuring that AI-driven science contributes meaningfully to global knowledge and sustainable development.



Democratising AI Resources

Equitable Access to Foundational AI Resources & Infrastructure

The Democratising AI Resources Chakra addresses these fundamental imbalances by fostering multilateral frameworks for equitable access to the essential building blocks of AI innovation. This chakra aims to transform AI development from a privilege of the few into a collaborative endeavour that harnesses collective intelligence and serves collective needs. By ensuring that all nations have access to the tools necessary for AI innovation, this chakra will enable the emergence of diverse AI solutions that reflect the full spectrum of human experience and contribute to truly inclusive technological progress.



Social good

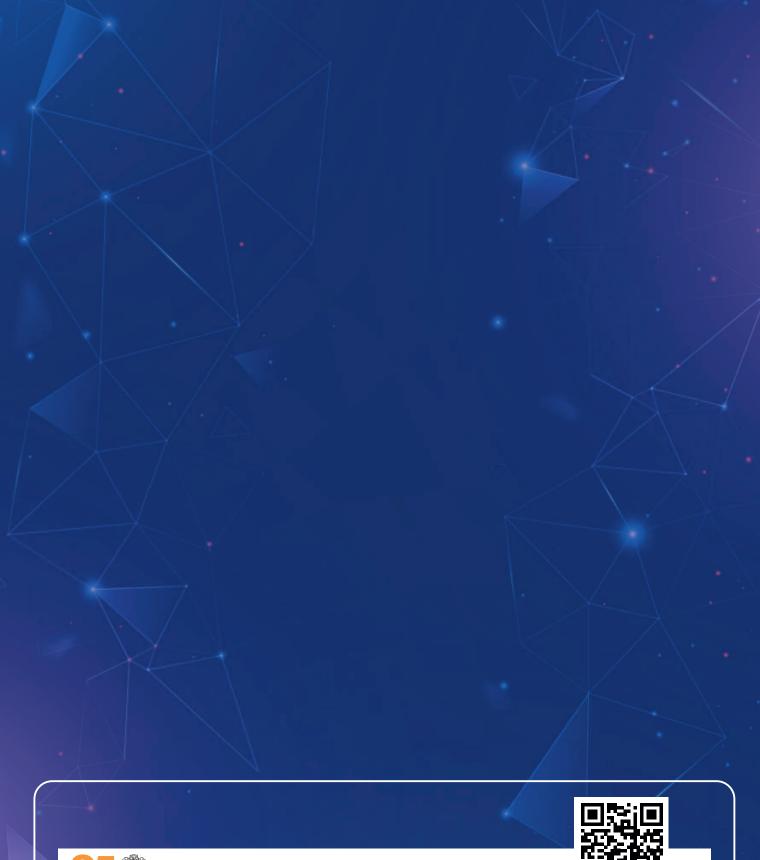
Al to Advance Development Goals

The Social Good Chakra addresses these challenges by enabling the identification and scaling of AI applications that advance public interest goals across critical sectors. This chakra focuses on creating frameworks and platforms that enable successful approaches to be adapted and replicated across different geographies. By establishing mechanisms for knowledge and resource sharing and cross-border collaboration, this chakra ensures that AI-driven innovations for social good can be more widely and equitably realised, allowing all countries to benefit from proven development solutions.

Each Chakra represents a distinct yet interconnected sphere of multilateral engagement through which the transformative potential of AI can be harnessed to deliver tangible benefits aligned with the three Sutras. Together, they offer a comprehensive framework that enables countries, international organisations, and stakeholders to coordinate AI deployment strategies that enable the maximisation of collective benefits while still addressing shared challenges.

The AI Impact Summit embodies a commitment to amplify the positive impact of AI and aims to build consensus so that AI's Impact is realised equitably by all. Through the three foundational Sutras and seven focused Chakras, the global community will forge concrete pathways that ensure AI serves the Planet, empowers all People, and accelerates Progress for every nation.

The AI Impact Summit embodies a commitment to amplify the positive impact of AI and aims to build consensus so that AI's Impact is realised equitably by all. Through the three foundational Sutras and seven focused Chakras, the global community will forge concrete pathways that ensure AI serves the Planet, empowers all People, and accelerates Progress for every nation.







www.psa.gov.in

Room No. 33027, 3rd Floor, Kartavya Bhavan 3, Central Secretariat, New Delhi - 110001 📞 011 2469 5488



f /prinsciadvoff

in @prinSciAdvOff

prinSciAdvOff

