

Is Generative AI beginning to deliver on its promise in India?

Aldea of India update

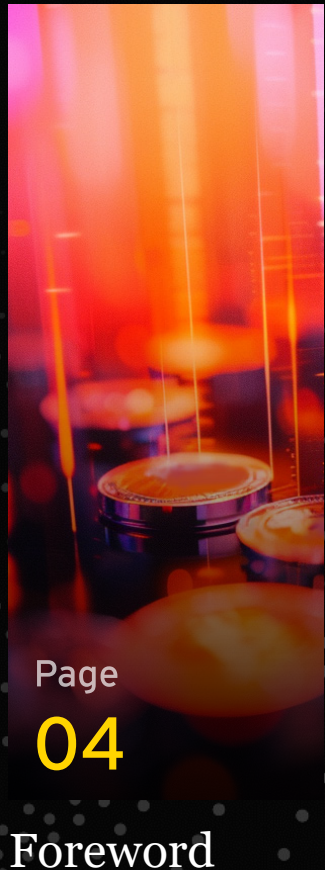
May 2024

■ ■ ■
The better the question. The better the answer.
The better the world works.

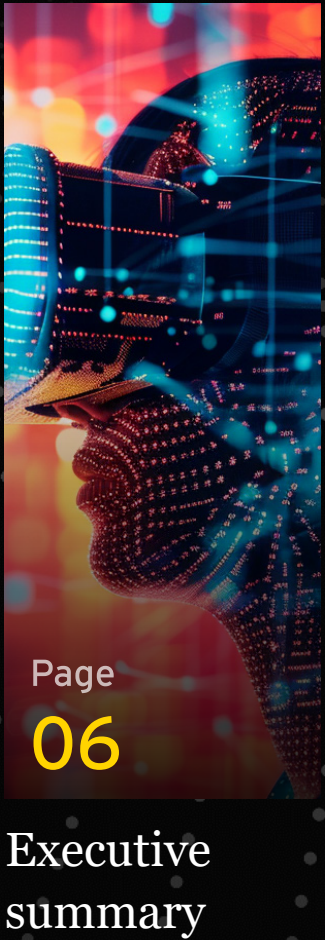
EY

Building a better
working world

Contents



Page
04
Foreword



Page
06
Executive
summary




Page
10
An evolving
value-cost
equation




Page
16

The race for AI
innovation



Page
22

Indic AI
roundup



Page
28

AI transaction
monitor



Page
32

India policy
matters



Foreword

Innovation in Generative AI (GenAI) continued at a frenzied pace in 2024.

Many companies in India launched proofs of concept (POCs) while deploying their first GenAI-powered apps into production. Big tech vendors brought their models to India with a range of options for enterprise users. Start-up innovation was rapid across a range of domains, including language, which saw the introduction of a raft of Indic models to cater to our unique needs and culture. We also saw the first graphics processing unit (GPU) clouds ready to offer GPUs on tap to enterprises.

However, against this vibrant backdrop, the pace of delivery of value for enterprises is still slow. Problems of hallucination and alignment of LLM responses with enterprise policies are still not fully resolved. Data privacy and data sovereignty concerns prevail and are especially acute in more regulated industries. Also, the cost equation has been quite dynamic for Indian enterprise users of LLM services, slowing down investment decisions.

This report probes the value-cost question for GenAI and recommends a two-speed approach. On one hand, companies need to continue rapid experimentation but in a more programmatic and guided fashion.

They will have to structure their innovation programs, by moving beyond a slew of shiny POCs to focus on those that can, if successful, really drive full scale functional transformation. The mandate is clear - get POCs out of labs and into useful production deployments.


However, in parallel, there is also a lot of heavy lifting to be done to build the enterprise AI platform to support transformation. This platform integrates GenAI approaches with the front-end experience platform, the automation framework, 'classical AI' approaches and a modern data stack.

We had published our report, "The A/idea of India: Generative AI's potential to accelerate India's digital transformation", in December but given the scorching pace of change, we did not want to wait another year to bring to you the latest edition. We hope you find this insightful.

Happy reading!



Mahesh Makhija
Partner and
Technology Consulting
Leader, EY India



Executive Summary

GenAI enters the enterprise

GenAI is transitioning from the initial hype to real-world applications as organizations progress from demos to rolling out POCs. EY analysis suggests that the adoption of GenAI varies widely across different types of enterprises. About 15% to 20% of POCs by domestic enterprises have been rolled out into production while among Global Capability Centers (GCCs), this number is between 30% and 40%. While GCCs are swiftly developing their capacities, domestic Indian enterprises are proceeding with more caution.

Progress has been rapid in a few key areas, especially in internal employee facing processes. Our analysis reveals approximately one-third of the use cases are focused on utilizing GenAI to execute point tasks using intelligent assistants while another 25% are related to marketing automation using various tools. Document intelligence is emerging as a key opportunity, with ~20% of the use cases focusing on document summarization, enterprise knowledge management, and search. Increasingly, GenAI assistants are powering enterprise intranets. Companies are also interested in using GenAI to build customer-facing chatbots for service and/or enhanced experience and coding assistants as well as for automating internal processes.

There are, however, several challenges that currently hinder the widespread enterprise-grade application of GenAI. Issues of reliability and functionality, with hallucination being a primary concern, have not been completely sorted yet. There are also limitations with respect to orchestrating actions - tracking multiple, long running transactions. While regulatory frameworks evolve, concerns around data privacy and data sovereignty are significant barriers to enterprises investing further in GenAI. While the situation is changing, many proprietary models are not available on local infrastructure in India today.

Another key focus area is enterprise readiness. To be able to leverage GenAI, enterprises need to significantly overhaul their digital and data architectures. Finally, the changing cost dynamics associated with adopting GenAI are causing Indian

companies to reassess their strategies. Today, it could cost almost as much as INR0.09 per word/token processed by an LLM (this is for GPT-4; GPT4o costs will be less) and this can lead to processing costs ballooning.

Innovation continues at a frenzied pace

The first few months of 2024 saw rapid innovation in GenAI across a range of domains. Foundation models evolved rapidly, demonstrating capabilities in advanced language understanding, photorealistic video generation and robotics integration. The open-source movement was a strong driver of innovation and is now backed by the largest companies in the world. Enterprises learned to leverage new techniques that harness the power of this new technology while recognizing its limitations. There was rapid introduction of new products from start-ups and new GenAI releases from the large cloud and platform players.

The capabilities of GenAI models have significantly advanced in various aspects. AI capabilities have expanded beyond text to include images and videos, with models like OpenAI's Sora enabling the creation of photorealistic videos from text prompts. This opens up new possibilities in content creation, smart virtual assistants, and virtual reality.

Enterprises are enhancing the functionality of language models through Retrieval Augmented Generation (RAG), which integrates company data into prompts, improving the relevance and accuracy of outputs. Companies like Cohere are advancing this with models that support complex data structures and multilingual content.

AI is progressing towards handling complex, multi-domain tasks using agent frameworks. These agents can operate autonomously, like in healthcare, where they can diagnose based on lab reports and patient history, blending multimodal data and enterprise resources to function effectively across various workflows.

GenAI is also challenging traditional search engines by providing direct, coherent, and contextually relevant answers, potentially transforming how users interact with and consume online information.

The drive for sustainability is pushing the development of smaller, more efficient language models like Microsoft's Phi and Apple's MM1. These models are designed for energy efficiency and privacy, suitable for on-device applications, and optimized for performance with innovations in architecture and hardware.

Cloud and platform services are increasingly integrating AI into their products, indicating a trend where enterprises will likely rely on these AI-enhanced cloud services for implementing AI strategies. For example, Google Gemini 1.5 Pro has enhanced breakthrough long context understanding and multimodality. Meta's Llama 3 is deployed across Facebook, Instagram, WhatsApp, and Messenger. SAP's GenAI copilot Joule integrates into its cloud enterprise portfolio and is adept at navigating features and streamlining tasks. Adobe GenStudio is designed to support marketing teams.

India charts its own path

In India, there is an increasing focus on developing Indic Large Language Models (LLMs) that cater to the diverse linguistic landscape of India, as a substantial portion of the Indian population accesses digital and internet interfaces in native languages. Domestic companies across the spectrum are innovating across a range of domains while Indian GCCs are leveraging GenAI to establish themselves as global COEs and there is rapid innovation. Furthermore, efforts by the government, large business groups, and start-ups to build robust domestic GPU infrastructure are set to bolster enterprise adoption of GenAI.

M&A activity continues to be strong in the GenAI space with Technology, Industrial and Energy Utilities, and Consumer Products showing increased activity. Some of the top deals include Krutrim SI Designs, AiDash, and DecisionPoint. Along with various developments in the private sector, government-driven AI initiatives are also on the rise.



The Indian government is striving to strike a balance between promoting AI development and mitigating associated risks. It has outlined plans to operationalize the IndiaAI Mission, which includes investments across all foundational aspects of an AI ecosystem. Additionally, it has issued advisories urging digital platforms to exercise due diligence when dealing with AI-generated content, emphasizing the need for responsible AI utilization and governance.

Key takeaways for enterprises

To enhance the delivery of enterprise value from GenAI, companies need to focus on a few key areas:



Transition from ad-hoc AI experiments to scalable programs, emphasizing the delivery and measurement of value from POCs. Expand the scope from specific use cases to comprehensive, agent-based applications for transformative impact in key business functions like drug discovery, content creation, and coding.



Monitor the financial implications and accessibility of AI models. They can opt for a hybrid approach in selecting language models based on specific needs and cost-effectiveness, while keeping abreast of significant cost differences among leading models like GPT-4 and its alternatives.



Develop a holistic enterprise AI platform, integrating with existing architecture rather than focusing solely on AI models. Leverage both leading closed-source models like GPT-4 for initial development and mature open-source options for sustainable growth, ensuring alignment with traditional AI methods, enterprise automation technologies and GenAI functionalities of COTS platforms.



Reinforce AI governance frameworks by setting clear priorities and policies across various AI use cases. Enhance data security measures to protect against leaks and ensure compliance with data sovereignty. Reassess approaches to AI transparency, reliability, and bias, and consider appointing a dedicated AI officer to steer AI strategy and governance.



An evolving value-cost
equation



An evolving value-cost equation

Early GenAI adopters have made slow but definite progress, moving **from demos and POCs to production releases**.

EY analysis suggests that **enterprises are at different levels of adoption maturity**. GCCs are moving fast to establish their capability as COEs, but domestic India enterprises are more cautious. While **some are prioritizing internal apps** due to lower risk and verification ability, others are using customer-facing GenAI apps.

With limitations in current LLM tech stacks, data privacy and sovereignty issues and a lack of preparedness among enterprises, **broader functional transformation is slow**.

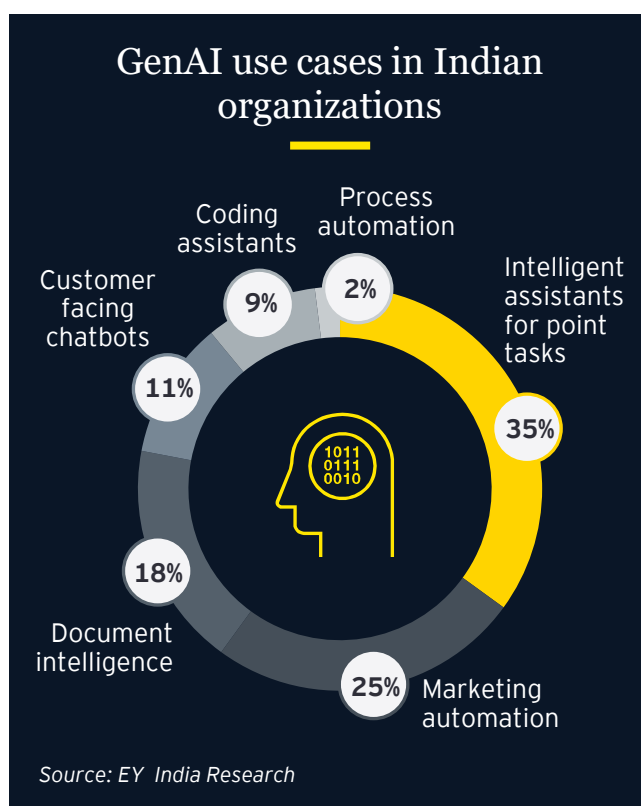
Companies need to adopt a two-speed approach - deploying 'fit for purpose' use cases that generate immediate value while preparing for functional transformation that the next wave of tech innovation will enable.

Enterprises go live with their first GenAI-powered apps

GenAI is slowly moving from hype to reality in 2024 as organizations deploy GenAI applications beyond POCs and into real-world production environments.

In India, companies have slowly started adopting GenAI. Enterprises have different levels of maturity with some segments, especially digitally savvy sectors, showing greater progress. Our estimates suggest that 30% to 40% of POCs by Global Capability Centers (GCCs) have been rolled out into production while among domestic enterprises this number is 15% to 20%.

Progress has been rapid in a few key areas. Approximately one-third of the use cases are focused on utilizing GenAI to execute point tasks using intelligent assistants. Another 25% are related to marketing automation using various tools; enabled by text generation and multimodal



capabilities, such as text-to-image and text-to-video. Document intelligence is emerging as a key opportunity, with ~20% of the use cases focusing on document summarization, enterprise knowledge management, and search. Increasingly, GenAI assistants are powering enterprise intranets. Companies are also interested in using GenAI to build customer-facing chatbots for service and/or enhanced experience and coding assistants as well as for automating internal processes.

Challenges with current GenAI implementations

While companies are adopting GenAI platforms for specific tasks, progress on broad functional transformation remains slow. There are several reasons for this, as detailed below.

01 GenAI's capability is yet to reach enterprise-grade levels of functionality and reliability

Hallucination continues to be an issue in many enterprise use cases. LLMs hallucinate where they need to navigate across a large document corpus, web pages or structured data stores. Several techniques have led to significant progress in this area. A widely used technique is Retrieval Augmented Generation (RAG) where a search is done on an enterprise corpus and results are passed to the LLM in the prompt, thus reducing hallucination. Significant progress in this area has been made in storing, indexing and retrieving enterprise data. Also, one can now pass more information in the prompt through longer context windows, which helps reduce hallucination.

Another issue with LLMs is their current inability to orchestrate actions. In the context of an enterprise, one needs to automate a process across multiple steps, systems and personas. Also, many transactions in an enterprise are long running and use workflow systems to help track and manage the state of items. This is a challenge with the current technology. However, there are many emerging innovations, especially in the area of agentic approaches, and 'Large Action Models,' which can learn to use different enterprise tools and coordinate transactions across systems. This is an emerging area of research, far from mature, but with significant implications for enterprises.

Last but not the least is the 'black box' nature of the technology. Enterprises need to understand why LLMs respond in a certain way. This is being addressed via a range of new research into explainability including through techniques like 'chain of thought prompting' where the LLM works through a guided series of mini prompts that help in explaining the steps in the reasoning process of LLMs.

02 Data privacy and data sovereignty are holding back GenAI investments from enterprises

Many companies are still thinking through their security and compliance approach to GenAI implementations. Key questions relate to the treatment of non-public private information, company specific data and the movement of data outside India.

03 Companies need to transform their digital and data architectures to leverage the benefits of GenAI

A prerequisite to being able to take advantage of GenAI is a fully digital customer and employee journey, automated processes and an evolved data landscape. While digital natives and large conglomerates are well on their way on this journey, many enterprises still need to do the hard work of building these foundational building blocks before they can leverage the power of GenAI.

04 A changing cost equation is slowing adoption

Indian enterprises embarking on GenAI projects are realizing the significant lifetime ownership costs involved, sparking concerns over affordability and tangible business benefits compared to traditional methods. Large-scale GenAI deployments often require committed spends to cloud providers, but many companies are hesitant due to the substantial monthly consumption commitments or expensive pay-as-you-go options.

Cost implications of GenAI in production

While many Indian enterprises have GenAI projects, they are increasingly realizing the true lifetime ownership costs associated with GenAI models. This has raised concerns about the ROI of these applications and their actual business benefits compared to traditional practices. In numerous instances, the deployment of large-scale GenAI applications necessitates a reserved processing capacity approach for cloud consumption, offering discounted pricing contingent upon the enterprise committing to a minimum number of consumption units per month. However, many enterprises are hesitant to commit to the substantial monthly consumption amounts typically required for such contracts. The popular pay-as-you-go pricing can also prove to be expensive.

We anticipate the cost challenges to be addressed in the coming quarters as open-source models reach enterprise grade levels and closed models increase capability and get commoditized for basic features. Since a large part of the cost associated with GenAI applications stem from model API usage by tokens, it is crucial to make appropriate choices regarding which GenAI model to deploy (bearing in mind that models with larger context windows and parameters entail higher costs), which country-based API instances to utilize (with data localization being a key consideration), and which Deep Learning technique, LLM or otherwise, to employ where. For instance, a simple named entity extraction task could be handled by NLP techniques, whereas an intelligent

Cost of running different operations using GenAI based on file type

Multimodal*			
Typical token range	Input cost	Output cost	
1,000	\$0.005	\$0.015	
Text based file			
Typical token range	Estimated input cost	Estimated output cost	
1,000	\$0.01	\$0.03	
Text-to-speech/ Audio		Image based file	
Usage (no. of characters)	Estimated usage cost	Usage (no. of characters)	Estimated usage cost
1,000	\$0.02	1,000	\$0.04

Costs are illustrative and based on a recent state-of-the-art Gen AI model for text-based, text-to-speech and image-based files

*GPT 4o pricing (multimodal) as of 13 May 2024


The costs in GenAI applications typically arise from the following areas

1



Documents (PDF, images, text files, etc.) uploaded into the application (like copies of government ID documents, application forms, supporting documents such as address proofs)

2



Words inserted into the query boxes (prompts) on the input side. Words from answers generated on the output side

3



Other computational running costs, like speech to text and text to speech conversions and local language translations, if required. Besides these, there would be periodic costs for fine tuning the GenAI model to a specific domain knowledge as well

summarization task might necessitate the use of an advanced LLM. With careful selection of tools and applications, GenAI has the potential to enhance business outcomes while delivering a good return on investment (ROI).

In case of running costs, every word is converted into tokens (one token equals 0.75 words). The current cost is US\$0.01 and US\$0.03 per 1,000 tokens (or 750 words) on the input and output sides, respectively, for one of the state-of-the-art closed-source LLMs. However, model innovations and pricing are very dynamic in the world of GenAI. For example, OpenAI launched its GPT 4o model on 13 May 2024, which has lower latency, is at half the price of GPT 4 Turbo and comes with multimodal capabilities like text, voice and image. It can respond to audio input in 232 milliseconds to match human level responses.

Hence, when used at scale, the running costs of GenAI applications may spiral up quickly. Enterprises may proactively use FinOps tools offered by hyperscalers to keep an eye on these costs.

Key takeaways for enterprises

Shift from ad hoc experiments to programs that provide functional transformation

Many POCs do not survive past the boardroom presentation.

- ▶ Focus on production deployment of projects in flight. Deliver and track value delivered by all POCs.
- ▶ Move beyond discrete use cases to broader agent-oriented implementations aimed at functional transformation.
- ▶ In the medium term, identify and focus on implementing core use cases that will add disproportionate value - drug discovery in pharma, multimodal content creation in media, or code generation in tech services, for example.

Keep a close eye on the cost, commercial structure and the India availability of models

Many POCs get stranded due to high inference costs and a lack of India-based availability of a pay-as-you-go commercial structure.

- ▶ Finetuning models is expensive, so be clear on the value and deliverable.
- ▶ Adopt a hybrid approach for LLM selection based on use case. Employ large, closed LLMs like GPT4 Turbo for prompts needed for complex requirements like financial advice and large context windows. Utilize open-source LLMs like Llama3 and Mistral for simpler prompts like employee chat queries. Use NLP to determine the appropriate LLM for the prompt, balancing higher accuracy with lower costs when necessary.
- ▶ Note that GPT4, for example, is 20 times more expensive than GPT 3.5T and even more expensive for Llama3 in inferencing.

Models are not all you need.

Think in terms of platform and not model. Envision your enterprise AI platform grounded in your specific enterprise context and integrated with your architecture.

LLM capabilities need to be extended in the context of your enterprise architecture.

- ▶ Adopt an iterative approach to LLM choice. First build on GPT-4, Gemini or other advanced closed-source models. Keep a parallel track going to increase enterprise skills with open-source to migrate when feasible.
- ▶ Do not forget that most value today is still generated with structured data and classical AI approaches. Long proven technologies of machine learning that evolved and matured well before the generative era, are still the go-to method for majority of enterprise use cases that involve large volumes of structured data including those involving forecasting, prediction and optimization.
- ▶ Integrate your GenAI architecture with the enterprise automation approach. GenAI will need to work with the enterprise workflow and automation technologies like OCR and RPA to increase straight through processing.
- ▶ Watch for the coming platform AI wave. Leverage commercial off-the-shelf (COTS) platforms that are rapidly releasing GenAI functionality in their transactional capabilities.

Re-evaluate your approach to AI and data governance

- ▶ At the organization level, put in place the basics - an approach to prioritize investments and policies for simple use cases (like ChatGPT usage) and advanced use cases (like customer facing implementations).
- ▶ Strengthen your data and cybersecurity posture to account for data leakage to foundation models and data sovereignty.
- ▶ If you are an advanced user of AI, rethink the approach to explainability, reliability and bias given current advances.
- ▶ Consider appointing an AI officer to guide your company through the impending AI transition. ● ● ●



The race for AI innovation



The race for AI innovation

Model capabilities have rapidly evolved across several dimensions.

There is significant momentum towards building agents that can work beyond point tasks and across multiple domains to accomplish a given goal.

Models have become smaller, more energy efficient and more performant at specific tasks including on device use cases.

Cloud and platform players are rapidly advancing AI offerings.

Evolving capabilities

In 2024, model capabilities evolved rapidly, and there has been an explosion of innovation along several dimensions.

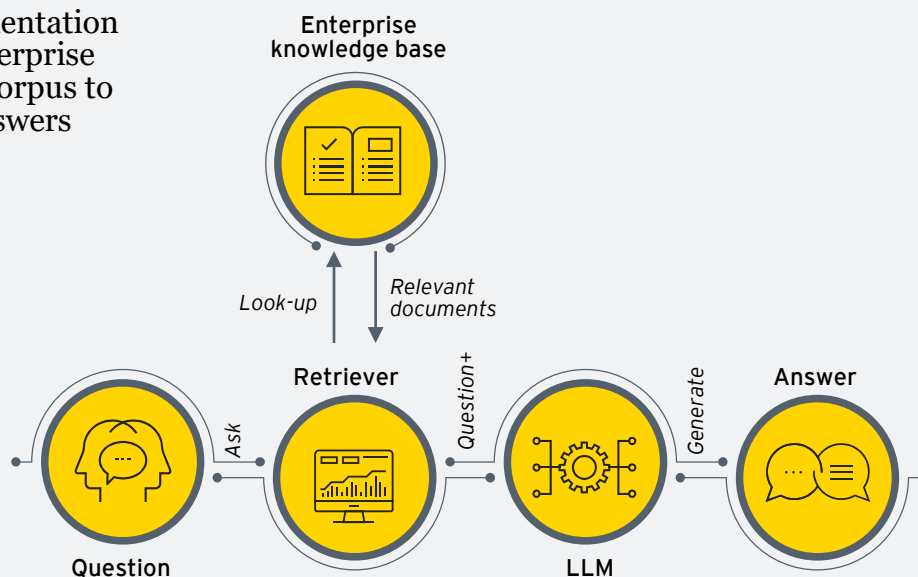
Moving beyond text: Most advanced models today are multimodal. This means they have been trained not just on text but simultaneously on images and increasingly on video. This broadens the scope of AI applications into innovative uses in content creation, smart virtual assistants, and virtual reality enhancements. For example, OpenAI introduced Sora, a text-to-video model that generates mesmerizing, photorealistic videos up to a minute long while adhering to the user's prompt.

Embedding enterprise context: Retrieval Augmented Generation (RAG) is helping businesses leverage LLMs by grounding the prompt to an LLM in the enterprise context. RAG workflows retrieve data from company databases and add it to the prompt sent to an LLM. This approach not only enhances the depth and accuracy of generated content but also enables businesses to tap into vast repositories of knowledge to inform their decisions.



The text-to-video model Sora by OpenAI can generate one-minute-long photorealistic videos based on user prompts (screen grab from Ships in Coffee video)

Using augmentation from an enterprise document corpus to improve answers



Several companies are now targeting RAG use cases and enterprise workflows. For example, Cohere's advanced language model Command R+ and the new Rerank 3 have improved enterprise search and retrieval systems. These enhancements allow for handling longer texts, support over 100 languages, and enable more effective searches on different and complex data structures.

A world of autonomous agents: LLMs are great at accomplishing point tasks but to be useful in an enterprise context they need to be able to complete long-running, complex tasks that often span multiple domains. This is done using an 'agent framework'.

An agent operates in a defined environment and seeks to accomplish a specific goal. For example, an autonomous medical assistant can analyze lab reports, diagnose symptoms and also recommend medication. This is difficult and needs an understanding of the patient's history, access to medical records, current state-of-the-art approaches to diagnosis and treatment and the nuances of managing patient interactions.

Significant progress has been made, but a lot remains to be done to bring this to enterprise use cases. Emerging approaches include:

- ▶ **Merging LLMs with multimodal inputs like visual and auditory data:** This integration facilitates sophisticated, real-time interaction strategies.
- ▶ **Action-oriented LLMs (Large Action Models):** Action-oriented LLMs seek to move LLMs into the territory of actions and workflows. In many cases this is accomplished by leveraging technology like Robotics Process Automation (RPA) to accomplish tasks. While using new form factors and hardware to access LLMs is a promising direction, devices like Rabbit and the Humane Pin

(also released this year) have yet to demonstrate feature completeness. Today the smartphone is the go-to device for most and this is unlikely to change soon.

- ▶ **LLMs learning to use enterprise resources:** In another approach, the LLM learns to use enterprise resources (applications, APIs, data) to accomplish tasks in much the same way as ChatGPT knows it has to browse the web when asked a question on current events.

An emerging view of LLMs is that they will become 'mini enterprise brains' which can comprehend requests and direct the flow of work to accomplish customer needs. These approaches are nascent but hold promise as we attempt to move from point tasks to full functional transformation in the enterprise.



Customer service agent: Klarna¹, a Swedish FinTech firm, utilizes GenAI agents for customer service, managing two-thirds of customer inquiries across 23 markets and 35 languages. Within four weeks of deployment, 2.3 million customer chats were handled, there was 25% decrease in repeat calls, and a reduction in average handling time from 11 minutes to two minutes. Closer home, Dukaan, an ecommerce platform that enables merchants to set up their own website, announced that they had laid off 90% of their related support staff after implementing an AI chatbot to answer customer queries.



Using safety-focused LLMs and connected to Avatar Cloud Engine, Nvidia's healthcare agent can conduct tasks, such as appointment scheduling and even video interact with patients through empathetic conversations.



Healthcare agent: Nvidia's healthcare agent, developed in collaboration with Hippocratic AI, is a GenAI-powered assistant that surpasses human nurses in patient interaction on video calls, offering empathetic conversations at a fraction of the cost. Utilizing safety-focused LLMs and connected to Nvidia's Avatar Cloud Engine (ACE), these agents conduct various healthcare tasks, including appointment scheduling, pre-operative outreach, and post-discharge follow-ups.



Coding assistant: Devin, an AI-driven tool launched by US-based start-up Cognition, revolutionizes end-to-end development, seamlessly translating project concepts into functional applications with autonomous learning and problem-solving capabilities. Acting as a collaborative coding partner, it empowers users to dictate its level of autonomy while continuously improving through self-learning from GitHub repositories, ultimately reshaping traditional coding practices with its dynamic approach.



Personal agent: Rabbit.ai introduced the concept of a 'Large Action Model'. It learns the UI and uses a device to be able to orchestrate actions for users, for example, playing a Spotify song with a simple voice command.

Smaller and more efficient models: The need for more sustainable and accessible AI technologies is driving the trend of smaller LLMs. Smaller models consume less energy and are more feasible for a variety of applications, including on-device implementations, which offer benefits, such as enhanced privacy and reduced data transmission costs.

Recent examples of this trend include Microsoft's Phi and Apple's MM1. These models represent significant advancements in making LLMs more efficient without sacrificing the quality of output. They are designed to provide robust AI capabilities while addressing the practical limitations of deploying large-scale models across various platforms, especially mobile and edge devices.

The efficiency and size reduction in models like Phi and MM1 are achieved through several key innovations including new model architectures, quantization to reduce the precision of numerical data used in computation and hardware optimization.

New ways to search the web: GenAI poses a challenge to traditional search engines like Google Search. AI models, by directly generating coherent and contextually appropriate responses, provide a user experience that can be more intuitive and efficient than sifting through pages of search results. Additionally, as models become better at understanding context and nuance, they can offer personalized responses that are immediate and highly relevant, further enhancing their appeal over standard search queries. This capability could redefine how information is queried and consumed, potentially reshaping the landscape of online search.

An example of this is Perplexity AI, an AI-first search engine. By focusing on conversational search experiences, Perplexity leverages natural language processing to provide more context-aware and interactive search results. This fresh approach, along with strategic innovations, has allowed it to carve out a niche, which is reflected in its growing popularity and strong backing from investors like Jeff Bezos and Nvidia.

Platform vendors gear up their offerings to enterprises

Platform and cloud players continue to embed AI into all their offerings. This is a key trend and it is likely that most enterprises will heavily leverage cloud and platform offerings as they implement AI strategies.

Enterprise COTS

SAP Joule

Joule, a GenAI copilot, integrates into SAP's cloud enterprise portfolio as an AI capability, adept at navigating features, providing advice, and streamlining tasks for a smoother workflow.

SAP Generative AI Hub

The hub is an AI supermarket for businesses to select tools easily, without requiring technical expertise.

Salesforce Einstein 1 Studio

This low-code conversational AI family integrates into Salesforce via Prompt, Model, and Copilot Builder, facilitating templated AI solutions.

ServiceNow Now Assist for TSM

ServiceNow enhances its workflow platform with more GenAI features for IT Operations Management AIOps and a GenAI-enabled Virtual Agent Designer, along with Impact AI Accelerators to expedite companies' return on investment.

Adobe GenStudio

The new GenAI-first offering supports marketing teams in planning, creating, managing, activating and measuring on-brand content.

Platforms

Google Gemini 1.5 Pro

Gemini 1.5 Pro, Google's latest model, excels with enhanced performance and breakthrough long context understanding. Built for multimodality, it seamlessly reasons across text, images, audio, video and code.

Google Gemma

Google's open-source GenAI aids rapid-response applications, integrated across various offerings. Google Cloud Next-April 2024 focused on leveraging GenAI for real-world problem-solving.

Apple MM1

Apple's MM1 model integrates advanced text-image understanding with 'in-context' learning for efficient user interaction. This is vital as AI embeds into everyday devices.

Meta

Meta introduced Llama 3, its latest open-source highly performant model. Llama 3 has now been deployed in 'Meta AI' across Facebook, Instagram, WhatsApp, and Messenger for tasks, learning, creation, and connection.

Oracle Cloud Infrastructure (OCI) Generative AI Service

Oracle combines LLMs from Cohere and Meta Llama 2 for diverse business use cases, offering multilingual support, enhanced GPU cluster management, and flexible fine-tuning options.

Microsoft Copilot

Microsoft embedded AI across products and launched Copilot Pro, offering priority access to the latest models like OpenAI's GPT-4 Turbo. ● ● ●



Indic AI roundup



Indic AI roundup

Indic LLMs supporting Indian languages are being developed as the large Indian population engages with internet and digital interfaces in local languages.

India Inc is **adopting GenAI with POCs going into production** for internal business use cases.

Indian GCCs are unlocking the potential of GenAI by positioning themselves as **COEs at a global scale within their parent enterprises**, particularly in the financial services, retail and consumer products sectors.

The government, **large business groups and start-ups are building crucial domestic and localized GPU infrastructure**, which will boost enterprise adoption of GenAI.

Progress on Indic LLMs

As the impact potential of GenAI becomes clearer to Indian enterprises, start-ups and investors alike, there is increased interest and effort in developing 'Indic LLMs' that support more than 20 key Indian languages. This need arises from the large Indian population that engages with internet and digital interfaces in local languages and can lead to mass adoption of GenAI the way some other digital technologies like UPI and social media apps have been adopted.

In this space, India has seen multiple Indic LLM initiatives like:

OpenHathi [Sarvam.ai]	Krutrim.ai [Ola]
Ambari Kannada LLM [Cognitive Lab]	Navarasa [Telugu LLM Labs]
Bhashini [Government of India]	Pragna [Socket Labs]
The Indus Project [Tech Mahindra]	BharatGPT [Corover.ai]
Airavata and IndicBART [AI4Bharat]	Tamil Llama [Abhishek Balachandran]
OdiaLlama [OdiaGenAI.org]	Hanooman [SML - Seetha Mahalakshmi Healthcare]
Kannada Llama [Tensoic]	

Some of these are wrappers/ finetuned models built around open-source foundational models built outside India. However, there are also others who are attempting to build foundational Indic LLMs as well.

Moreover, many LLMs built outside India, like Meta Llama 3, Google Gemini, OpenAI GPT and Claude Anthropic, are already supporting many Indian languages or have announced upcoming default support for major Indian languages.

Developing specific use cases

Internal applications are often prioritized for their lower risk and the ability to verify accuracy before making final decisions. In some cases organizations are also developing or using customer-facing applications of GenAI. Here are a few use cases:



IndiGo

IndiGo implemented GenAI, leveraging GPT-4 technology, to introduce an AI chatbot, 6Eskai. Instead of adopting a standard, off-the-shelf solution, the airline partnered with a hyperscaler platform to develop a customized version. The AI model, equipped with 1.7 trillion parameters, was programmed through prompt engineering, mimics human behavior, responds to emotions, and can add humor. Operating as a specialized customer service chatbot, 6Eskai manages customer queries, facilitates ticket bookings, and delivers personalized recommendations.²



Innovaccer

Indian healthcare cloud company Innovaccer employs AI to connect healthcare data from many sources, allowing doctors and other healthcare workers to make better patient care decisions. The AI-powered solutions allow healthcare organizations to aggregate and analyze patient data, improve care coordination, and drive population health management. AI capabilities assist health-care providers offer personalized and data-driven care while improving operational efficiency.³



Tata Steel

Tata Steel has partnered with an AI tech platform to leverage AI to produce Green Steel, reduce emissions, and enhance production quality. AI-powered Digital Twin and advanced analytics optimize pig iron unit operations, offering cloud-based recommendations for sinter and blast furnaces.⁴



Ola

Mobility major Ola implemented GenAI to offer dynamic ride pricing and accurately forecast demand and traffic. They introduced Krutrim AI, a voice-based AI assistant to improve rider and driver experiences. Employing Natural Language Processing, Machine Learning (ML), and Deep Learning, it recognizes patterns, analyzes complex data, and operates in multiple Indian languages. AI-powered route optimization algorithms enhance driver efficiency by suggesting an optimal route and provides personalized ride recommendations.⁵



Flipkart

Flipkart's knowledge assistant Flippi utilizes GenAI and LLMs to offer personalized product recommendations and enhance shopping efficiency. Semantic search technology deepens user intent, and the assistants' understanding goes beyond keyword matching. Exploring image-led discovery and multi-modal search diversifies options for intuitive product discovery. AI in fashion design generates innovative trends, enabling sellers to stay ahead in the market.⁶



HealthifyMe

Fitness app HealthifyMe has used GenAI to offer a multilingual conversational virtual coach, Ria 2.0. The AI coach responds to texts and voice commands, providing personalized health coaching. It tailors nutrition plans based on user data like age, dietary restrictions, lifestyle, health condition, and food preferences.⁷

GCCs unlocking GenAI potential

Over the years, GCCs have evolved from being cost centers to value-added centers and are now embracing GenAI as the next disruptor to enter a new era of innovation. Indian GCCs have stepped onto the global stage with a GenAI-first approach and have a pivotal role to play in shaping responsible usage. With their deep talent pool in AI/ML, GCCs hold the key to unlocking a treasure trove of AI-powered opportunities.

In Financial Services, GCCs are strategically intensifying the focus on AI technologies to drive product development, foster innovation, and optimize operational efficiency. By doing so, they position themselves as Centers of Excellence (COEs) within their parent enterprises. The use cases are focused on revolutionizing financial services by automating tasks, improving compliance, enhancing investment recommendations, and creating personalized experiences across diverse use cases, including risk assessment, customer service automation, fraud detection, and knowledge management.

GCCs in retail and consumer products sectors have taken a lead in exploring GenAI, followed by GCCs in Life Sciences and Pharmaceuticals. The objective is primarily to impact productivity and employee experience. Some of the common use cases being implemented are related to service desk related knowledge management for self-serve and self-solve, contract comparisons to identify discrepancies and niche areas in supply chain and warehousing. Here are a few examples.

Fidelity Investments

The company is actively leveraging GenAI to enhance services and streamline operations. Some notable use cases are 'Saifr' for Compliance Management and Fidelity AMP Platform for investment recommendations. While details are not publicly disclosed, the company is likely exploring more use cases, including in risk management, personalized policy recommendations, and dynamic pricing models. It is using GenAI to improve compliance, enhance investment recommendations, and explore innovative applications.⁸

American Express

The company is applying GenAI to financial services and credit cards. It is testing ideas for personal and small business use to improve customer behavior prediction, which will help with financial planning and decision-making. They are also exploring ways to improve services and enhance customer experiences.⁹

Schneider Electric

It has partnered with a hyperscaler platform to develop text and code generating solutions to enhance operational processes. It uses a Resource Advisor Copilot to streamline tasks, optimize resource allocation, and improve efficiency through enhanced data analysis and decision support. GenAI Finance Advisor gives analysts precise information, and Knowledge Bot assists customer care representatives with internal documentation.¹⁰

PepsiCo

PepsiCo is betting big on GenAI in the areas of productivity, marketing, customer service and feedback to develop new products and launch faster. Its experiment with GenAI in marketing campaigns led to the campaign cycle reducing from six to nine months to three to four months. With its snack Cheetos, the company utilized GenAI to give the product its "perfect shape, perfect flavor", on the basis of relevant customer feedback. This allowed the company to sustain margins on the product while driving revenue.¹¹

Bayer

Bayer has launched a pilot GenAI system to aid farmers and agronomists, utilizing proprietary data to train LLM. Developed by its crop science division, which has a significant presence in India, the project aims to benefit millions of smallholder farmers. The company aims to advance Indian and global agricultural solutions through integration of the GenAI system with digital offerings.¹²

Novo Nordisk

The pharmaceutical company leverages data and AI across the value chain, with access to billions of clinical data points driving AI-based innovation in R&D and helping overcome many challenges. The AI-powered NovoScribe platform enables automated authoring of structured documents, reducing the time needed for regulatory submissions. The company has been investing in a wide range of projects, including an AI Centre of Excellence and GenAI capabilities.¹³

India's start-up landscape

EY India's analysis of the publicly disclosed adoption of AI/GenAI by India's top 50 most valued Indian unicorns shows that approximately 66% are already using AI or GenAI technology, indicating that start-ups have been more proactive in terms of adopting AI technology as compared to the traditional players.

The start of 2024 has been positive for India's GenAI start-ups. The Indian start-up ecosystem produced its first AI unicorn (Ola Krutrim) in January, and early-stage start-up investors in India are optimistic about AI/GenAI initiatives.

While the Indian start-up community has been optimistic about Indic LLMs for some time, the ecosystem is also gearing up to focus on enterprise applications of GenAI and a few key themes are emerging in this space as the technology has evolved from pure LLMs to Large Vision Models (LVMs) and multimodal models.

The key themes and some start-ups in this space are:

- ▶ **Marketing and sales using GenAI:** VoiceOwl, Pixis AI, Rephrase AI, InVideo, UnScript AI, Hexo
- ▶ **Customer support using GenAI:** Haptik, Yellow AI, Gupshup, Dave AI, ClearFeed
- ▶ **Media creation, content generation and editing using GenAI:** Beatoven.ai, Visual Dub (Neural Garage), Dubdub.ai, Vitra.ai, LongShot, Instoried, Dubverse
- ▶ **DevOps, Infra and LLMops:** Maxim.ai, Portkey, Truefoundry, AuraML, dowhile.ai, inferless.com
- ▶ **Sectoral/domain specific GenAI apps and vertical LLMs:** Boltzmann (drug discovery), OnFinance (financial services), Expertia.ai (HR and recruiting), SydeLabs (cybersecurity)

The GenAI start-up environment in India is maturing towards building functional and domain specific enterprise applications and we can expect some marquee winners to emerge in these areas soon. As these use cases mature and scale up in production, end consumers and enterprises will start seeing the GenAI-driven magical experiences in daily life.

Availability of GPUs: local infrastructure for data privacy and regulations

The availability of compute is a critical requirement both to train models and to incorporate GenAI into applications. GPUs (the chips which are used for AI) have become very expensive as demand has skyrocketed in the GenAI era. Below are several developments that will ease out this constraint:

- ▶ All cloud providers have started offering GPUs as part of their infrastructure as a service offering. While GPUs are expensive, availability does not seem to be hindering current corporate GenAI roadmaps as many companies are still in the experimentation stage and will hit GPU constraints when they either do significant training of open-source models or deploy scaled GenAI apps into production.
- ▶ GPU clouds: Companies like Yotta are developing high-performance computing capabilities from data centers in India to enable the country's enterprises, start-ups and researchers to set up AI services. Yotta has partnered with NVIDIA to supercharge its data centers with best-in-class AI hardware. On the GPU infrastructure side, Yotta Data Services received its first tranche of 4000 NVIDIA H100 Tensor core GPUs this March. Yotta Data Services plans to take this number to 32,768 units by the end of 2025.
- ▶ Reliance Industries and Tata Group have also announced strategic partnerships with NVIDIA to develop India focused GPU infrastructure on top of the GH200 Grace Hopper Superchip.
- ▶ The Government of India has also made a provision of INR10,000 crore as part of the IndiaAI Mission, either through rent and sublet model or through a marketplace model supported by government incentives scheme. ● ● ●




AI transaction monitor



AI transaction monitor

 New focus sectors such as **technology, industrial, energy and consumer product goods (CPG)** have emerged in the M&A landscape in 2024.

 Broader solution offerings vis-à-vis conventional applications of GenAI have led to **increased M&A activity, also highlighting investor confidence** in India's AI prowess.

Year 2024 has been active in terms of M&A deals in the Indian AI market. This is indicative of a maturing and dynamic industry, driven by heightened investor interest, technological advancements and the increasing recognition of AI's potential across various sectors.


In 2023, investments were primarily focused on healthcare, life sciences and financial services. Investments have, however, diversified into various sectors such as technology, industrial and energy utilities, and consumer products goods (CPGs). This shift signifies a change in investment strategies compared to a year ago. Along with various developments in the private sector, government-driven AI initiatives are also on the rise.

A significant portion of investments have been in early-stage companies. Leading venture capital firms in India such as Lightspeed and Peak XV Partners have pioneered early-stage investments across various use cases and business models. SenseAI Ventures has announced an INR200 crore SenseAI Fund to invest in about 20 AI-first start-ups at the seed or pre-series A stage. Technology giants such as Google are also investing in Indian AI-first start-ups via their start-up accelerator program. Investors are expecting most value generation from companies focusing on AI tooling and infrastructure, development of GenAI-based personalized content generation and vertical AI companies building for specific domains.

While the Indian GenAI ecosystem has witnessed rapid uptick in the application and services space, there are significant opportunities for building native foundation models. Players such as Sarvam and Krutrim are working towards combining model innovation and application development to build population-scale solutions for India. Investments have primarily been directed towards four distinct business landscapes:

Top AI investment deals in 2024 in India

Target company	Target sector	Bidder company	Deal value (US\$ million)
Krutrim SI Designs	Technology	Matrix Partners	50
AiDash	Industrial, Energy and Utility	Lightrock	50
Decision Point (70% stake)	CPG	Latent View Analytics	39.1
Nanonets	Technology	Accel Partners, Elevation Capital, Y Combinator	29
Ema	Healthcare	Accel, Section 32, and Prosus Ventures	25
Neysa	HiTech	Matrix Partners India, Nexus Venture Partners, NTTVC, Anchorage Capital Group, Blume	20
Attentive AI	Enterprise Applications	Info Edge Ventures, Vertex Ventures, Surge, RedStart Labs, Peak XV Partners, Tenacity Ventures	7
PO Security	IT Security Company	Lightspeed Venture Partners with LPs (Pantheon, 57Stars, and Unigestion), Filter Capital, Innoven Capital	6.5
Louisa AI	FS	Palm Drive VC, Evolution VC, Nucleus VC, and Gaingels (Oxford University's investment arm)	5
RagaAI	HiTech	pi Ventures, Arka Venture Labs, Anorak Ventures, Mana Ventures, Exfinity Venture Partners, TenOneTen	4.7
Upliance.ai	Manufacturing	Khosla Ventures	4
SydeLabs	IT Security Company	RTP Global	2.5
Beatoven.AI	Media and Entertainment	Capital 2B, IvyCap Ventures, Upsparks and Rukam Capital	1.3



- ▶ Firms specializing in AI model training and optimization
- ▶ IT infrastructure entities, notably those involving data centers that strengthen AI applications
- ▶ Software developers, data scientists and machine learning engineers engaged in creating streamlined solutions to enable AI-powered operations
- ▶ Clients that procure software and cloud infrastructure services

Investing in AI companies necessitates a three-faceted evaluation. These are estimating performance and its subsequent impact, understanding the cost implications of AI, and considering AI compliance factors. Therefore, assessing potential AI capabilities of investment targets is becoming increasingly critical.

India's AI investment landscape is rapidly evolving, marked by strategic emphasis on AI, IoT, and government-driven AI initiatives. With substantial investments in R&D and talent development, India's potential as an AI hub continues to expand, offering fertile ground for innovation and skill enhancement. ● ● ●



India policy matters



India policy matters

● The Indian government has announced plans to **operationalize IndiaAI Mission**, including investing in all the building blocks of an AI ecosystem.

● It also issued **advisories for digital platforms and intermediates** to do the due diligence on AI-generated content.

Policymakers in India continue to balance promotion of AI and safeguarding against technology-related risk. Over the last few months, there have been several announcements related to both promotion and deployment of AI as well as regulating it to minimize the spread of misinformation, unlawful content and harmful AI generated content. For example, various advisories have been issued for social media platforms and intermediaries.

The Ministry of Electronics and Information Technology (MeitY) issued an advisory on 26 December 2023 for digital intermediaries and platforms. All social media platforms were required to comply with existing IT rules¹⁴ specifically focused on misinformation powered by AI and deepfakes. The announcement led to more detailed advisories of 2024, which are discussed below.

● 1 March 2024 ●

MeitY issued an advisory for intermediaries and platforms to conduct due diligence under the existing IT Act and rules¹⁵. Under the advisory:

- ▶ The use of under-testing or unreliable AI model(s), Large Language Models (LLMs), Generative AI (GenAI), software(s) or algorithm(s) and its availability to users on Indian internet must be done with explicit permission of the Government of India and be deployed only after appropriately labeling the possible and inherent fallibility or unreliability of the output generated.
- ▶ Intermediaries were required to submit an action taken cum compliance report within 15 days of the issue of this advisory.

A new advisory dated 15 March 2024 (discussed next) replaced this one.

● 15 March 2024 ●

The ministry advised intermediaries and platforms under the existing IT Act and Rules¹⁶ to ensure due diligence on the following in the usage of algorithms, AI models, LLMs, GenAI and software:

- ▶ Users should not be able to host, display, upload and share unlawful content.
- ▶ The AI platform/software should not permit biases or discrimination or threaten the integrity of the electoral process.
- ▶ The users should be made aware that they are working with an AI software and this may be enabled through a pop-up mechanism.
- ▶ Where synthetic content can be created using AI, it should be appropriately labeled as AI generated.

● 20 March 2024 ●

MeitY notified the Press Information Bureau (PIB) as the fact checker for the central government¹⁷. The press release added that since November 2019, PIB's Fact Check Unit has been tackling fake news related to policies, schemes, rules and regulations, programs, initiatives, etc.

The recent regulatory announcements reflect the government's focus on the impact of AI-generated content in the public domain and its approach whereby the onus has been put on the platforms and intermediaries.

Promoting development and adoption of AI in India

In September 2023, the government had released the India AI Report¹⁸, which details an action plan for the development and deployment of AI in India. Since then, the most significant development has been the announcement on 7 March 2024 on how the government plans to operationalize the IndiaAI Mission.



The announcement recognizes all the building blocks of an AI ecosystem i.e., computational power, data, algorithms, skills and entrepreneurship, though the level of details with regard to each varies.



The government has made a specific financial commitment toward AI innovation funding in India of INR10,372 crore over a period of five years.



It envisages enabling the development of indigenous LLMs and domain specific foundational models. Though not covered as a part of the AI Mission announcement, the government has released Bhashini, which provides for real-time translation of text and audio in Indian languages. The focus is to incorporate voice-based functions and promote content creation in different Indian languages.



The plan provides for setting up of compute infrastructure encompassing a minimum of 10,000 GPUs for access to start-ups and a research ecosystem for the next five years. The deployment and access particulars are yet to be announced.



The announcement also reiterates the government's intent to facilitate access to non-personal data; prioritize development of applications for important government sectors; talent development; and facilitate access to financing for start-ups. In the Budget speech, the Finance Minister had announced that a corpus of INR100,000 crore with a 50-year interest free loan would be set up to provide long-term financing or refinancing.

While the budget speech did not mention AI explicitly, this corpus is targeted at the private sector to scale up research and innovation significantly in sunrise domains.

Globally, AI regulation is at the top of policymakers' agenda. However, countries are developing distinctly different approaches keeping in view their legal and judicial systems. Within regions also there have been debates on issues such as usage of facial recognition software by police in EU. While approaches may range from voluntary guidance to mandatory rules, each jurisdiction is pursuing a risk-based approach. As policies evolve, there will be greater clarity.

In India, the recent announcements and actions reflect the government's focus on developing AI infrastructure, funding innovation and making benefits of AI accessible in Indian languages. However, the policy initiatives are evolving and we expect that over the next few months, there would be greater clarity on implementation, such as deployment of GPUs, and the plans for other building blocks, such as making non-personal data available. ● ● ●

Bibliography

1. <https://www.klarna.com/international/press/klarna-ai-assistant-handles-two-thirds-of-customer-service-chats-in-its-first-month/>
2. <https://www.indiatoday.in/technology/news/story/indigo-introduces-chatgpt-based-ai-chatbot-for-queries-ticket-booking-2469409-2023-11-30>
3. <https://innovaccer.com/resources/news/innovaccer-launches-industry-leading-healthcare-AI-solution-suite>
4. https://cms.eugenie.ai/eugenie_resources/from-davos-to-the-factory-floor-how-eugenie-ai-is-revolutionizing-industrial-sustainability/
5. <https://www.hindustantimes.com/technology/meet-krutrim-india-s-own-ai-developed-by-ola-ceo-bhavish-aggarwal-led-venture-101702638061355.html>
6. https://www.business-standard.com/companies/news/flipkart-bets-big-on-generative-ai-to-improve-customer-seller-experience-124042501023_1.html
7. <https://indianexpress.com/article/technology/social/healthifyme-wants-to-improve-your-diet-with-its-ria2-0-ai-assistant-here-is-how-5544698/>
8. <https://newsroom.fidelity.com/pressreleases/fidelity-investments--launches-innovative-new-regulatory-technology--regtech--business-to-help-finan/s/b776e743-d2cc-4f46-a730-1ae9cb5d2e07>
9. <https://timesofindia.indiatimes.com/business/india-business/american-express-india-hub-builds-ai-tech-for-global-operations/articleshow/106683653.cms>
10. <https://www.se.com/ww/en/about-us/alliances/microsoft.jsp>
11. <https://www.forbes.com/sites/bernardmarr/2019/04/05/the-fascinating-ways-pepsico-uses-artificial-intelligence-and-machine-learning-to-deliver-success/?sh=59cc50d5311e>
12. <https://www.bayer.com/media/en-us/bayer-pilots-unique-generative-ai-tool-for-agriculture/>
13. <https://www.novonordisk.com/news-and-media/news-and-ir-materials/news-details.html?id=134465>
14. pib.gov.in/PressReleaseframePage.aspx?PRID=1990542
15. Advisory_MeitY_01Mar2024 final (regmedia.co.uk)
16. Advisory 15March 2024.pdf (meity.gov.in)
17. pib.gov.in/PressReleaseframePage.aspx?PRID=20157865
18. IndiaAI-Expert-Group-Report-First-Edition.pdf (meity.gov.in)

Acknowledgements



Core team

- ▶ Mahesh Makhija
- ▶ Rakesh Kaul
- ▶ Rajnish Gupta
- ▶ Prashant Garg
- ▶ Gaurav Goyal
- ▶ Vijay Shankar
- ▶ Rohit Pandharkar
- ▶ Santosh Tiwari
- ▶ Shivani Nagpaul
- ▶ Kumar Abhishek
- ▶ Aditi Gokhale
- ▶ Nishant Kinra
- ▶ Ankan De
- ▶ Abhilash Shukla
- ▶ Sachin Premnath

Editorial

- ▶ Prosenjit Datta
- ▶ Radhika KTP
- ▶ Kaveri Nandan
- ▶ Vikram D. Choudhury

Design

- ▶ Rachita Gupta
- ▶ Ritika Saini

Our Offices

Ahmedabad

22nd Floor, B Wing, Privilon
Ambli BRT Road, Behind Iskcon
Temple, Off SG Highway
Ahmedabad - 380 059
Tel: + 91 79 6608 3800

Bengaluru

12th & 13th floor
"UB City", Canberra Block
No.24 Vittal Mallya Road
Bengaluru - 560 001
Tel: + 91 80 6727 5000

Ground Floor, 'A' wing
Divyasree Chambers
11, Langford Gardens
Bengaluru - 560 025
Tel: + 91 80 6727 5000

Chandigarh

Elante offices, Unit No. B-613 & 614
6th Floor, Plot No- 178-178A
Industrial & Business Park, Phase-I
Chandigarh - 160 002
Tel: + 91 172 6717800

Chennai

Tidel Park, 6th & 7th Floor
A Block, No.4, Rajiv Gandhi Salai
Taramani, Chennai - 600 113
Tel: + 91 44 6654 8100

Delhi NCR

Ground Floor
67, Institutional Area
Sector 44, Gurugram - 122 003
Haryana
Tel: +91 124 443 4000

3rd & 6th Floor, Worldmark-1
IGI Airport Hospitality District
Aerocity, New Delhi - 110 037
Tel: + 91 11 4731 8000

4th & 5th Floor, Plot No 2B
Tower 2, Sector 126
Gautam Budh Nagar, U.P.
Noida - 201 304
Tel: + 91 120 671 7000

Hyderabad

THE SKYVIEW 10
18th Floor, "SOUTH LOBBY"
Survey No 83/1, Raidurgam
Hyderabad - 500 032
Tel: + 91 40 6736 2000

Jamshedpur

1st Floor, Fairdeal Complex
Holding No. 7, SB Shop Area
Bistupur, Jamshedpur - 831 001
East Singhbhum Jharkhand
Tel: + 91 657 663 1000

Kochi

9th Floor, ABAD Nucleus
NH-49, Maradu PO
Kochi - 682 304
Tel: + 91 484 433 4000

Kolkata

22 Camac Street
3rd Floor, Block 'C'
Kolkata - 700 016
Tel: + 91 33 6615 3400

Mumbai

14th Floor, The Ruby
29 Senapati Bapat Marg
Dadar (W), Mumbai - 400 028
Tel: + 91 22 6192 0000

5th Floor, Block B-2
Nirlon Knowledge Park
Off. Western Express Highway
Goregaon (E)
Mumbai - 400 063
Tel: + 91 22 6192 0000

Pune

C-401, 4th floor
Panchshil Tech Park, Yerwada
(Near Don Bosco School)
Pune - 411 006
Tel: + 91 20 4912 6000

Proud to be part of
the face of the future

ey.ai unifying
platform

www.ey.com/en_in/ai/platform

Ernst & Young LLP

EY | Building a better working world

EY exists to build a better working world, helping to create long-term value for clients, people and society and build trust in the capital markets.

Enabled by data and technology, diverse EY teams in over 150 countries provide trust through assurance and help clients grow, transform and operate.

Working across assurance, consulting, law, strategy, tax and transactions, EY teams ask better questions to find new answers for the complex issues facing our world today.

EY refers to the global organization, and may refer to one or more, of the member firms of Ernst & Young Global Limited, each of which is a separate legal entity. Ernst & Young Global Limited, a UK company limited by guarantee, does not provide services to clients. Information about how EY collects and uses personal data and a description of the rights individuals have under data protection legislation are available via ey.com/privacy. EYG member firms do not practice law where prohibited by local laws. For more information about our organization, please visit ey.com.

Ernst & Young LLP is one of the Indian client serving member firms of EYGM Limited. For more information about our organization, please visit www.ey.com/en_in.

Ernst & Young LLP is a Limited Liability Partnership, registered under the Limited Liability Partnership Act, 2008 in India, having its registered office at Ground Floor, Plot No. 67, Institutional Area, Sector - 44, Gurugram - 122 003, Haryana, India.



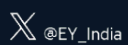
© 2024 Ernst & Young LLP. Published in India.
All Rights Reserved.

EYIN2405-013

This publication contains information in summary form and is therefore intended for general guidance only. It is not intended to be a substitute for detailed research or the exercise of professional judgment. Neither EYGM Limited nor any other member of the global Ernst & Young organization can accept any responsibility for loss occasioned to any person acting or refraining from action as a result of any material in this publication. On any specific matter, reference should be made to the appropriate advisor.

RS2

ey.com/en_in



Images generated using AI tools