



AI NEWS
OVERSIGHT OVERLOAD

AWS Launches AuditLLM, a Multiprobe Approach Tool for LLMs

AuditLLM is set to provide a comprehensive audit trail of all activities performed by LLMs.

Tanya | DECEMBER 23, 2025, 8:38 PM

Happy Llama 2026 March 13, 2026 Bangalore

Follow Preferred Source Google News WhatsApp Telegram

Register MS+

THE GCC ADVANTAGE FOR US MID-SIZED RETAILERS
ANALYSIS REPORT FOR SCALING ANALYTICS & AI IN RETAIL
JANUARY 2026
DOWNLOAD FREE REPORT

PRINT EDITION
February 2026
Subscribe
All Issues

AIM hosts 30+ AI conferences worldwide.
View Events Calendar

LATEST VIDEOS
LECUR DEFIES SILICON VALLEY WITH AIM
AI Workforce Era Begins | NVIDIA's AI Workers, Claude's \$2.5B Surge & India's Job...

AIM GAMES
Play Brain Games

KEY TAKEAWAYS

What Actually Matters.

AWS launches AuditLLM, a new tool for auditing large language models (LLMs). Utilize a multi-probe approach to provide comprehensive audit trails of LLM activities.

More

AWS has recently unveiled a tool named **AuditLLM**. By announcing their demo paper at **CIKM 2024**, they introduce a novel tool for auditing LLMs using a multi-probe approach.

This solution is designed to streamline the auditing process for LLMs by focusing on activity-based auditing.

"AuditLLM is set to redefine the way auditing is performed in the digital age," said **Aman Chadha**, researcher from Stanford University & AWS Leadership Member.

Authored by **Maryam Amirizani**, **Elias Martin**, **Tanya Roosta**, **Chirag Shah**, and Chadha this is a collaboration with the **University of Washington**, where the primary objective of AuditLLM is to provide a comprehensive audit trail of all activities performed by LLMs.

The tool leverages advanced technology to meticulously scrutinise the activities of LLMs, ensuring transparency, accuracy, and

Despite its sophisticated functionality, AuditLLM is easily navigable, allowing auditors to perform tasks without extensive technical knowledge.

As LLMs gain prominence across sectors, effective auditing tools like AuditLLM are increasingly crucial. With its launch, it addresses this need, setting a new benchmark for digital-era auditing.

Meet the AI Expert Building Indic LLMs with IITs

AIM got in touch with Chadha, who was working on building a **medical large language model** for India on top of Sarvam AI's OpenHathi, and released a research paper around it.

Given the amount of Indic languages speakers all over the world, Chadha expressed his happiness that models like Bharat GPT, Sarvam AI and Kissan AI are coming up. "But there's nothing on the healthcare or the medical side," he added, saying that he has been tracking all the recent announcements.

Join the Discussion
Be the first to share your thoughts
Open

Advertise with Us
Our Events Calendar
Our AI Trainings

Our Coverage of AI News

NVIDIA Posts \$68.1 Bn Q4 Revenue, Cashing in on Data Centre Surge

- OpenAI Releases Codex App for Windows, Brings AI Coding Agents to Developers**
 The new Windows app expands access to OpenAI's coding assistant as developers can now run AI agents, automate coding... [Read more →](#)
- The ePlane Company Begins Prototyping, Testing Electric Aircraft at IIT Madras**
 The company's new facility at IIT Madras will accelerate the development of electric air taxis while working with... [Read more →](#)
- Atlassian Introduces AI Agents in Jira to Enable Enterprise-Scale Human-AI Collaboration**
 With agents in Jira, teams can elevate AI from ad-hoc experiments to accountable teammates. [Read more →](#)
- Optimum Solutions & ABeam Consulting Launch India GCCs to Accelerate AI-Led Enterprise Transformation**
 The GCCs are expected to scale to 2,000 professionals by 2030. [Read more →](#)
- Karnataka Budget Proposes Robotics Innovation Zone in Bengaluru, AI Centre in Raichur**
 An Artificial Intelligence and Robotics Technology Park will be developed under the Indian Institute of Science to... [Read more →](#)

ABOUT THE AUTHOR

[Follow](#)

Tarunya
Contributor

Got a tip? Share confidential information with AIM.

[Editorial Standards](#) | [Reprints & Permissions](#)

What to Read Next



NVIDIA Posts \$68.1 Bn Q4 Revenue, Cashing in on Data Centre Surge



OpenAI Releases Codex App for Windows, Brings AI Coding Agents to Developers



The ePlane Company Begins Prototyping, Testing Electric Aircraft at IIT Madras



Atlassian Introduces AI Agents in Jira to Enable Enterprise-Scale Human-AI Collaboration



Optimum Solutions & ABeam Consulting Launch India GCCs to Accelerate AI-Led Enterprise Transformation



Karnataka Budget Proposes Robotics Innovation Zone in Bengaluru, AI Centre in Raichur

Trending

- So, Sam Altman Was Right About Indian AI Startups
- In Just 243 Lines of Python Code, Andrej Karpathy Recreates GPT From Scratch
- Cognizant Risks Recouping AI Investments as Competitors, Clients Tap Automation
- Claude Code vs Codex: Developers are Choosing Sides
- The Real Story Behind GCC Layoffs
- Figma Wants to Break the Fourth Wall in Design
- Texas Instruments Opens New R&D Centre in Bengaluru
- Google Unveils 10-Hour AI Certificate in Latest Learning Push
- LTS Sees Wave of Senior Leadership Reshuffle Over 8 Months
- In 630 Lines of Code, Andrej Karpathy Builds AI Research System Running on a Single GPU

Explore our newsletters

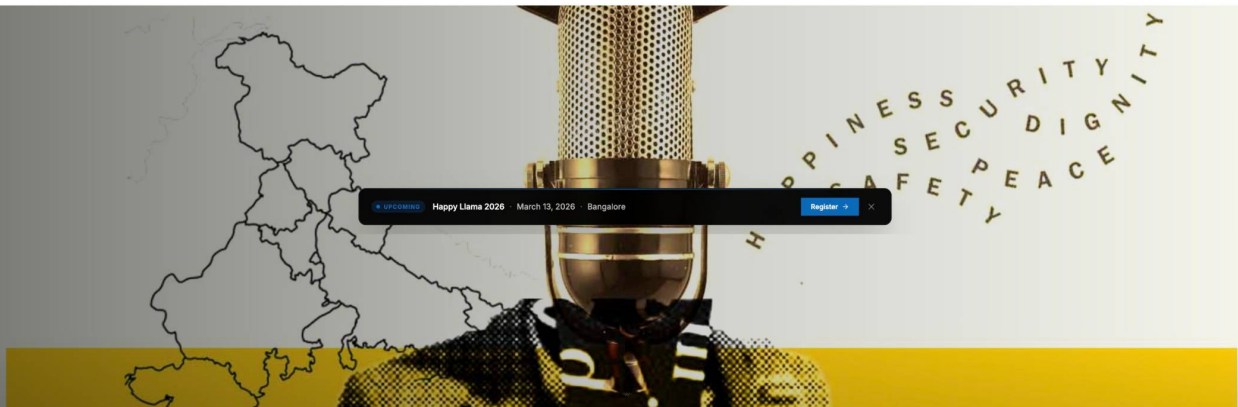
UPCOMING: Happy Llama 2026 March 13, 2026 Bangalore [Register →](#)

<p>WAKE UP INFORMED</p> <p>Make sense of the day's AI news and breakthroughs with our morning briefing.</p> <p>WEEKLY Belamy See the latest</p>	<p>INDUSTRY INTELLIGENCE</p> <p>Receive a roundup of AI adoption stories by industry vertical, curated for professionals.</p> <p>BI WEEKLY Sector6 See the latest</p>
--	--

aman.x64@gmail.com [Subscribe](#)

By signing up, you agree to our [Privacy Policy](#).

CONTINUE READING



UPCOMING: Happy Llama 2026 March 13, 2026 Bangalore [Register →](#)

Google Introduces IndicGenBench to Benchmark Indic LLMs Across 29 Languages

A benchmark to help in evaluating the generative capabilities of Indic LLMs, IndicGenBench is part of a slew of India-centric updates released during Google I/O Bengaluru 2024.

Donna Eva | DECEMBER 23, 2024, 8:30 PM

Listen 9 min

SHARE Save T

Follow: Preferred Source Google News WhatsApp Telegram

KEY TAKEAWAYS

What Actually Matters.

Launch IndicGenBench to evaluate generative capabilities of Indic LLMs across 29 languages.

Open source CALM enables developers to enhance language models for specialized applications.

More -

As part of several initiatives that Google has taken up in India to improve Indic LLM capabilities, Google Pay vice president and GM [Ambarish Kenghe](#) announced the launch of IndicGenBench.

A benchmark to help in evaluating the generative capabilities of Indic LLMs, IndicGenBench is part of a slew of updates released during Google I/O Bengaluru 2024. Kenghe said that the benchmark covers as many as 29 languages, including several Indian languages that do not currently have benchmarks.

Speaking to AIM, Google Cloud director of customer engineering and field CTO [Subram Nataraj](#) said, "In India, there are two main areas of focus: Addressing language-related issues, while the second involves large-scale transformations across various industries, be it in customer engagement or addressing the broader needs of the Indian population."

With a focus on improving language-related issues, Kenghe announced the open sourcing of DeepMind's Composition to Augment Language Models (CALM), allowing developers to combine specialised language models with Google's Gemma models. Interestingly, research on CALM had been done specifically by the Google DeepMind and Google Research teams in India, with the [paper](#) released earlier this year.

"Let's say you're building a coding assistant that can converse in English. Now, by composing a Kannada specialist model with CALM, you may be able to offer coding assistance to Kannada users as well," explained Kenghe.

This focus on Indic language LLMs comes as DeepMind expands [Project Vaani](#), a collaborative effort between Google and the Indian Institute of Science (IISc), wherein over 14,000 hours of speech data in 58 languages, has been made accessible to developers. This data was collected from over 80,000 speakers in 80 districts across the country.

As previously covered by AIM, this is being open-sourced as part of MeitY's flagship AI initiative, [Bhashini](#). These capabilities are also soon to be expanded as Bhashini also launched an initiative called [Bhasha Daan](#), to help crowdsource voice and text data in multiple Indian languages.

ABOUT THE AUTHOR

Donna Eva
Contributor

Follow

Got a tip? Share confidential information with AIM.

Editorial Standards | Reprints & Permissions

SYSTEMC+

THE GCC ADVANTAGE FOR US MID-SIZED RETAILERS

A NEW PLAYBOOK FOR SCALING SERVICES & AI IN RETAIL

DOWNLOAD FREE REPORT

PRINT EDITION

February 2026

Subscribe

All Issues

AIM hosts 30+ AI conferences worldwide.

View Events Calendar ->

LATEST VIDEOS

HEAL, WOMEN, REAL BENCHMARKS.

AI Workforce Era Begins | NVIDIA's AI Workers, Claude's \$2.5B Surge & India's Job...

Watch on YouTube ->

AIM GAMES

Trend Timeline ->

Chess

Play Brain Games ->

SHARE

WhatsApp

Facebook

LinkedIn

Print

Email

More


CONTINUE READING



Chinese Company SenseTime Releases SenseNova 5.5, Beats OpenAI's GPT-4o

SenseTime has expanded its suite of applications with the release of Vimi, an AI avatar video generator capable of creating short video clips.

Siddharth Jindal | DECEMBER 23, 2023, 8:38 PM



Listen 3 min | SHARE | Save | T

Follow | Preferred Sources | Google News | WhatsApp | Telegram

KEY TAKEAWAYS

What Actually Matters.

SenseTime launches SenseNova 5.5, outperforming OpenAI's GPT-4o with enhanced multimodal capabilities.

Achieves a 30% performance boost in mathematical reasoning, English proficiency, and command following.

More

SenseTime, a Chinese AI company, has released its upgraded **SenseNova 5.5 LLM** at the 2024 World Artificial Intelligence Conference & High-Level Meeting on Global AI Governance (WAIC 2024).

The new release includes China's first real-time multimodal model, SenseNova 5o, which offers streaming interaction capabilities comparable to OpenAI's GPT-4o.

The SenseNova 5.5 upgrade boasts a 30% improvement in overall performance compared to its predecessor, SenseNova 5.0, released just two months ago. Key enhancements include improved mathematical reasoning, English proficiency, and command following abilities, putting it on par with GPT-4o in terms of interactivity and core indicators.

Dr. Xu Li, Chairman of the Board and CEO of SenseTime, emphasised the significance of this release, saying, "This is a critical year for large models as they evolve from unimodal to multimodal. In line with users' needs, SenseTime is also focused on boosting interactivity."

The company has also introduced a cost-effective edge-side large model, reducing the cost per device to as low as RMB 9.90 per year. This move aims to facilitate widespread deployment across various IoT devices, including smartphones, tablets, and in-vehicle computers.

SenseTime has expanded its suite of applications with the release of Vimi, an AI avatar video generator capable of creating short video clips with precise control over facial expressions and upper body movements from a single photo. The company has also upgraded its SenseTime Raccoon Series, improving coding precision and response speed in the Code Raccoon tool.

To lower entry barriers for enterprise users, SenseTime launched the "Project \$o Go" scheme, offering a free onboarding bundle for new enterprise users migrating from the OpenAI platform.

The SenseNova Large Model has already been deployed at more than 3,000 government and corporate customers across various industries, including technology, healthcare, finance, and programming.

RECOMMENDED: Happy Liama 2026 March 13, 2026 Bangalore Register

healthcare, aiming to boost productivity and cost-efficiency in these sectors.

Kyutai, a French non-profit AI research laboratory, has introduced **Moshi**, a real-time native multimodal foundational AI model. This open-source project features a voice-enabled AI assistant offering capabilities that rival OpenAI's GPT-4o and Google's Astra.

Meanwhile, Anthropic's latest model, Claude Sonnet 3.5, continues to challenge GPT-4o by recently dethroning it and securing the top spot in both the Coding Arena and Hard Prompts Arena.



ABOUT THE AUTHOR

Siddharth Jindal

Tech Journalist

Siddharth is a media graduate who loves to explore tech through journalism and putting forward ideas worth pondering about in the era of artificial intelligence.

Follow

Got a tip? Share confidential information with AIM.

Editorial Standards | Reprints & Permissions

Download Free Report: THE GCC ADVANTAGE FOR US MID-SIZED RETAILERS

Print Edition: February 2026


AIM hosts 30+ AI conferences worldwide.

Latest Videos: AI Workforce Era Begins | NVIDIA's AI Workers, Claude's \$2.5B Surge & India's Job...

AIM Games: Trend Timeline, Chess

Vertical share menu with icons for WhatsApp, Telegram, Email, Print, and others.

Vertical share menu with icons for WhatsApp, Telegram, Email, Print, and others.



AI NEWS

MANUFACTURING MIRAGE

UptimeAI Secures \$14 Million in Series A Funding to Revolutionise Manufacturing with AI

UptimeAI has successfully raised \$14 million in a Series A funding round to boost AI in Manufacturing led by WestBridge Capital, with participation from existing investors.

MIA_AI_Reporter | DECEMBER 23, 2025, 8:37 PM

KEY TAKEAWAYS

What Actually Matters.

UptimeAI raises \$14 million in Series A funding to enhance AI solutions in manufacturing. Investment led by WestBridge Capital with participation from Emergent Ventures and Aditya Birla Ventures.

UptimeAI, a startup specialising in AI-powered solutions for manufacturing plants, has successfully raised \$14 million in a Series A funding round. The investment was led by WestBridge Capital, along from existing investor Emergent Ventures and new investor Aditya Birla Ventures.

UptimeAI Expanding AI Solutions for Industry

Founded in 2019 by Jagadish Gattu and Vamsi Yalamanchili, UptimeAI has developed an AI-driven platform that monitors and optimises manufacturing plant operations. The company's technology uses advanced machine learning algorithms to predict equipment failures, improve efficiency, and reduce downtime in industrial settings.

Strategic Growth and Market Expansion

With this new influx of capital, UptimeAI plans to:

- Accelerate product development
Expand its market presence in North America and Europe
Strengthen its team by hiring top talent in AI and machine learning

UptimeAI's solutions have already been adopted by several Fortune 500 companies across various sectors, including energy, chemicals, and heavy industries. The platform has demonstrated significant improvements in plant reliability and performance, with some clients reporting up to 20% reduction in unplanned downtime.

Future Outlook

This investment will allow UptimeAI to further enhance our AI capabilities and bring cutting-edge solutions to more manufacturing plants globally. We're excited to partner with WestBridge Capital and our existing investors to drive the next phase of our growth. The funding round highlights the growing importance of AI in industrial operations and positions UptimeAI as a key player in the digital transformation of manufacturing.

ABOUT THE AUTHOR

MIA_AI_Reporter

Contributor

Follow

UPCOMING Happy Llama 2026 March 13, 2026 Bangalore Register

Got a tip? Share confidential information with AIM.

Editorial Standards | Reprints & Permissions

CONTINUE READING



PRINT EDITION



AIM hosts 30+ AI conferences worldwide.

View Events Calendar

LATEST VIDEOS



Watch on YouTube

Trend Timeline

Chess

Play Brain Games

UPCOMING Happy Llama 2026 March 13, 2026 Bangalore Register

AI FEATURES

RECENT ARTICLES

Top 10 Uncensored LLMs You Can Run on a Laptop

UPCOMING Happy Llama 2026 March 13, 2026 Bangalore Register

Page 3/3