Vita Digital Journal

Life through Technology & Innovation

Computer Science Department Monthly

Vol. 1 | Issue 7 | August 2026

VITA DIGITAL JOURNAL

Life through Technology & Innovation

Computer Science Department Monthly

Volume 1 | Issue 7 | August 2026

Published by:
Department of Computer Science,
Marian College of Arts and Science,
Thiruvananthapuram.

Chief Editor: Livin M Miranda
Student Editors:
Santhipriyan M
Joel George Mathew

mcasmedia123@gmail.com
https://www.mcas.ac.in/

EDITORIAL

Technology continues to move beyond the boundaries we once imagined. August 2026 has highlighted an exciting direction in computing—the growing integration of Artificial Intelligence with scientific research, robotics, advanced hardware, and real-world systems. AI is increasingly moving from simply generating information to assisting with experiments, analysing complex data, and interacting with physical technologies.

For students of Computer Science and other science disciplines, this transformation offers both opportunities and responsibilities. The technologies we learn today can become the tools that solve tomorrow's challenges. What matters is not only how advanced technology becomes, but how creatively, ethically, and responsibly we use it.

As we present the August issue of Vita Digital Journal, let us remain curious, keep learning, and embrace emerging technologies with an open and innovative mind. The future of technology is being shaped today—and our students can be among those who shape it.

- Chief Editor
Mr. Livin M Miranda
Head, Department of Computer Science

INDEX

Enhance The Problem-Solving Skills and Creativity With Python
Eric Bobby Fernandez 1

6G Communication Technology: The Next Generation of Wireless Connectivity
Sivakrishna T 4

How Deepfakes Are Created and Detected Using Artificial Intelligence
Abhijith Anil 11

Physical AI and Robotics : A Comparative Study of
NVIDIA Isaac, Google DeepMind Gemini Robotics, and Boston Dynamics
V Sai Sesh 16

Brain-Inspired AI and Cognitive Robotics
Praveena M P 21

Measurable Sustainable Technology and ESG-Driven IT Strategy :
A Metrics-Based Framework for Green Digital Transformation
Ms. Anila Aneesh 25

Cloud 3.0 : The Next Evolution of Cloud Architecture
Ms. Mubeena J 30

ARPPO 2.0 36

Tech News 38

Editorial Board 42

Enhance The Problem-Solving Skills and Creativity With Python

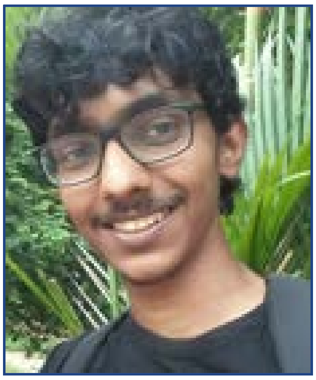

Eric Bobby Fernandez
S1 BSc Computer Science

Before we move onto the contents, here's one eyecatching quote on Python:

“Python is an experiment in how much freedom programmers need. Too much freedom and nobody can read another's code; too little and expressiveness is endangered."

- Guido Van Rossum

THE MINDSET: WHY PYTHON HELPS YOU THINK BETTER

- In languages like C++ or Java, you have to write a ton of background setup code just to print a simple message. You end up spending half your mental energy worrying about strict memory rules and compiler errors instead of the actual logic.
- Python strips all that noise away. Its syntax looks almost like plain English. When you don't have to fight the language itself, you can focus on algorithmic thinking--breaking a huge, scary problem down into tiny, manageable steps.
- It also lowers the penalty for making mistakes. When code is fast to write, you aren't afraid to try a risky idea. If it fails, you throw it out and try another approach in two minutes. That quick trialand-error loop is where real creativity happens.

HOW TO ACTUALLY IMPROVE: AN ADVICE FOR THE PYTHON DEVELOPERS

- Don't Code First—Solve First: The biggest mistake beginners make is opening an editor and typing immediately. Step away from the keyboard. Grab a notepad, sketch out a flowchart, or write down the steps in plain English. If you can't explain the logic in plain human words, writing it in Python won't magically make it work.
- Learn to Write "Pythonic" (Clean Over Clever): As you get better, resist the urge to write overly complicated solutions to prove you're smart. The core philosophy of Python is simple: Readability counts. Use built-in features like list comprehensions ([x for x in list]) or standard functions like zip() and enumerate() to keep your logic clean and concise, rather than nesting four for loops together.
- Test Small to Move Fast: Don't write a 200-line script and run it all at once--you'll end up spending three hours hunting down a bug on line 12. Use Python's interactive terminal (IDLE/REPL) or Jupyter Notebooks to test tiny pieces of your logic as you build. Getting instant feedback keeps your momentum going and builds confidence.
- Build Stuff You Actually Care About: Solving textbook math problems gets boring fast. Use Python to automate something annoying in your personal life--like scraping data for a passion project, sorting your downloads folder, or building a simple bot. When you care about the outcome, you stick around long enough to solve the tricky bugs.
- Embrace the Art of Refactoring: Getting your code to work is only step one. The real learning happens when you look back at working code and ask, " Can I make this easier to understand?" Cleaning up your own code after it works is how you transition from someone who just writes scripts to a real Software Engineer.

THE BIG PICTURE

At the end of the day, tools will come and go, but the ability to take an intimidating problem, dismantle it, and build a clean solution is a skill that lasts forever. Python just happens to be the canvas that lets you practice that craft with the least amount of friction. As the community often says: Life is short, use Python.

REFERENCES

- Python Software Foundation. (2026) - Python Documentation (https://wiki.python.org/moin/BeginnersGuide)

- Hunt, A., & Thomas, D. (2019) - The Pragmatic Programmer: Your Journey to Mastery [2nd Ed.] (Addison-Wesley)

- Miller, B. N., & Ranum, D. L. (2011) - Problem Solving with Algorithms and Data Structures Using Python [2nd Ed.] (Franklin, Beedle & Associates)

6G Communication Technology: The Next Generation of Wireless Connectivity

Sivakrishna T
S3 BSc Computer Science

INTRODUCTION

Communication technology has transformed the way people live, work, learn, and interact. From the early generations of mobile communication to today's advanced 5G networks, each generation has brought improvements in speed, connectivity, reliability, and accessibility. The next major step in this evolution is 6G (Sixth-Generation Wireless Communication Technology).

6G is expected to go beyond simply providing faster internet. It aims to create an intelligent, highly connected digital environment in which humans, machines, sensors, vehicles, robots, and artificial intelligence systems can communicate seamlessly. Research into 6G is already underway around the world, although commercial 6G networks are not yet available.

6G is generally expected to support extremely high data rates, ultra-low latency, massive numbers of connected devices, advanced artificial intelligence, integrated sensing, and highly reliable communication. It could become an important foundation for future technologies such as autonomous systems, immersive extended reality (XR), digital twins, smart cities, advanced robotics, and connected healthcare.

WHAT IS 6G COMMUNICATION TECHNOLOGY?

6G stands for Sixth-Generation Mobile Communication Technology. It is the proposed successor to 5G and is being developed to meet communication requirements that are expected to emerge in the 2030s.

While 5G focuses strongly on enhanced mobile broadband, massive machine-type communication, and ultra-reliable low-latency communication, 6G is expected to combine communication with capabilities such as artificial intelligence, sensing, computing, positioning, and automation.

6G should therefore not be viewed merely as "faster 5G." Instead, it is envisioned as a comprehensive platform for connecting the physical, digital, and intelligent worlds.

EVOLUTION FROM 1G TO 6G

1G – First Generation: Introduced in the 1980s, 1G provided basic analog voice communication. The networks had limited capacity and relatively poor voice quality.

2G – Second Generation: 2G introduced digital communication. It enabled improved voice quality, SMS messaging, and basic data services.

3G – Third Generation: 3G significantly improved mobile data communication and made mobile internet, video calling, and multimedia applications more practical.

4G – Fourth Generation: 4G brought high-speed mobile broadband and supported applications such as HD video streaming, social media, online gaming, and mobile cloud services.

5G – Fifth Generation: 5G introduced higher speeds, lower latency, massive device connectivity, and support for applications such as smart factories, connected vehicles, and advanced IoT.

6G – Sixth Generation: 6G is being designed to support an even more intelligent and integrated communication environment. It is expected to combine communication with AI, sensing, computing, positioning, and advanced automation.

WHY IS 6G NEEDED?

The number of connected devices is continuously increasing. Smartphones, smartwatches, vehicles, smart appliances, medical devices, sensors, and robots are becoming part of connected digital ecosystems.

Future applications may require communication capabilities beyond what current networks can efficiently provide. Examples include immersive virtual and augmented reality, autonomous vehicles, remote industrial operations, smart cities with huge numbers of sensors, AI systems, and future holographic or highly immersive communication.

MAJOR FEATURES OF 6G

- Extremely High Data Rates: 6G research targets data rates significantly higher than today's mobile networks. Potential applications include ultra-high-resolution video, immersive XR, real-time 3D communication, advanced cloud gaming, and rapid transfer of large scientific and industrial datasets.
- Ultra-Low Latency: Latency is the delay between sending information and receiving a response. 6G aims to provide extremely low latency for autonomous transportation, industrial robots, remote control systems, interactive XR, and collaborative robotics.
- Artificial Intelligence-Native Networks: AI is expected to be deeply integrated into 6G for network optimization, traffic management, energy management, fault detection, security monitoring, resource allocation, predictive maintenance, and automatic configuration.
- Integrated Sensing and Communication: Integrated Sensing and Communication (ISAC) is a major research area. Network signals could potentially help detect objects, monitor movement, support autonomous vehicles, enable positioning, monitor infrastructure, and assist industrial automation.
- Massive Internet of Things: Future 6G networks could support enormous numbers of sensors and connected devices, including environmental sensors, agricultural sensors, smart-home devices, industrial equipment, wearable devices, and transportation systems.

TECHNOLOGIES EXPECTED TO SUPPORT 6G

Terahertz Communication: 6G research is exploring very high-frequency spectrum, including terahertz frequencies. These frequencies may offer very large bandwidths, but they also face challenges such as propagation loss, atmospheric absorption, shorter transmission distances, and complex hardware requirements.

Reconfigurable Intelligent Surfaces: RIS consists of controllable elements that can manipulate radio-signal propagation. It could potentially improve coverage, redirect signals around obstacles, reduce dead zones, and improve energy efficiency.

Satellite and Non-Terrestrial Networks: Future 6G networks are expected to integrate terrestrial networks with satellites and high-altitude platforms, improving connectivity in rural, remote, ocean, mountain, and disaster-affected areas.

Edge Computing: Edge computing brings processing closer to users and devices. This can support applications requiring fast responses, such as autonomous vehicles and industrial systems.

Digital Twins: A digital twin is a digital representation of a physical object, system, or environment. Future networks could support real-time digital twins of factories, cities, buildings, vehicles, machines, and energy systems.

APPLICATIONS OF 6G

Smart Cities: Smart traffic management, intelligent street lighting, waste management, environmental monitoring, public safety, and smart energy distribution.

Healthcare: Remote health monitoring, connected medical devices, real-time medical information, AI-assisted healthcare, advanced telemedicine, and remote robotic assistance.

Autonomous Vehicles: Vehicle-to-vehicle and vehicle-to-infrastructure communication, real-time traffic information, collision avoidance, autonomous navigation, and intelligent traffic management.

Education: Virtual laboratories, 3D educational environments, augmented-reality lessons, remote practical training, and interactive simulations.

Industry 5.0 and Smart Manufacturing: Collaborative robots, automated production, predictive maintenance, digital twins, real-time quality control, and intelligent supply chains.

Agriculture: Connected sensors can monitor soil, temperature, humidity, crop health, irrigation requirements, and weather conditions, while AI can assist with agricultural decision-making.

Immersive Communication: Future communication may involve extended reality, 3D environments, spatial communication, virtual collaboration, and highly immersive remote experiences.

ADVANTAGES OF 6G

- Much higher data capacity
- Extremely low latency
- Support for massive numbers of devices
- Greater integration of AI
- Improved sensing and positioning
- Integration of terrestrial and satellite networks
- Improved network automation
- Potentially better energy efficiency
- Support for advanced industrial and scientific applications

CHALLENGES OF 6G

Technical Complexity: Higher frequencies and advanced architectures require sophisticated hardware and signal processing.

Infrastructure Cost: New infrastructure could require substantial investment.

Energy Consumption: Supporting huge numbers of devices and advanced AI systems could create significant energy demands.

Security and Privacy: More connected devices and AI systems could create new cybersecurity and privacy risks.

Spectrum Availability: Identifying and managing suitable spectrum is a major international challenge.

Digital Divide: Affordable and accessible connectivity will be important to ensure that advanced technology does not benefit only highly developed regions.

WHEN WILL 6G BECOME AVAILABLE?

6G is still in the research, standardization, and development stage. International organizations and industry groups are working toward defining requirements and technical standards.

DIFFERENCE BETWEEN 5G AND 6G

The ITU's framework for future mobile communication is known as IMT-2030, which provides a global framework for the development of 6G-related systems.

Commercial deployment of 6G is generally expected around 2030 or later, but the exact timeline will depend on international standards, technological development, spectrum availability, testing, investment, and regulatory decisions.

Feature	5G	6G
Generation	Fifth generation	Sixth generation
Main focus	High-speed connectivity and diverse services	Intelligent, integrated connectivity
Data capacity	Very high	Expected to be significantly higher
Latency	Very low	Expected to be even lower
AI integration	Increasing	Expected to be deeply integrated
Sensing	Emerging capability	Expected to be more integrated
Spectrum	Sub-6 GHz and mmWave	Research includes higher-frequency and THz bands
Satellite integration	Developing	Expected to be more deeply integrated
Applications	IoT, smart factories, connected vehicles	AI-native systems, immersive communication, advanced robotics, digital twins, and more

The exact technical specifications of 6G are still being developed and standardized by researchers and international organizations. Therefore, the capabilities and features discussed here should be viewed as research goals, expected advancements, and likely future directions rather than finalized specifications for commercial 6G networks.

CONCLUSION

6G Communication Technology represents a major vision for the future of wireless connectivity. It is expected to move beyond the traditional idea of mobile internet by combining very high-capacity communication with artificial intelligence, sensing, computing, positioning, and automation.

The technology could support a wide range of future applications, including smart cities, autonomous transportation, advanced healthcare, intelligent agriculture, immersive education, smart manufacturing, robotics, and digital twins.

At the same time, 6G faces major challenges involving infrastructure, spectrum, energy use, security, privacy, affordability, and technical complexity. Since the technology is still under development, many of its final specifications and capabilities have not yet been established.

Ultimately, the aim of 6G is to create a digital world that is more intelligent, reliable, sustainable, and seamlessly connected. If it is developed responsibly, 6G could become one of the major technological foundations of the 2030s and beyond.

REFERENCES

- International Telecommunication Union (ITU), IMT-2030 (6G) Framework and Overall Objectives, ITU-R Recommendation M.2160.
- International Telecommunication Union (ITU), IMT-2030: Future Development of International Mobile Telecommunications.
- 3GPP, 3GPP and the Development of Future Mobile Communication Systems.
- European Commission, 6G Research and Innovation / Smart Networks and Services.
- Hexa-X, 6G Vision and Research Programme, European 6G research initiative.
- NGMN Alliance, 6G Position Statement and Research Perspectives.
- IEEE, 6G Wireless Communications: Research Directions and Emerging Technologies.

How Deepfakes Are Created and Detected Using Artificial Intelligence

Abhijith Anil
S3 BSc Computer Science

INTRODUCTION

Artificial Intelligence (AI) has completely changed the way digital images, videos, and audio can be created and edited. One of the most interesting examples of this is the deepfake. A deepfake is digitally generated or manipulated media created using AI, particularly deep learning techniques. It can make a person's face, voice, or actions appear different from what actually happened. While these technologies can be useful in areas such as entertainment and creative media, they can also create serious concerns related to misinformation, privacy, identity, and trust in digital content. Because of this, researchers are working on AI systems that can understand synthetic media and determine whether digital content has been altered.

DEEPFAKE CREATION

Creating a deepfake mainly involves deep learning models that are trained using large amounts of images, videos, or audio. During training, the system learns different patterns, including facial features, expressions, movements, and voice characteristics. Earlier deepfake technologies often relied on autoencoders and similar neural network techniques to manipulate faces.

Generative Adversarial Networks (GANs) later became an important technology for generating realistic synthetic media. A GAN consists of two main parts: a generator that creates artificial content and a discriminator that tries to identify whether the content is real or generated. By repeatedly training against each other, the generator becomes better at producing realistic results. More recently, diffusion-based generative AI models have further improved the quality and flexibility of synthetic media.

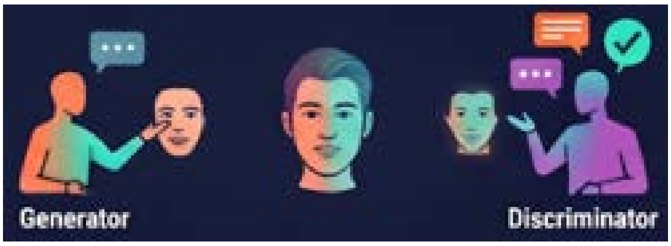

TYPES OF DEEPFAKES

Deepfakes can be created in several ways, depending on the type of manipulation. Face swapping replaces one person's face with another, while face reenactment can alter facial expressions, movements, or other aspects of appearance. These techniques can produce realistic changes that are often difficult to notice.

AI can also generate realistic faces of people who do not actually exist, making it possible to create artificial identities without using a real person's face. Similar techniques can be used to generate speech that imitates a person's tone, pronunciation, rhythm, and other characteristics of their voice. Deepfakes can also combine manipulated visuals with AI-generated audio, making the final content highly realistic and convincing. This makes modern deepfakes harder to identify than older versions that often contained obvious visual mistakes. Detection systems may need to examine facial features, movements, lighting, audio patterns, and the synchronization between speech and lip movements. Analysing these clues together can provide a more reliable way of determining whether digital content is genuine or manipulated.

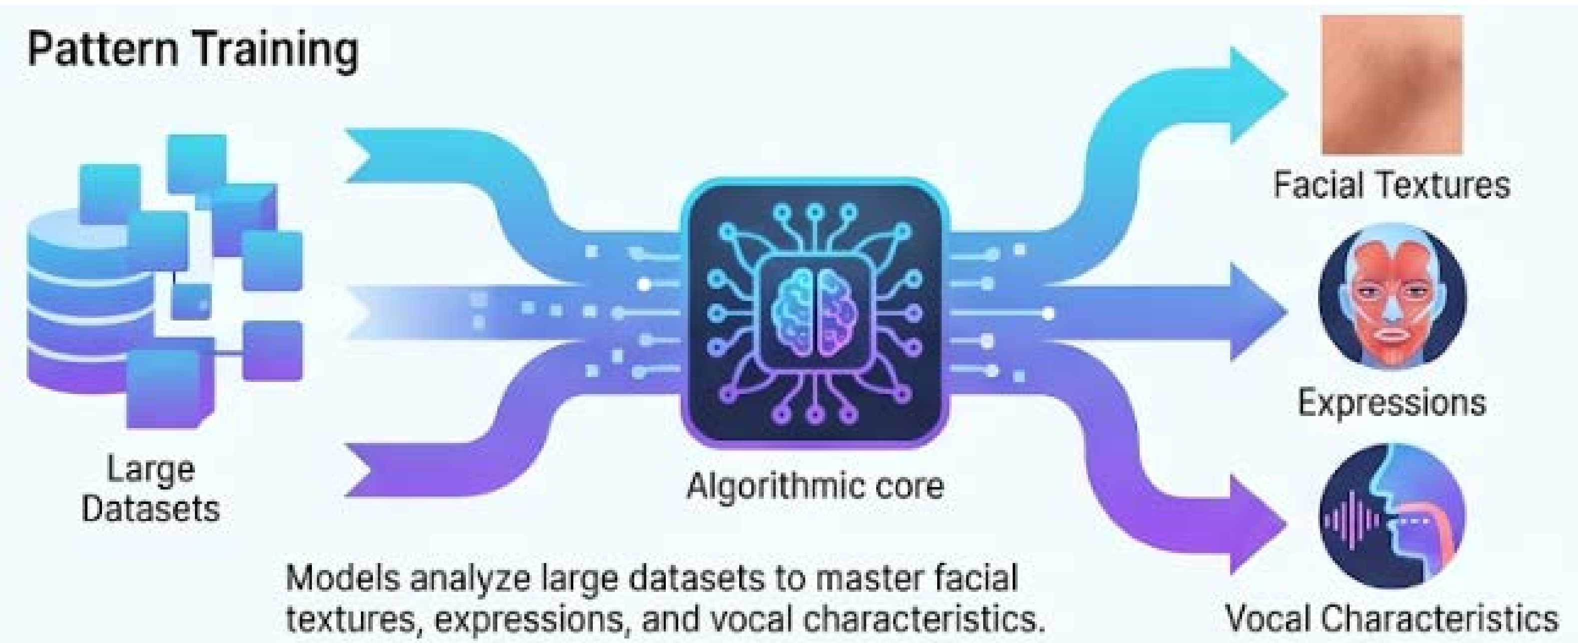

DEEPFAKE DETECTION

becomes better at producing realistic results. More recently, diffusion-based generative AI models have further improved the quality and flexibility of synthetic media.

Technologies such as Convolutional Neural Networks (CNNs), vision transformers, and other machine-learning models can identify subtle patterns linked to manipulation.

These patterns may include unusual facial textures, unnatural-looking lighting, blending artifacts, compression patterns, or other subtle details that can be difficult for people to notice. The FaceForensics++ research introduced a large benchmark of manipulated facial images, showing how automated detection systems could be tested and evaluated against a variety of different manipulation techniques.

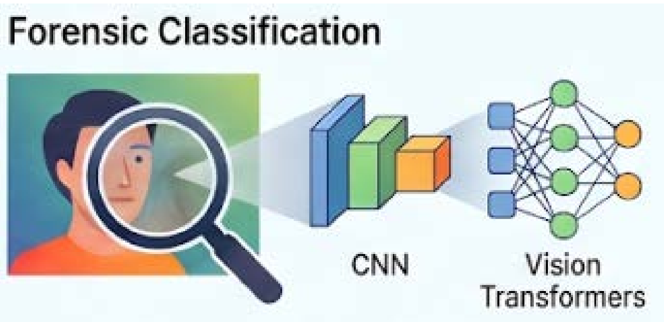

VIDEO AND AUDIO-VISUAL DETECTION

When dealing with videos, detection systems can also examine how information changes over time. Rather than analysing only one frame, an AI system can study how facial features, expressions, and movements develop from one frame to the next

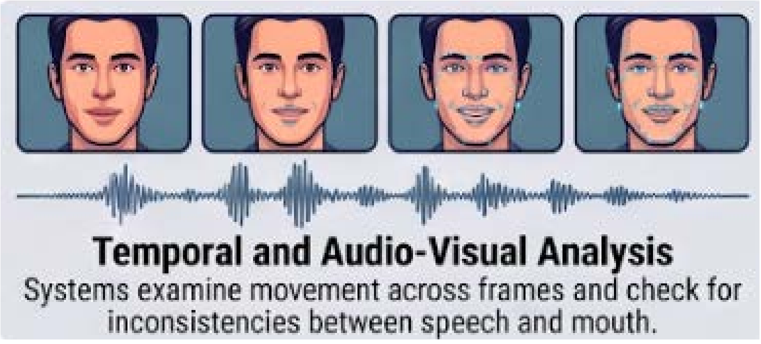

Audio and video can also be analysed together by comparing a person's speech with their mouth movements, facial expressions, and other visual information. If these elements do not match naturally, the mismatch can indicate that the content has been manipulated or artificially generated. This is becoming increasingly important as modern deepfakes become more advanced, realistic, and capable of combining multiple forms of media, making them harder to detect through normal observation.

CHALLENGES IN DEEPFAKE DETECTION

However, detecting deepfakes is far from easy. As generation techniques continue to improve, the obvious visual mistakes found in older deepfakes are becoming less common. A detector that performs well on one type of manipulation may also struggle when it encounters a completely new technique that was not included in its training data. This is known as the generalization problem.

To address this, researchers have developed large benchmark datasets and standardized testing methods. DeepfakeBench, for example, provides a unified framework for comparing different detection techniques, datasets, and evaluation methods.

Real-world conditions make the problem even harder. Social media platforms and messaging applications often resize, compress, or re-encode images and videos before they reach other users. These changes can remove or weaken some of the forensic clues that detection systems depend on, making manipulated content more difficult to identify accurately.

NIST's current work on deepfake forensics highlights the difference between how detectors perform in controlled academic experiments and how well they work in real-world situations. This shows why realistic and challenging benchmarks are important when evaluating deepfake detection systems.

IMPORTANCE OF DEEPFAKE DETECTION

Deepfake detection matters because digital media is now widely used as a source of information and, in some situations, as evidence. Manipulated content can damage a person's reputation, invade privacy, spread false information, and reduce people's trust in what they see online.

Because of these risks, researchers are exploring approaches that combine AI-based detection with other techniques such as digital forensics, provenance information, watermarking, and content-authenticity systems. It is also important not to treat a single AI detector as a completely reliable authority. In many cases, confirming whether digital content is genuine may require several independent sources of evidence.

CONCLUSION

Deepfakes show both the impressive capabilities and the limitations of modern computer science. Deep learning models can now create highly realistic synthetic images, videos, and audio, while other AI systems are being developed to identify the subtle traces left behind by manipulation.

This has created an ongoing technological race between systems that generate synthetic media and systems that try to detect it. As generative AI continues to develop, deepfake detection will remain an important area of computer science, cybersecurity, and digital forensics. Building stronger detection techniques, realistic testing environments, and reliable methods for proving digital authenticity will be essential for maintaining trust in a world where AI-generated content is becoming increasingly common.

REFERENCES

- Rössler, A., Cozzolino, D., Verdoliva, L., Riess, C., Thies, J., & Nießner, M. (2019). FaceForensics++: Learning to Detect Manipulated Facial Images. IEEE/CVF International Conference on Computer Vision (ICCV).
- Yan, Z., Zhang, Y., Yuan, X., Lyu, S., & Wu, B. (2023). DeepfakeBench: A Comprehensive Benchmark of Deepfake Detection. Advances in Neural Information Processing Systems (NeurIPS).
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., et al. (2014). Generative Adversarial Nets. Advances in Neural Information Processing Systems, 27.
- National Institute of Standards and Technology (NIST). (2026). GenAI: Deepfakes 2026: Deepfake Forensics.

Physical AI and Robotics : A Comparative Study of NVIDIA Isaac, Google DeepMind Gemini Robotics, and Boston Dynamics

V Sai Sesh
S5 BSc Computer Science

INTRODUCTION

Robotics research is going through a major transition that is similar in some ways to the transformation caused by large language models in natural-language processing. Rather than relying entirely on hand-written controllers designed for specific tasks, researchers are increasingly training general-purpose "vision-language-action" (VLA) policies using large and varied collections of demonstrations, simulated experiences, and large-scale video data. These models can then be deployed on physical robots. This approach, commonly referred to as physical AI or embodied AI, aims to bring perception, reasoning, and action into one learned system instead of treating them as completely separate engineering problems.

Three organizations demonstrate different parts of this broader shift. NVIDIA focuses on computing, simulation, development tools, and an open robot foundation model, effectively providing infrastructure that other robot developers can build on. Google DeepMind focuses on advanced multimodal reasoning models adapted for physical environments, with particular attention to planning, tool use, and transferring capabilities between different robot platforms. Boston Dynamics brings decades of experience in mechanical design and robot control, while increasingly combining that expertise with learned perception and planning systems. Its Atlas humanoid is also moving from laboratory demonstrations toward real factory applications. This paper examines these contributions and the way they are beginning to connect.

NVIDIA ISAAC AND THE GROOT PLATFORM

NVIDIA's robotics work is built around the Isaac platform and Isaac GROOT. Isaac provides an ecosystem of simulation, perception, and manipulation tools supported by CUDA-accelerated libraries, while GROOT is NVIDIA's effort to create an open, general-purpose foundation model for humanoid robots. Rather than being limited to one model, GROOT is presented as an end-to-end development pipeline. It covers the collection of synthetic and human demonstration data through Isaac Teleop and Isaac Lab, physics-based simulation and evaluation through Isaac Sim, and deployment on NVIDIA's Jetson Thor edge computer for real-time inference

The GROOT N vision-language-action model is trained using both real human demonstrations and simulated trajectories. More recent versions combine the model with a smaller reasoning component intended to improve long-horizon task decomposition and make it easier to transfer skills between different robot bodies. NVIDIA has also introduced an open humanoid reference design that combines the GROOT software stack with hardware from third-party manufacturers, including a commercially available bipedal platform and dexterous hands. This gives universities and industry research teams a way to experiment without having to build a complete robot from the ground up. Organizations working in areas such as manipulation, logistics, and mobile robotics have adopted parts of the GROOT ecosystem, while NVIDIA's simulation-ready datasets have reportedly reached millions of downloads. These developments reinforce NVIDIA's position primarily as an infrastructure and tooling provider rather than a traditional robot manufacturer.

GOOGLE DEEPMIND GEMINI ROBOTICS

Google DeepMind has extended its Gemini family of multimodal models into robotics through two closely related model lines. Gemini Robotics-ER, or embodied reasoning, is a vision-language model designed for spatial understanding, task planning, and detecting whether an action has succeeded. It can also use external tools, including search and humanoid reference, allowing it to operate as a high-level "strategist" that determines the next step for a robot. Gemini Robotics serves as the corresponding vision-language-action model. It takes visual information and instructions and translates them into motor actions, effectively acting as the "executor".

Later releases have expanded this two-model approach. Gemini Robotics 1.5 added explicit step-by-step reasoning before actions are performed, making the robot's intermediate plan easier to inspect. It also introduced cross-embodiment learning, allowing skills developed on one robot to be transferred to another [5]. A later generation introduced what Google describes as "whole-body intelligence," extending control beyond individual manipulators to full humanoid systems. This includes coordination of the feet, torso, and dexterous hands. Using motion-transfer techniques from earlier work, the system can also adapt to new bi-arm robot platforms with only a few hundred examples of task-specific data [6]. Updates to the embodied-reasoning model improved spatial reasoning, video-based monitoring of task progress, and coordination between multiple robots.

A distilled on-device version also makes VLA capabilities available on more limited local hardware, which can be useful when low latency is important [13]. Google DeepMind has additionally entered a direct hardware partnership with Boston Dynamics to develop embodied AI for Atlas, which is discussed in the following section.

BOSTON DYNAMICS ATLAS

Boston Dynamics, which is majority-owned by Hyundai Motor Group, has been developing humanoid robots since the original hydraulic Atlas was introduced as part of a DARPA research project. At CES 2026, the company presented a completely redesigned, all-electric version of Atlas aimed at commercial use. The move from hydraulic actuation to an electric drivetrain is intended to improve efficiency, reduce noise, and make maintenance simpler. The new Atlas has 56 degrees of freedom, a reach of approximately 7.5 feet, and the ability to lift loads of up to 110 pounds for short periods. It also includes autonomous battery swapping designed to support continuous operation.

Unlike earlier prototypes, the redesigned Atlas is aimed at enterprise use in material handling, order fulfillment, and other industrial tasks. It can work autonomously, through teleoperation, or via a tablet interface, and can be integrated into factories using Boston Dynamics' Orbit software. Reports said the entire 2026 production run was committed before its general release, with early deployments planned for Hyundai and Google DeepMind. Later in 2026, Boston Dynamics simplified Atlas's control software and demonstrated self-supervised abilities, including complex kicking motions without motion-capture scripting, showing continued progress in learned whole-body control.

COMPARATIVE DISCUSSION

The three efforts occupy different positions in the physical AI stack, although those positions are becoming increasingly connected. NVIDIA mainly provides horizontal infrastructure through simulation, synthetic data generation, and an open foundation model that can be adopted by different robot manufacturers.

This places NVIDIA more in the role of an ecosystem supplier than a conventional robot vendor. Google DeepMind focuses on advanced reasoning and control models, with an emphasis on generalization, tool use, and fast adaptation to new robot bodies. Its technology is now being integrated directly with Boston Dynamics hardware. Boston Dynamics, meanwhile, brings decades of experience in mechanical engineering and low-level robot control and is the closest of the three to large-scale factory deployment. At the same time, it is increasingly using external foundation models for higher-level intelligence instead of developing the entire cognitive stack internally.

This reflects the wider AI industry, where different companies specialize in different parts of the technology stack. Future advantages in physical AI may depend more on the quality of data, advanced simulation, and effective skill transfer than on hardware alone. All three efforts also face important challenges, such as collecting high-quality training data, safely transferring learned skills, and ensuring reliable operation and safety when robots work around people.

CONCLUSION

Physical AI is moving quickly beyond isolated laboratory demonstrations and toward integrated systems intended for real commercial environments. NVIDIA provides simulation and open-model infrastructure that can support a broad range of robot developers. Google DeepMind contributes advanced reasoning and action models designed to generalize across tasks and robot platforms. Boston Dynamics provides a mature humanoid platform that is now being connected with these AI capabilities. The direct partnership between Boston Dynamics and Google DeepMind is a clear example of this convergence. It suggests that future progress in embodied AI will depend not only on improvements within individual technologies, but also on how effectively the different layers of the physical AI ecosystem work together.

REFERENCES

- NVIDIA Corp., "NVIDIA Announces NVIDIA Isaac GR00T Reference Humanoid Robot for Academic Research," NVIDIA Newsroom, Jun. 2026.
- E. Llontop and B. Neel, "Develop Humanoid Robot Policies End-to-End with NVIDIA Isaac GR00T," NVIDIA Technical Blog, Jul. 2026.
- NVIDIA Corp., "NVIDIA Announces Project GR00T Foundation Model for Humanoid Robots and Major Isaac Robotics Platform Update," NVIDIA Newsroom, Mar. 2024.
- "NVIDIA Unveils Full-Stack Robotics Platform at GTC 2026," Blockchain.News, Mar. 2026.
- Google DeepMind, "Gemini Robotics 1.5 Brings AI Agents into the Physical World," DeepMind Blog, Sep. 2025.
- Google DeepMind, "Gemini Robotics 2 Brings Whole Body Intelligence to Robots," DeepMind Blog, 2026.
- "Google DeepMind Releases Gemini Robotics-ER 1.6: Bringing Enhanced Embodied Reasoning and Instrument Reading to Physical AI," MarkTechPost, Apr. 2026.

Brain-Inspired AI and Cognitive Robotics

Praveena M P
S5 BSc Computer Science

ABSTRACT

Most AI systems learn from huge datasets and consume large amounts of power, yet still struggle with things brains do easily: learning continuously, adapting quickly, and running on little energy. Brain-inspired AI closes this gap by copying real neuroscience mechanisms rather than the loose idea of a "neural network." Cognitive robotics extends this into physical robots whose sensing, thinking, and movement are modeled on brain-body coordination. This article explains these ideas simply, gives real examples, and reviews the main open challenges.

Keywords: brain-inspired AI, cognitive robotics, spiking neural networks, neuromorphic computing, embodied cognition

INTRODUCTION

Most AI is trained using backpropagation on massive datasets — very different from how brains learn. A human brain runs on about 20 watts, yet learns new skills from a few examples and adapts instantly without forgetting old ones. AI systems typically need far more data and energy, and often forget old knowledge when learning something new.

Brain-inspired AI borrows directly from neuroscience: how neurons signal, how the brain predicts what it expects to sense, and how it learns from reward. Applied to physical robots that must see, move, and react in real time, this becomes cognitive robotics. Robots face constraints a data-center computer does not so researchers look to biology, which already solves these problems continuously.

KEY IDEAS

Spiking Neural Networks

Ordinary AI neurons send constant, smooth signals. Real neurons stay silent and only “fire” a brief spike once enough input builds up — the basis of spiking neural networks (SNNs). Since SNNs communicate only when a spike occurs, they are naturally energy-efficient, like a doorbell that draws power only when pressed — ideal for battery-powered robots. The trade-off is that spikes are harder to train, since standard backpropagation doesn't apply directly; researchers use adapted training methods or biological learning rules, such as strengthening a connection only when one neuron reliably fires just before another.

Neuromorphic Chips

SNNs reach full efficiency only on chips built for them, such as Intel's Loihi or IBM's TrueNorth, which activate only where a spike passes through — similar to how only relevant brain regions activate when you catch a ball. This lets a robot process sensing and adjust movement instantly on its own body instead of waiting on a distant server, which matters for tasks like balancing or surgical precision.

Cognitive Architectures

Cognitive architectures model the brain at a higher level — how memory, attention, and decision-making are organized. SOAR breaks difficult tasks into sub-problems when stuck, like a person pausing to plan when facing something unfamiliar. ACT-R mimics brain regions directly, with separate modules for memory, skill, vision, and movement that communicate through limited channels, mirroring the brain's own bottlenecks. In robots, such architectures typically handle slow, deliberate planning, while spiking or simpler networks handle fast, reflexive movement.

PUTTING IT TOGETHER IN A ROBOT

Brains constantly predict what they expect to sense and react strongly when reality doesn't match. Robots copy this with “forward models” that predict the sensory result of a planned movement, allowing near-instant correction. For example, tightening a grip the moment an object starts to slip.

The basal ganglia are believed to signal when an outcome is better or worse than expected.

in AI, where a robot tries actions and improves from reward or penalty.This overlap, found independently in neuroscience and AI, is widely used to teach robots to walk or grasp objects through trial and error.

AI systems often “forget” old skills when learning new ones, a problem known as catastrophic forgetting. The brain helps avoid this through two memory systems: a fast hippocampus-like store for new experiences, and a slower neocortex-like store that gradually absorbs important patterns, including during sleep. Robots can mimic this process through “memory replay,” periodically revisiting past experiences to retain older skills while continuing to learn and adapt to new tasks.

REAL-WORLD EXAMPLES

System	Brain-Inspired Mechanism	Application
iCub	Infant-like developmental learning	Humanoid learns object recognition and movement
Loihi robots	Spiking, low-power chips	Real-time, low-power arm and mobile control
SpiNNaker robots	Large-scale spiking simulation	Insect-like visual navigation
RL-based robots	Basal ganglia-like reward learning	Walking/grasping via trial and error

REMAINING CHALLENGES

- **Training difficulty**: spiking networks are still harder to train accurately than standard AI networks.
- **Limited hardware**: neuromorphic chips remain small-scale for rich vision or hearing tasks.
- **Sim-to-real gap**: robots trained in simulation often behave differently in the real world.
- **Fast-slow integration**: combining reflexive movement with deliberate planning remains unsolved.
- **No standard benchmark**: there is no widely accepted way to measure how “brain-like” a system is.

CONCLUSION

Brain-inspired AI and cognitive robotics combine neuroscience and engineering to build machines that use less energy, adapt faster, and learn continuously. No single mechanism solves this alone, progress comes from combining efficient spiking computation, structured cognitive architectures, predictive processing, and reward-based learning into layered systems, mirroring the brain's own reliance on many specialized systems working together.

REFERENCES

- Anderson, J. R. (2007). How Can the Human Mind Occur in the Physical Universe? Oxford University Press.
- Davies, M., et al. (2018). Loihi: A neuromorphic manycore processor with on-chip learning. IEEE Micro, 38(1), 82–99.
- Laird, J. E. (2012). The Soar Cognitive Architecture. MIT Press.
- Maass, W. (1997). Networks of spiking neurons. Neural Networks, 10(9), 1659–1671.
- Metta, G., et al. (2010). The iCub humanoid robot. Neural Networks, 23(8–9), 1125–1134.
- Rao, R. P. N., & Ballard, D. H. (1999). Predictive coding in the visual cortex. Nature Neuroscience, 2(1), 79–87.

Measurable Sustainable Technology and ESG-Driven IT Strategy : *A Metrics-Based Framework for Green Digital Transformation*

Ms. Anila Aneesh
Asst. Prof.,
Dept of Computer Science

ABSTRACT

Environmental, Social, and Governance (ESG) mandates are pushing Information Technology (IT) organizations beyond qualitative sustainability statements toward quantifiable, auditable performance. This paper proposes a metrics-based framework that links measurable sustainable technology (Green IT) indicators—energy efficiency, carbon intensity, water usage, e-waste diversion, and circularity—to enterprise ESG reporting obligations such as GRI, SASB, and TCFD/ISSB. We define a layered architecture spanning data collection, normalization, and governance, and propose a composite IT Sustainability Index (ITSI) that aggregates operational and embodied-impact metrics. A qualitative case discussion illustrates how the framework supports cloud, data-center, and hardware lifecycle decisions. The paper concludes with implementation guidance and open research challenges, including data quality, Scope 3 attribution, and greenwashing risk in IT-specific ESG claims.

INTRODUCTION

iDigital infrastructure now accounts for a measurable and growing share of global electricity use and greenhouse-gas (GHG) emissions, placing Information Technology (IT) squarely inside corporate Environmental, Social, and Governance (ESG) agendas. Investors, regulators, and customers increasingly expect ESG disclosures to be quantitative, comparable, and auditable rather than narrative. At the same time, IT leaders are expected to demonstrate that "sustainable technology" initiatives—cloud migration, hardware refresh cycles, data-center consolidation—produce measurable environmental and governance outcomes, not only cost savings.

This creates a gap: sustainability frameworks (GRI, SASB, TCFD/ISSB) are written at the enterprise level, while IT operational metrics (e.g., Power Usage Effectiveness, server utilization) are written at the infrastructure level. Few organizations have a structured bridge between the two. This paper addresses that gap by proposing a metrics-based, ESG-aligned IT strategy framework that inventories measurable sustainable-technology indicators, maps them to recognized ESG disclosure standards, and aggregates them into a composite index usable in IT governance and board-level reporting.

BACKGROUND AND RELATED CONCEPTS

ESG-driven IT strategy sits at the intersection of three established bodies of practice. First, corporate carbon accounting, formalized by the GHG Protocol, distinguishes Scope 1 (direct), Scope 2 (purchased energy), and Scope 3 (value-chain) emissions—IT is predominantly a Scope 2 and Scope 3 contributor through electricity purchases and hardware manufacturing/disposal. Second, data-center and infrastructure efficiency standards, notably ISO/IEC 30134 and The Green Grid's Power Usage Effectiveness (PUE) metric, provide operational-level measurement conventions. Third, enterprise sustainability disclosure standards such as GRI 302 (Energy), the SASB Software & IT Services standard, and TCFD/ISSB climate-risk disclosure define what enterprises must report externally.

Existing Green IT literature has largely focused on individual levers —virtualization, renewable-powered data centers, e-waste recycling— without a unifying measurement layer that connects operational telemetry to ESG line items.

Our framework treats measurement, rather than technology choice, as the central organizing principle for evaluating and guiding the overall approach.

MEASURABLE SUSTAINABLE TECHNOLOGY INDICATORS

Energy and Carbon Metrics

Power Usage Effectiveness (PUE = Total Facility Energy / IT Equipment Energy) remains the primary infrastructure efficiency metric [3]. Carbon intensity of compute (kg CO_2e per kWh, weighted by grid mix and Power Purchase Agreements) converts energy metrics into GHG Protocol Scope 2 figures, reportable under GRI 305.

Resource and Circularity Metrics

Water Usage Effectiveness (WUE), hardware utilization rate, average server refresh cycle length, and e-waste diversion rate (percentage of retired assets reused, resold, or responsibly recycled versus landfilled) quantify circular-economy performance, aligned with emerging circularity disclosure expectations under ISSB and EU CSRD-adjacent frameworks.

Governance and Data-Quality Metrics

Because ESG figures face increasing audit scrutiny, governance metrics—percentage of IT assets with automated telemetry, data completeness rate, and third-party assurance coverage—determine whether the above metrics are decision- and disclosure-grade rather than estimated.

Composite IT Sustainability Index (ITSI)

We propose a normalized composite score:

ITSI = w1·N(PUE) + w2·N(CI) + w3·N(CircR) + w4·N(DataQ) where N(·) normalizes each indicator to a 0–1 scale against sector benchmarks, CI is carbon intensity, CircR is the circularity rate, DataQ is data-quality/assurance coverage, and weights w1..w4 are set by governance policy. ITSI enables trend tracking and board-level comparison without exposing raw operational data externally.

ESG-DRIVEN IT STRATEGY FRAMEWORK

The framework operates in four layers, executed as a continuous cycle rather than a one-time project.

Collect

Instrument data centers, cloud accounts (via provider carbon/sustainability APIs), and asset-management systems to capture energy, utilization, and lifecycle data at source, minimizing manual estimation.

Normalize and Aggregate
Convert heterogeneous telemetry into the standard units defined in Section III, apply grid-emission factors, and compute ITSI at the workload, business-unit, and enterprise level.

Govern
Assign accountability (e.g., a Green IT/ESG steering committee), set targets against science-based or sector benchmarks, and define assurance procedures so metrics can withstand external audit.

Disclose and Decide
Feed validated metrics into enterprise ESG reports (GRI/SASB/TCFD-ISSB) and into IT decisions—procurement, cloud-region selection, refresh-cycle policy—closing the loop between measurement and action.

ILLUSTRATIVE APPLICATION

Consider an enterprise consolidating three legacy data centers into a hyperscale cloud region selected partly on published provider carbon intensity. Applying the framework: PUE and carbon-intensity metrics quantify the operational efficiency gain; e-waste diversion metrics track the decommissioned hardware; data-quality metrics confirm cloud-provider telemetry meets assurance requirements; and the resulting ITSI improvement is reported as a quantified contribution to the enterprise's GRI 305 energy-emissions disclosure. This demonstrates how a single migration decision produces auditable ESG evidence rather than a qualitative sustainability claim.

CHALLENGES AND LIMITATIONS

Three challenges limit measurable ESG-IT alignment today. First, Scope 3 attribution for hardware manufacturing and cloud supply chains relies on industry-average emission factors with limited granularity. Second, cross-provider comparability is weak because cloud vendors report carbon and efficiency data using differing methodologies. Third, the absence of independent assurance for many IT sustainability metrics creates greenwashing risk, where reported improvements are not verifiable. Addressing these gaps requires standardized, provider-agnostic telemetry APIs and third-party assurance schemes analogous to financial audit.

CONCLUSION

Measurable sustainable technology is the operational foundation on which credible ESG-driven IT strategy must be built. By mapping infrastructure-level metrics to enterprise disclosure standards and aggregating them into a composite index, organizations can move from aspirational Green IT statements to auditable, decision-relevant sustainability performance. Future work includes empirical validation of the ITSI weighting scheme across industry sectors and integration with automated ESG assurance tooling.

REFERENCES

- World Resources Institute and World Business Council for Sustainable Development, "The Greenhouse Gas Protocol: A Corporate Accounting and Reporting Standard" .
- ISO/IEC, Information technology Key performance indicators — Part 2: Power usage effectiveness (PUE), International Organization for Standardization, Geneva, Switzerland.
- The Green Grid, "PUE: A Comprehensive Examination of the Metric," White Paper #49, The Green Grid Association.
- Global Reporting Initiative, "GRI 302: Energy 2016" and "GRI 305: Emissions 2016," GRI Standards, Amsterdam, The Netherlands.

Cloud 3.0 : The Next Evolution of Cloud Architecture

Ms. Mubeena J
Asst. Prof.,
Dept of Computer Science

ABSTRACT

Cloud computing has evolved significantly over the past two decades. It began with virtualization-focused infrastructure, commonly referred to as Cloud 1.0, and later moved toward elastic, container-based, and API-driven platforms associated with Cloud 2.0. Today, a new stage of development is emerging, referred to in this paper as Cloud 3.0. This model combines artificial intelligence, edge computing, autonomous infrastructure, and sustainability-focused resource management. Cloud 3.0 aims to create systems that can make operational decisions with less human involvement while responding to changing workloads, locations, and environmental conditions. This paper examines the development of cloud architecture, explains the major characteristics of Cloud 3.0, presents a layered

reference architecture, and discusses the key challenges that may affect its adoption. The goal is to provide researchers and practitioners with a clear framework for understanding the technologies and ideas shaping the next generation of cloud computing.

Index Terms: cloud computing, edge computing, distributed systems, artificial intelligence, autonomous infrastructure, serverless computing, sustainable computing.

INTRODUCTION

Cloud computing has passed through several major stages of architectural development over the last two decades. The first stage, known as Cloud 1.0, was mainly built around server virtualization and the transition from privately owned

data centers to on-demand computing and storage services. The second stage, Cloud 2.0, introduced technologies such as containers, microservices, elastic scaling, and API-based architectures. These developments made it easier to scale applications and support faster software delivery through DevOps practices.

A third stage is now beginning to take shape, which this paper refers to as Cloud 3.0. It is driven by the growing combination of artificial intelligence, edge computing, distributed infrastructure, and increasing pressure to reduce energy consumption and carbon emissions. As a result, cloud infrastructure is becoming more intelligent, decentralized, and capable of responding to its surroundings and workload conditions without constant human intervention.

The purpose of this paper is to describe Cloud 3.0 as a new stage in cloud architecture, identify its major technical characteristics, and present a layered reference model that can help guide future cloud system development and research.

EVOLUTION OF CLOUD ARCHITECTURE

Cloud 1.0 ,Virtualized Infrastructure

The early generation of cloud platforms focused on abstracting physical hardware through virtualization. Hypervisors made it possible to divide physical resources into virtual machines, supporting the delivery of Infrastructure as a Service. During this period, scaling was generally handled manually, and applications were closely tied to individual virtual machine instances.

Cloud 2.0 , Elastic, Service-Oriented Cloud

The second generation of cloud computing introduced containers, orchestration platforms such as Kubernetes, and microservice-based application architectures. These technologies made applications more modular and allowed individual services to scale according to demand. Platform as a Service offerings and serverless functions also reduced the amount of infrastructure that developers needed to manage directly. Together, these changes improved

resource utilization and made cloud application development faster and more flexible.

Cloud 3.0 ,Autonomous, Distributed, Intelligent Cloud

Cloud 3.0 builds on the capabilities of earlier cloud generations while adding a greater level of intelligence and automation. Machine learning can be integrated into cloud operations to support decisions related to workload placement, scaling, monitoring, and system recovery. At the same time, computing resources are increasingly distributed across traditional data centers, regional infrastructure, and edge devices.

Another important aspect of Cloud 3.0 is its focus on energy efficiency and sustainability. Cloud systems can consider factors such as energy consumption and carbon intensity when deciding where and when workloads should run. This represents a shift from simply optimizing performance and cost toward balancing performance, reliability, cost, and environmental impact.

EVOLUTION OF CLOUD ARCHITECTURE

AI-Native Operations: Machine learning models embedded in the control plane perform predictive scaling, anomaly detection, and self-healing with minimal human intervention.

Edge-to-Core Continuum: Compute and storage are distributed across core data centers, regional nodes, and edge devices, reducing latency for real-time applications such as IoT and AR/VR.

Serverless and Function-Centric Design: Event-driven, stateless execution models mature into the default abstraction for application deployment, abstracting servers entirely from developers.

Distributed and Multi-Cloud Fabric: Workloads are portable across providers through open standards, reducing vendor lock-in and improving resilience.

Sustainability-Aware Scheduling: Workload placement incorporates carbon intensity and energy cost signals to minimize environmental impact.

Zero-Trust and Quantum-Resilient Security: Identity-centric security models and post-quantum cryptography are integrated at the infrastructure level to address evolving threat landscapes.

PROPOSED REFERENCE ARCHITECTURE

A possible reference architecture for Cloud 3.0 can be organized into four major planes.

The Physical and Edge Plane provides the underlying computing resources across central data centers, regional infrastructure, and edge locations. These resources may include different types of processors, storage systems, and connected devices.

The Orchestration Plane manages applications and workloads running through containers and serverless platforms. AI-assisted scheduling can be used to decide where workloads should run and how resources should be allocated according to current conditions.

The Intelligence Plane contains the machine learning and artificial intelligence components responsible for tasks such as workload forecasting, anomaly detection, performance optimization, and automated remediation. This layer allows the infrastructure to learn from operational data and adjust its behavior over time.

The Governance Plane manages policies related to security, compliance, cost, and sustainability. It provides a common layer of control across the underlying infrastructure and ensures that automated decisions remain within organizational and regulatory requirements.

Separating these responsibilities into different layers makes it possible to introduce autonomous capabilities gradually. Organizations can adopt intelligent management features without having to completely replace their existing service-oriented cloud infrastructure.

CHALLENGES AND OPEN ISSUES

Interoperability: Cloud and edge environments often use different APIs, platforms, and data formats. Creating consistent ways for these systems to communicate remains a major challenge, particularly in multi-cloud environments.

Trust in Autonomous Systems: As more infrastructure decisions are made by artificial intelligence, organizations need to understand why those decisions are being made. Explainability, transparency, and reliable control mechanisms will become increasingly important when autonomous systems are responsible for critical workloads.

Sustainability Measurement: There is still a lack of consistent standards for measuring carbon emissions and energy usage across different cloud providers. This makes it difficult to accurately compare the environmental impact of workloads running in different environments.

Evolving Security Requirements: Cloud security must continue to adapt to new threats, including those that could emerge from quantum computing. At the same time, security mechanisms need to maintain the low latency required by edge applications. Finding the right balance between strong security and high performance remains an open problem.

CONCLUSION

Cloud 3.0 represents a significant change in the way cloud infrastructure is designed and managed. Earlier cloud generations focused mainly on virtualization, scalability, and service-oriented architectures. The emerging model moves further toward infrastructure that is autonomous, distributed, intelligent, and environmentally conscious.

By placing artificial intelligence directly within the infrastructure control plane and extending computing capabilities from centralized data centers to the network edge, Cloud 3.0 can potentially provide faster responses, greater resilience, better resource utilization, and improved energy efficiency.

However, turning this concept into a widely adopted architecture will require further work. Interoperability between different cloud and edge platforms, explainable autonomous by edge applications. Finding the right balance between strong security and high performance remains an open problem.

Addressing these issues will be essential for building cloud systems that are not only more intelligent, but also reliable, secure, and sustainable.

REFERENCES

- M. Armbrust et al., "A view of cloud computing," Commun. ACM, vol. 53, no. 4, pp. 50-58, 2010.
- L. M. Vaquero and L. Rodero-Merino, "Finding your way in the fog: Towards a comprehensive definition of fog computing," ACM SIGCOMM Comput. Commun. Rev., vol. 44, no. 5, pp. 27-32, 2014.
- R. Buyya et al., "A manifesto for future generation cloud computing: Research directions for the next decade," ACM Comput. Surv., vol. 51, no. 5, pp. 1-38, 2019.
- S. Dustdar, Y. Guo, B. Satzger, and H.-L. Truong, "Principles of elastic processes," IEEE Internet Comput., vol. 15
- W. Shi, J. Cao, Q. Zhang, Y. Li, and L. Xu, "Edge computing: Vision and challenges," IEEE Internet Things J., vol. 3
- A. N. Toosi, R. N. Calheiros, and R. Buyya, "Interconnected cloud computing environments: Challenges, taxonomy, and survey," ACM Comput. Surv., vol. 47, no. 1, pp. 1-47, 2014.
- A. Radovanovic et al., "Carbon-aware computing for datacenters," IEEE Trans. Power Syst., vol. 38, no. 2, pp. 1270-1280, 2023.
- National Institute of Standards and Technology, "Post-quantum cryptography standardization," NIST, Gaithersburg, MD, USA, Rep. NIST IR 8413, 2022.

ARPPO 2.0

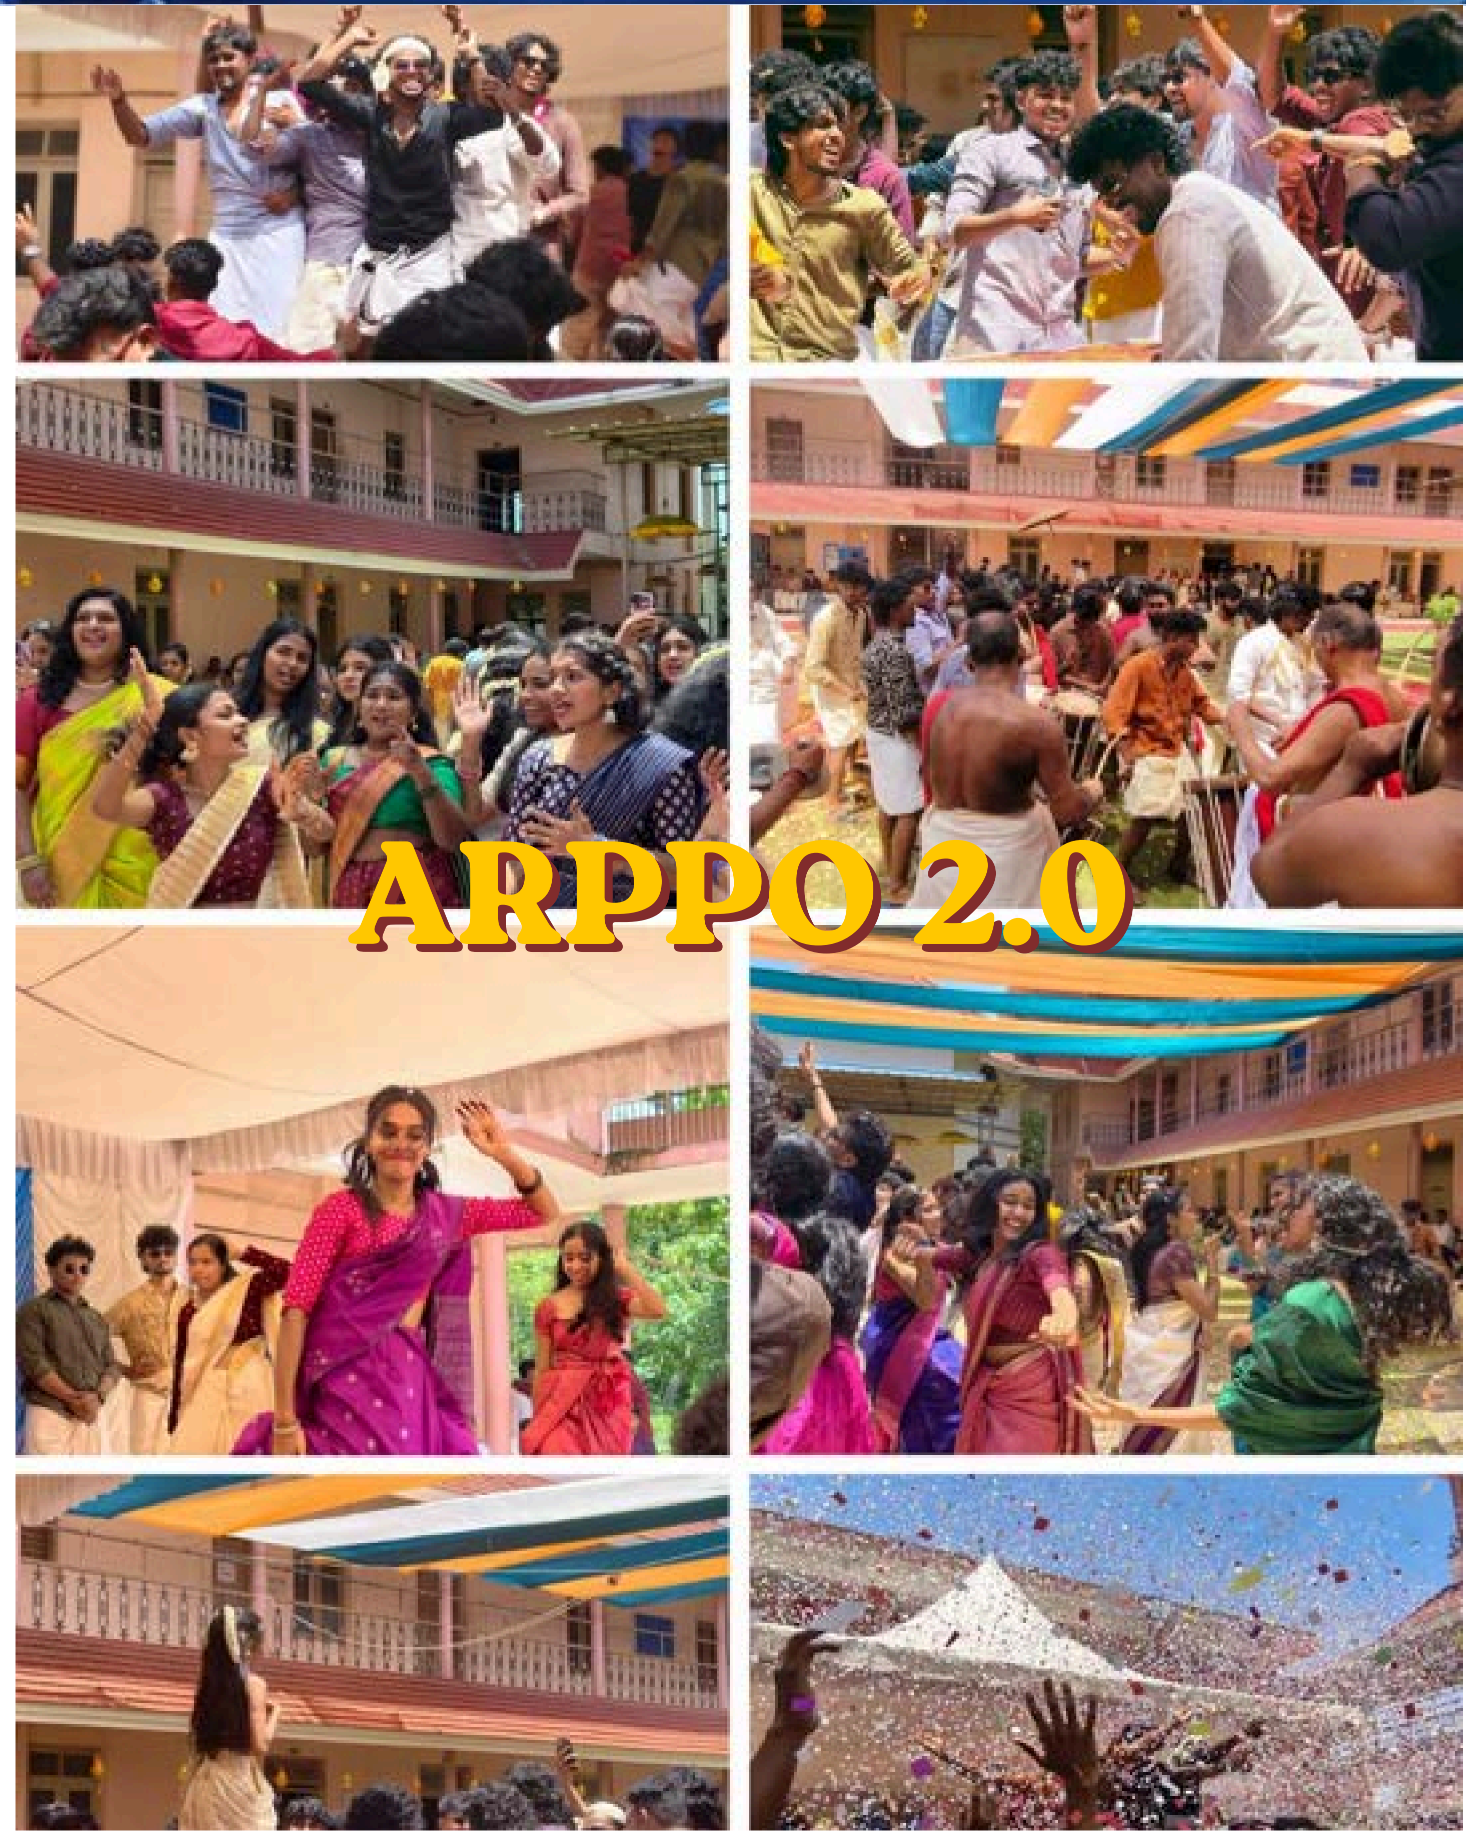

The college celebrated the spirit of Onam through a series of festive events, culminating in the grand Onam Celebration held on August 20.


10
24

TECH NEWS

GOOGLE BRINGS GEMINI ENTERPRISE TO LAWYERS WITH NEW AI TOOLS, AUTOMATION PLUG-INS

Google has expanded its Gemini Enterprise platform into the legal sector with the launch of Gemini Enterprise for Legal, a specialised set of AI tools designed for law firms and corporate legal departments. Announced on August 25, 2026, the new offering is aimed at helping lawyers automate routine and complex legal work while allowing them to spend more time on higher-value tasks. Google says the system provides a governed environment in which organisations can use AI while maintaining control over their data and legal workflows.

The platform is designed to handle a range of legal activities, including contract review and negotiation, monitoring regulatory changes, data discovery, responding to Data Subject Access Requests (DSARs), preparing documents for motions to seal, and drafting non-disclosure agreements (NDAs). Instead of simply generating suggestions, Google's AI agents can carry out parts of these workflows.

This represents a shift from traditional generative AI, where a lawyer asks a chatbot a question and receives an answer, toward agentic AI, where AI systems can perform multiple steps of a task with less manual intervention.

The development is significant because the legal profession handles highly confidential information and requires accuracy, security and regulatory compliance. AI can potentially reduce the time lawyers spend on repetitive work, but legal professionals still need to review AI-generated or AI-performed work carefully because mistakes can have serious consequences. Google's expansion therefore represents an important step in the development of AI-powered professional services, where AI is moving from being a simple assistant to becoming a system capable of participating directly in business workflows.

Source : ***The Indian Express***

ANTHROPIC WANTS AI AGENTS TO CONTROL PHYSICAL MACHINES WITH NEW HARDWARE STANDARD

Anthropic, the company behind Claude, has introduced the Model Hardware Standard (MHS), a framework designed to help AI agents safely interact with physical machines. The early research preview is being tested with selected research labs and manufacturers, with applications including microscopes, robotic arms, lasers, and quantum-computing equipment.

MHS aims to solve the problem of building separate software and hardware integrations for every machine, which can take weeks or months. Instead, it provides a common interface that allows AI agents to understand a machine's capabilities, controls, and safety limits. Similar to Anthropic's Model Context Protocol (MCP), MHS is designed as a common language, but for physical equipment.

Early tests have included AI-assisted drug discovery at Genentech, imaging experiments at the Howard Hughes Medical Institute's Janelia Research Campus, and laser stabilisation on QuEra quantum computers. These examples show how AI agents could eventually run experiments, handle errors, and adjust equipment with less human intervention.

Safety remains a major concern because AI-controlled machines can cause physical damage or injury. Anthropic is therefore conducting additional safety evaluations before wider release. MHS is also intended to be model-agnostic, allowing AI systems beyond Claude to use it.

The project represents a step toward physical AI, where intelligent software can directly interact with the real world. However, MHS remains in an early research stage, with Anthropic planning further testing and safety improvements before broader availability.

Source : ***The Indian Express***

GTA VI EXTENDED LOOK RELEASES ON NETFLIX, YOUTUBE

Rockstar Games has released Grand Theft Auto VI: An Extended Look, its most detailed official preview of the highly anticipated game so far. The presentation premiered on Netflix on August 27, 2026, through a partnership with Rockstar, before becoming available for free on Rockstar's YouTube channel and GTA VI website six hours later. The footage was captured entirely from the PlayStation 5 version.

Running for around 26 minutes, the preview focuses on protagonists Jason Duval and Lucia Caminos, whose criminal activities lead them into a wider conspiracy across the fictional state of Leonida and its returning city, Vice City. It also showcases gameplay including character switching, police chases, shootouts, driving, scuba diving, sports, and other activities.

Rockstar highlights a detailed and interactive open world, improved animations, and the evolving relationship between Jason and Lucia. The Netflix premiere marked an unusual move for Rockstar, giving subscribers early access before the wider public.

Rockstar highlights a detailed and interactive open world, improved animations, and the evolving relationship between Jason and Lucia. The Netflix premiere marked an unusual move for Rockstar, giving subscribers early access before the wider public.

Source : ***Rockstar Games***

OPENAI AND 100 OTHERS WARN THAT WINDOW TO DEFEND AGAINST AI ATTACKS IS NARROWING

More than 100 technology companies and organisations, including OpenAI, Anthropic, Google, Microsoft, Amazon Web Services, CrowdStrike, Visa and Mastercard, have issued a joint warning about the growing threat of AI-powered cyberattacks. In an open letter published on August 27, 2026, they warned that AI could make cyberattacks faster, cheaper and more sophisticated, particularly against critical infrastructure such as hospitals, financial systems and water-treatment facilities.

AI can already automate tasks such as analysing systems, finding vulnerabilities, generating malicious code and coordinating attacks. This could allow attackers to launch more attacks with less technical expertise. The companies are therefore calling for stronger cybersecurity investment, international cooperation, threat-intelligence sharing, and greater use of AI by defenders to identify and fix vulnerabilities.

The warning highlights AI's dual-use nature, as the same technology can help both attackers and cybersecurity teams. OpenAI's "Defender's Window" initiative encourages organisations to adopt AI-based security tools before AI-driven attacks become more widespread.

The companies argue that governments, AI developers, cybersecurity firms and businesses must work together to strengthen defences while there is still time to prepare for increasingly capable AI-driven threats.

Source : ***Tech Runch***

EDITORIAL BOARD